



BLUE WATERS

SUSTAINED PETASCALE IN ACTION:
ENABLING TRANSFORMATIVE RESEARCH

2018 ANNUAL REPORT



BLUE WATERS

SUSTAINED PETASCALE IN ACTION:
ENABLING TRANSFORMATIVE RESEARCH

2018 ANNUAL REPORT

Editors

Barbara Jewett
Catherine Watkins

Creative Director

Steve Duensing

Graphic Designer

Megan Janeski

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 2018 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 (69)
GA	GPU-accelerated: written to run faster on XK nodes than on XE nodes (40)
TN	Thousand node: scales to at least 1,000 nodes for production science (72)
MI	Memory intensive: uses at least 50 percent of available memory on 1,000-node runs (22)
BW	Blue Waters: research only possible on Blue Waters (35)
MP	Multi-physics/multi-scale: job spans multiple length/time scales or physical/chemical processes (59)
ML	Machine learning: employs deep learning or other techniques, includes "big data" (16)
CI	Communication-intensive: requires high-bandwidth/low-latency interconnect for frequent, tightly coupled messaging (40)
IA	Industry applicable: researcher has private sector collaborators or results directly applicable to industry (10)
FS	Frontier Science: this science has produced a first of its kind outcome (21)
BI	Big idea: this research relates to the NSF's "10 Big Ideas" initiative (70)

TABLE OF CONTENTS

3	NCSA DIRECTOR BILL GROPP – BLUE WATERS CHANGED THE WORLD AND ILLINOIS
4	PROJECT DIRECTOR BILL KRAMER – ENABLING FRONTIER SCIENCE AND RESEARCH
6	PAST NCSA DIRECTOR THOM DUNNING – CONGRATULATIONS ON FIVE YEARS OF SERVICE TO THE NATION’S SCIENCE AND RESEARCH COMMUNITY
8	BECAUSE OF BLUE WATERS
10	BLUE WATERS MELTS DATA ANALYSIS WITH HIGH-PERFORMANCE COMPUTING
14	BLUE WATERS SYMPOSIUM KEYNOTES HIGHLIGHT SIGNIFICANT SCIENCE ADVANCES
18	EXTENDED ABSTRACTS
18	<i>space science</i>
70	<i>geoscience</i>
102	<i>physics & engineering</i>
178	<i>computer science & engineering</i>
200	<i>biology, chemistry & health</i>
262	<i>social science, economics, & humanities</i>
268	<i>graduate fellows</i>
318	BLUE WATERS FULL SERVICE ASSESSMENT SUMMARY
322	SCIENCE AND ENGINEERING ADVISORY TEAM COMMITTEE
323	OTHER BLUE WATERS PROJECTS
324	REFERENCES
348	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.

BLUE WATERS CHANGED THE WORLD AND ILLINOIS



For the past five years, the Blue Waters Project at the National Center for Supercomputing Applications has enabled remarkable work in biology, chemistry, physics, geosciences, cosmology and astrophysics, atmospheric science, and many other fields such as agriculture, economics, political science, and the social sciences. You'll find these remarkable discoveries in the pages that follow.

Blue Waters has always been about much more than aiding scientific discovery. Blue Waters was the first petascale system for academic use, continuing NCSA's long history of doing things that have not been done before. And Blue Waters is not just a peak petascale system—it is a true sustained-petascale system, specifically designed to be well-balanced and deliver petascale computing power across a variety of applications.

Those pioneering efforts helped us attract the best talent to not only NCSA, but also the University of Illinois. The Project helped recruit and retain 20 professors at the UI by designating them as "Blue Waters Professors," an honor that comes with substantial computing and data resources intended to boost their research efforts.

This amazing marvel of technology and engineering is one of our area's biggest attractions, drawing visitors from across the country and around the world. In the last five years, more than 12,000 people—ranging in age from 2 weeks to 95—have toured NCSA's National Petascale Computing Facility, which includes a visit to NCSA's machine room and an opportunity to be up close to Blue Waters. Our outreach staff has standing annual reservations from K-12 schools, colleges, and universities to bring students as part of their curriculum, often traveling hundreds of miles to do so.

An often overlooked aspect of the Blue Waters Project is the economic impact the Project has had on the state of Illinois. A 2017 study conducted by Dr. Sandy Dall'erba, director of the Regional Economics Applications Laboratory and associate professor of

Agricultural and Consumer Economics (ACES) at the University of Illinois at Urbana-Champaign, found that from October 2007 to June 2019 the Project has a projected \$1.08 billion direct economic impact on Illinois' economy. Furthermore, Blue Waters is the catalyst for many of NCSA's engagements with Industry and other research organizations even if they do not use Blue Waters.

NCSA was founded in 1986 to enable discoveries not possible anywhere else. The Blue Waters Project has admirably continued and expanded that tradition.

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

ENABLING FRONTIER SCIENCE AND RESEARCH




Every year when I sit to write this letter and reflect on the doors we have opened to discovery, I feel the same wonder and delight that I did when we published our first annual report five years ago. As you read the reports in this book, I hope you share in this amazement at the incredible scientific discoveries that are possible as a result of the Blue Waters Project.

As part of its mission, Blue Waters enables computational and data analysis investigations that could not be done otherwise. In fact, you can see many reports carry the badge indicating “Only on Blue Waters” to signify that the research would not have been possible on other currently-deployed U.S. open-science resources. When I am asked why leadership computing systems like Blue Waters are required, I often refer to the cycle of discovery represented in the included diagram. Whether called Frontier Science or Grand Challenges, these “breakthroughs” are first-of-a-kind advances. These make the infeasible feasible; they raise the bar in a given field through higher fidelity, longer simulated time periods, better quality and accuracy, or lower cost. Frontier problems almost always require improvements in methods and algorithms, as well as increased computational and analysis power. To achieve their one or more orders of magnitude of more effective insight they increasingly also require the convergence of modeling and simulation with the best data analysis and machine learning.

Time and again, we have seen the solution of one “frontier problem” evolve. First the solution enables some teams to start solving more best-of-breed problems (BoB) at that scale once the first one is solved, and eventually these computing and data applications become more ubiquitous and the methods wide spread—what once was on the frontier becomes common practice. The role of Blue Waters is to accelerate advancing the frontier and to rapidly expand the adoption of these methods into other

best of breed and “out front” uses. This process is how the world moved from the first 2-D atmospheric circulation model to being able to accurately predict weather 10-15 days in advance, from studying basic fluid dynamics to using computing for complete aircraft and ship design, and going from studying small molecules to understanding the basic principles of life.



I am amazed that over just five years of service, Blue Waters and the science teams using it have repeatedly pushed the frontier into common practice process multiple times, in multiple disciplines. Examples include the first all-atom simulation of the HIV capsid (and now all-atom simulations of the flu and other diseases), to the first ever simulations of a billion years of evolution of the universe and galaxy formation, to the ability to improve digital elevation models by over six orders of magnitude. This particular example went from a proof of concept to a project to repeatedly map the entire earth in less than three years.

You may notice our fifth annual report contains 12 percent more high-impact summaries than the 2017 report, which itself represented a 225 percent increase in accomplishments over our first report in 2014. The fact that in five 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, bio-medicine, 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.

In the first 2014 report, the projects were all traditional large-scale, physics-based modeling and simulation efforts. In this report, 53% of the projects identify as “data intensive” processing and 11% as doing machine learning/artificial intelligence. This shift is another indication of the leadership and impact Blue Waters is making. Many projects are not only data intensive or machine learning, but instead are combining these methods with multi-scale and multi-physics simulation to achieve cutting edge results. Blue Waters is truly a catalyst for convergence, not just through supporting teams using different methods, but also helping science teams integrate different methods within their work to achieve remarkable new results.

The numbers are impressive, but the Blue Waters Project has always been about much more than just providing physical computing, memory capacity, and storage resources. While high performance resources are necessary for success, they are not sufficient for effectiveness.

Through the Project’s Petascale Application Improvement Discovery (PAID) program that concluded last year, the Project provided millions of dollars to science teams and computational and data experts to improve the performance of applications in a measurable manner. These efforts resulted in scientists advancing their research even further in the current reporting year.

Likewise, the Blue Waters Project continues to help develop our nation’s HPC workforce. The fourth “Best Practices for HPC Training” workshop at SC17 (the International Conference for High Performance Computing, Networking, Storage, and Analysis) with coordination led by Blue Waters staff, had over 100 attendees.

Making high performance computing available to more of our scientific community is important. To broaden participation in the utilization of the Blue Waters resources and services, a new category of Broadening Participation allocations was announced this past year, a component of the Innovation and Exploration allocation managed by NCSA. Through this, over 3.7 million node-hours of computing resources were awarded to 21 research

teams. Included among the principal investigators (PIs) are 10 females and two underrepresented minorities. In addition, there are four female and eight underrepresented minority colleagues listed as co-PIs. Among the lead institutions, five are Minority Serving Institutions and nine are within EPSCoR jurisdictions.

Our Blue Waters intern program for undergraduates has trained over 100 of the next generation HPC workforce. In addition, our highly-competitive Blue Waters Fellows program has helped 46 young researchers gain experience on advanced research methods to take full advantage of leadership-class computing.

At the suggestion of our NSF review panel, we’ve shared some “lessons learned” over the course of our five years. By sharing the more than 35 best practices we developed, our staff is distributing the expertise and experiences we have acquired from our accumulated knowledge. More importantly, our Blue Waters science teams have published more than 1,000 articles and papers in their respective domains and the project staff was also active delivering presentations and publishing papers on best practices developed through our operation of Blue Waters.

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

CONGRATULATIONS ON FIVE YEARS OF SERVICE TO THE NATION'S SCIENCE AND RESEARCH COMMUNITY



It is difficult to believe that it was ten years ago that NSF made the award for the Blue Waters project to the National Center for Supercomputing Applications at the University of Illinois at Urbana-Champaign. I was the NCSA Director at that time and the original PI for the Blue Waters project. It was a time of unprecedented opportunity as the promise of petascale computing was

almost palatable. However, it was also a time of unprecedented challenges. Deploying, operating and using a petascale computer to advance science and engineering was no simple task. But now, five years later, it is rewarding to see how the Blue Waters project has advanced our understanding of the world in which we live.

We recognized early on that new science and engineering codes were one of the keys to the success of the Blue Waters project. As Blue Waters was being developed by our technology partner, we worked intensely with 18 computational science and engineering teams from across a range of disciplines that were granted future access to Blue Waters to get their codes ready for the orders of magnitude increase in computational power provided by Blue Waters. We worked with these teams to rewrite and optimize their codes based on an in-depth knowledge of Blue Waters' hardware. We also worked with our technology partners to develop and enhance the hardware and software that would enable "sustained petascale computing" across a wide range of open science and engineering. This process later became known as "co-design." And, we had to adjust the project due to many unexpected changes, especially the loss of our first technology partner.

Reviewing this fifth annual report on the accomplishments enabled by Blue Waters, what is most remarkable to me is the breadth and depth of new understanding that the project has enabled, from cosmology to biomedicine to political science to severe weather to my own field of chemistry. Blue Waters has advanced almost every field of investigation that it has touched, enabling unprecedented new insights into the research topics being explored. For just a small set of examples, I refer you to the retrospective summaries in astronomy/cosmology, molecular dynamics, and weather/climate by Brian O'Shea, Rommie Amaro, and Susan Bates later in this report (you can see their video presentations at <https://bluewaters.ncsa.illinois.edu/symposium-2018-invited-speakers>). Today, there is research being done using Blue Waters that we never expected when we deployed the system. These include, to name just a few, digital elevation modeling of the entire Arctic and Antarctic, developing rational alternatives to gerrymandering, gaining new insights into economics, and predicting crop yields.

Another integral aspect of the Blue Waters project was, and continues to be, education of a new generation of researchers who know how to use high performance computing and data analysis to push the boundaries of their disciplines. The Blue Waters internship and fellowship programs have become models of how to introduce the next generation of researchers to these new technologies. It is wonderful to see these young scientists and engineers using Blue Waters under the mentorship of some of the most forward-looking senior researchers in the world. The thousands of people trained on Blue Waters, and the tens of thousands who have visited Blue Waters and the National Petascale Computing Facility and learned about the importance of research and computing are added bonuses.

Blue Waters has set landmarks in leadership computing. It is an exquisitely balanced system that

had more memory and memory bandwidth than any open system, the fastest and largest storage system in existence, and more general computing power than any other system in the world at the time. It is also a system for which the software has been optimized to take full advantage of the extraordinary capabilities provided by the hardware. There have been few investments in computing that have had such a rich return on investment.

Blue Waters established a momentum and enthusiasm across the nation's academic research community that, just like the universe, is measurably expanding and accelerating. It is critical that the nation build on this momentum and enthusiasm with the development and deployment of future generations of leading edge computing systems along with the investments in the expertise needed to take full advantage of these new technologies. It is investments such as these that will keep our nation competitive and secure.

In closing, I want to thank the National Science Foundation, the State of Illinois, the University of Illinois at Urbana-Champaign, and the national research community for putting its faith in NCSA and the Blue Waters Project. I would also like to thank the NCSA staff—it was their expertise and dedication that made this project successful. This project was a great and, ultimately, rewarding endeavor. I cannot wait to see what comes next.

A handwritten signature in cursive script that reads "Thom Dunning, Jr.".

Dr. Thom Dunning
Past Director of NCSA
Past Blue Waters Principal Investigator
Professor Emeritus, Department of
Chemistry
University of Illinois at Urbana-Champaign

BECAUSE OF BLUE WATERS

When NSF proposed an open research petascale supercomputer in 2007, they were interested in focusing on a few major areas of science. NCSA took NSF's vision and expanded it into a well-balanced, leadership-class system that supports research and engineering across all disciplines, from atmospheric science, biology, chemistry, physics, geosciences, astronomy, cosmology and astrophysics, to many other fields such as agriculture, economics, political science, and the social sciences.

Blue Waters was the first petascale system for completely open research use, continuing NCSA's long history of doing things that have not been done before. And Blue Waters is not just a petascale system—it is a true sustained-petascale system, specifically designed to be well-balanced and deliver unprecedented insights.

Because of our approach in creating Blue Waters, there have been thousands of research advances that would not have been possible on other currently-deployed U.S. open-science resources. The reports in this book summarize the latest research, and you can read past discoveries online at <https://bluwaters.ncsa.illinois.edu/annual-report>.

These scientific discoveries were made possible not just because of the hardware available, but because of Blue Waters' approach to supporting research teams who use Blue Waters. NCSA already had a reputation of providing high-quality "user" support to scientists running on our machines. The uniqueness of Blue Waters required we take that one step further to ensure our partner's success: in addition to 24/7/365 user support, each major project (e.g. PRAC—Petascale Resource Allocation), Blue Waters Professor, and graduate fellow was assigned an expert support staff who focus on what the research team needs to accomplish with Blue Waters to achieve their scientific outcomes, creates a joint work plan with the research teams and assists the team in making that happen.

Because of Blue Waters' 1.66 petabytes of memory, more main memory than any other system in the world, researchers such as Tiziana Di Matteo and Brian O'Shea were able to accurately simulate the first billion years of the evolution of the universe.

And because of Blue Waters' 36 petabytes of raw on-line storage, capacity for more than 250 petabytes of nearline storage, and an I/O bandwidth of more than 1.1 terabyte per second, researchers were able to make incredible advances in their fields. Researchers like P.K. Yeung, who uncovered new basic principles for turbulence, Susan Bates and Don Weubbles who advanced climate change research, and Leigh Orf and Bob Wilhelmson who advanced our knowledge of severe weather, particularly when it comes to tornado formation.

As astounding as the science discoveries have been, the response by the scientific and education communities, and the public, to this pioneering research instrument has been equally amazing.

Because of the annual Blue Waters Symposium, where scientists using the system and services gather for a week focused on the science enabled by Blue Waters, researchers came together and science communities were strengthened. And because of Blue

Waters, collaborations were established that would, quite frankly, probably not have happened otherwise. Many such collaborations happened because of Blue Waters' support for the NAMD software, a parallel molecular dynamics code designed for high-performance simulation of large biomolecular systems.

This year, in part because of their time using Blue Waters, 12 Blue Waters PRAC PIs are now also receiving Innovative and Novel Computational Impact on Theory and Experiment (INCITE) awards, giving them access to the nation's new fleet of leadership class computing resources run by the Department of Energy.

Because of Blue Waters, NCSA is one of the central Illinois area's most popular technology attractions, drawing visitors from across the country and around the world. In the last five years, more than 12,000 people—ranging in age from newborn to 95—have toured the University of Illinois' National Petascale Computing Facility to be up close to Blue Waters as well as viewing the power and cooling infrastructure that make it run. Our outreach staff has standing annual reservations from K-12 schools, community colleges, and universities to bring students as part of their curriculum, often traveling hundreds of miles to do so.

Because of Blue Waters, a new generation of students are developing the expertise required to make effective use of advanced computing and data analysis, helping to create an advanced workforce to help make America's research and industry competitive. NCSA and its partners, including the Southeast Universities Research Association, the Great Lakes Consortium for Petascale Computation and the Shodor Education Foundation, integrated petascale computing into pre-college and college education and developed a new virtual school of computational science and engineering that pioneered distance learning and has an impact on a national scale. Seven courses have been offered through the virtual school; there were 633 participants from 27 institutions, of which 18 are EPSCoR institutions and eight are Minority Serving Institutions.

Through our Blue Waters intern program for undergraduates we have trained over 100 of the next generation of the HPC workforce, in many cases interesting undergraduates in careers in HPC. And our highly-competitive Blue Waters Fellowship program has expanded and accelerated the research programs of 46 young researchers to take full advantage of leadership-class computing as they begin their careers as the next generation of scientists. Many of those graduate fellows credit Blue Waters with, as one put it, "changing my life." She went on to say, "Because of my fellowship, my research project changed directions. This literally changed my career path. I'll have a better career and a better life now because of Blue Waters."

BLUE WATERS MELTS DATA ANALYSIS WITH HIGH-PERFORMANCE COMPUTING

During its first five years of full service, the Blue Waters system has steadily expanded its reach from traditional modeling and simulation-based high-performance computing (HPC) to converge with a variety of data-processing and data-analysis workloads. These expanding data-oriented and machine learning workloads are being driven by two forces. First, the quality and quantity of sensors available to observational scientists continues to grow exponentially, fueled by technology improvements, as well as by economies of scale due to commercial applications such as small imaging satellites. Growth in scientific “big data” has in turn driven the development of efficient and automated methods of extracting quantitative and actionable information. The most powerful of these methods, such as deep neural networks, are highly demanding of both computation and storage, requiring the capabilities of leadership-class machines such as Blue Waters.

The second force driving data-oriented computing in HPC is the opportunity to extract additional insight from the data generated by the high-performance simulations that remain the bulk of the Blue Waters workload. With previous generations

of supercomputers, simulation output would be downloaded by science teams to their local workstations for visualization and analysis. The output produced by simulations on Blue Waters, however, may be too large for the science team to store locally, and the time needed to transfer the output could delay decisions on the direction of a computational campaign. The data can then be analyzed *in situ* on Blue Waters, either read back from high-performance storage or processed directly in memory during the running simulation. Having the massive computing power of Blue Waters available also increases the variety and complexity of analyses that can be applied, including communication-intensive graph analytics and floating-point-intensive deep neural networks (aka Machine Learning and AI).

The convergence of HPC and data analysis is, therefore, a two-way street, with observational and computational methodologies both enhancing and enabling each other. In fact, in this report, 53% of the projects identify as data intensive and 11% as using machine learning/AI, a larger increase from past years.

A DATA SCIENCE ENVIRONMENT

We can now explore the demands that this convergence places on HPC systems and how the hardware and software of Blue Waters meet these demands. A traditional supercomputer had three primary tasks: run simulations in parallel on large numbers of processors, write the resulting output to a few large files on a high-performance parallel filesystem, and transfer those files over the internet to researchers’ machines for visualization and analysis. Blue Waters, however, was envisioned as a complete problem-solving environment for high-end computational science that would enable researchers to remotely analyze simulation output, and to this end it has larger-memory nodes, software for GPU-accelerated remote visualization, and an increasing variety of tools from the broad data analytics community.

A traditional scientific data processing system such as a compute “farm” associated with a particle accelerator or other major instrument can be contrasted with a supercomputer in that it must reliably ingest large flows of data, transform and

reduce the incoming data through a set of independent or trivially parallel computations, and store only the valuable events or other smaller results of the analysis. A modern data science environment, however, provides storage for both ingested and transformed data that allows it to be repeatedly queried and analyzed as new data are acquired and as new insights and hypotheses arise. Data scientists are accustomed to working interactively with data sets, allowing numerous iterations on the design of data models and analysis methods than afforded by traditional batch processing systems.

The original “big data” architecture focused on deploying storage that was inexpensive, distributed, and redundant. This architecture was able to “scale out” by adding more server nodes, and the stored data could be queried efficiently by the “MapReduce” algorithm, in which data were first selected, filtered, and transformed (i.e., “mapped”) at the level of individual nodes and then the much smaller results aggregated (“reduced”) globally via internode communication. The dominant implementation of this model is the Apache Hadoop software, which is today synonymous with big data.

STORAGE AND FILE SYSTEMS

A primary enabling technology of Hadoop is the Hadoop Distributed File System (HDFS), which runs across the local disks of a Hadoop cluster. While highly scalable, HDFS lacks many of the capabilities of a true parallel filesystem, such as the Lustre filesystem used on Blue Waters. Most significantly, Lustre is deployed on dedicated storage servers that are connected to a much larger number of compute-node clients by a high-bandwidth, low-latency network, allowing any file to be accessed by any client with equal performance across the entire Blue Waters System. Lustre data are automatically distributed across servers at the level of complete files, or large individual files may be broken into stripes that are spread across servers for increased single-file bandwidth. Resilience to hardware failure is provided not by redundant copies of the data but by mathematical algorithms that allow data to be reconstructed after the loss of one or more disks with less than a 30% increase in disk usage. Finally, a Lustre filesystem is accessed by the same commands and function calls as local disk, requiring no special adaptation of existing software.

An effort in 2015 to deploy Hadoop on Blue Waters identified several challenges to using Hadoop on HPC platforms, including the need to use the shared parallel filesystem to emulate local storage and the need to interoperate with the batch scheduler. The initial results were considered encouraging but not officially supported. The results were presented at the 2015 International Conference on Advances in Big Data Analytics. Since this time, the Hadoop Adapter for Lustre (HAL) software has become available, providing an alternative to emulating local disks.

The bottleneck of Hadoop's distributed filesystem has also been relaxed by the development of Apache Spark, which reduces storage access by keeping as much in-flight data in memory as possible. As memory is much faster than disk, adding additional nodes to a memory-starved Spark cluster can provide a super-linear performance increase. Blue Waters has relatively large memory per-node for an HPC system, and with its large node count a user could run Spark with over a petabyte of memory. The transient nature of a Spark cluster deployed on Blue Waters further allows resources to be right-sized for the analysis at hand and returned to the general pool for other work when idle.

PROGRAMMING

The Python programming language is ubiquitous in data science, a testament to both its ease of use and its massive library of extensions. Traditional HPC programming languages such as Fortran, C and C++ are compiled from source code into a single binary executable file, which can then be efficiently launched in parallel across thousands of nodes. Python, however, is either interpreted or compiled automatically on every core where it is run, triggering access to a large number of library files. When Python is run on the Blue Waters parallel filesystem, the repeated reading of very small files leads to slow program startup even on a single node, becoming unusable at scale. This problem is addressed



on Blue Waters by a custom “bwpy” Python stack installed as a virtual read-only filesystem. Access to the many individual Python library files is managed separately by each node, while the Lustre filesystem sees only a single large file. Loading the bwpy module automatically makes the virtual filesystem available.

A highly popular tool for developing, sharing, and executing data-science analyses in Python is the interactive Jupyter notebook system, which integrates code, output, and commentary into a single editable document that is presented to the user via a web browser. Many data scientists today consider Jupyter to be a standard tool of the field that is expected to be available as needed. Blue Waters enables the Jupyter notebook server to run on a login node for light usage, on a dedicated compute node for heavier usage, or in parallel across multiple nodes for the most demanding analyses. In all cases a secure network “tunnel” is required to view the Jupyter notebook on the user's desktop or laptop.

While tools such as Jupyter notebooks enable interactive access to Blue Waters, only the few shared login nodes are available on demand. Access to the thousands of compute nodes is only allowed via a batch-resource scheduler, with large jobs possibly waiting for hours or days to start depending on job priority and competing workload in the queue. These delays, combined with unpredictable start times, are incompatible with extensive interactive use, so most nondevelopment work on Blue Waters is done in batch mode. Users can reserve a required node type and count at a specific date and time as needed to enable interactive runs at scale for development, debugging, and exploration.

GPUS ON BLUE WATERS

GPU acceleration has been a critical enabling technology for the growth of deep neural network machine learning applications in fields ranging from self-driving cars to gravitational wave detection. The decision to include over 4,000 GPUs in Blue Waters was motivated by the goal of accelerating traditional simulations but now makes the machine highly suited for designing, training, and deploying deep neural networks. The high degree of parallelism available on Blue Waters can be used to both train a single network faster and to tune the many “hyperparameters” of the underlying network architecture. The popular TensorFlow package is available for use on Blue Waters, and the continued growth of the machine learning field is inspiring the development of new tools designed to exploit the network, storage, and scaling capabilities of HPC systems.

DATA ANALYSIS

Modern data analysis software is assembled from multiple components of external packages. The software is debugged and validated for correctness using specific versions of these packages, and using a mismatched version can introduce incompatibilities, failures, or bugs leading to incorrect results. Container technology, in particular Docker, is the preferred solution to this versioning problem. A Docker image encapsulates and packages the complete software environment of an application, allowing the

program to run identically on any machine, even alongside other “containerized” applications with different environments. The original Docker software required administrator privileges to execute applications, which is an unacceptable security risk on any shared system. Blue Waters instead deploys the Shifter and Singularity packages, either of which allows users to safely execute Docker images on the compute nodes.

The final requirement for high-performance data analysis is to automate the management of the many data transfers and program executions of a modern data science workflow. Blue Waters supports multiple tools for scientific workflow management, including the Swift/T system for dynamically scheduling large numbers of serial tasks. Large data transfers to and from the Blue Waters filesystem and near-line tape archive are managed by the Globus Online system, which schedules and monitors transfers and can be accessed via a web browser, the shell command line, or Python. Users may also integrate Blue Waters into external workflows managed by a central server with assistance from the application support team.

CONCLUSION

In conclusion, modern data science increasingly requires HPC resources, and traditional fields of HPC are increasingly benefitting from data science techniques as both numerical simulation and data analysis continue to complement the classical domains of theory and experiment. Although the initial target applications for Blue Waters were all modeling and simulation-based, the machine's forward-looking balance of compute, interconnect, storage, and external networking have enabled it to support an increasing variety of data analysis workloads, making Blue Waters a true “go anywhere, do anything” universal instrument for computational science.

BLUE WATERS SYMPOSIUM KEYNOTES HIGHLIGHT SIGNIFICANT SCIENCE ADVANCES



Brian O'Shea



Susan Bates



Rommie Amaro

The Blue Waters Project can take credit for enabling thousands of scientific discoveries. However, there are domains where, taken together, these discoveries significantly accelerated advancing science across the field. Leaders in these fields gave keynote talks at the 2018 Blue Waters Symposium to share some of these extraordinary advances.




Image from Brian O'Shea.

ASTRONOMY, ASTROPHYSICS, AND SPACE SCIENCES

Since researchers in the astronomy, astrophysics, and space sciences area use about 25 percent of the time on Blue Waters, there have been hundreds of notable achievements in these domains over the past five years. Michigan State astrophysicist Brian O'Shea summarized the integrated impact of what Blue Waters has brought to this area as "the ability to go to three dimensions from 2D, and to do larger physical volumes like global simulations, and combining different simulations together to understand these phenomena over a large scale." O'Shea noted astronomical and space science relies on large scale computing to make progress, and Blue Waters has been a unique and indispensable resource for these communities over the past five years.

"Astronomy, astrophysics, and space sciences are fields that are inundated with data, and involve physical phenomena that are complex and cannot be explored via experiments on Earth. As a result, sophisticated simulation tools are required to make theoretical progress in understanding these phenomena, and large-scale computing is required to both analyze observational data and to execute the necessary simulations. Blue Waters has a great interconnect; it's high bandwidth and low latency... and the file system is both very fast and very large.

Both of these things are important because the kinds of simulations that we deal with typically generate vast volumes of data and you need to write that data to something really quickly and then you need to read it back in really quickly to analyze it and visualize it. Having access to Blue Waters has given me the ability to do astrophysics problems that I could not have tackled otherwise; ...having access to powerful and very capable computational resources is crucial to progress in my field."

O'Shea noted that fundamental advances have been made at all scales, "from the Earth's magnetosphere to the edge of the Universe!"

These include:

- modeling the formation of the terrestrial planets, (Matthew Clement, Blue Waters Fellow, University of Oklahoma);
- exoplanet discovery in the Kepler and K2 missions (Ethan Kruse, Blue Waters Fellow, University of Washington);
- greatly improved understanding of plasma turbulence (Vadim Roytershteyn, Space Science Institute);
- global simulations of the interaction of the solar wind with the interstellar medium (Nikolai Porgoloev, University of Alabama in Huntsville);
- the first 3D global studies of solar wind and magnetosphere interaction with MHD+kinetics together (Gabor Toth and Ward Manchester, University of Michigan);
- verifying LIGO's detection of a gravitational wave and LIGO's detection of a neutron star binary merger (Eliu Huerta and Roland Haas, NCSA/University of Illinois);
- showing that the most massive black holes in the universe grow in compact, spheroidal galaxies at high redshift (Tizianna DiMatteo, Carnegie Mellon University);
- understanding what happens when two black holes collide (Manuela Campanelli, Rochester Institute of Technology and Scott Noble, NASA's Goddard Space Flight Center);
- showing that super massive black holes (SMBH) are unlikely to be SMBH progenitors (Brian O'Shea, Michigan State University, John Wise, Georgia Tech, and Michael Norman, University of California, San Diego);
- identifying plausible mechanisms for the "direct collapse" scenario and the first prediction of time scales for close SMBH pairs following galaxy mergers (Tom Quinn, University of Washington); and
- "Zoom" galaxy simulations showing growth of SMBHs by short bursts of accretion (Philip Hopkins, California Institute of Technology and Claude-Andre Faucher-Giguere, Northwestern University).



Image from Leigh Orf.

ATMOSPHERIC SCIENCE, CLIMATE, AND WEATHER

While delivering her keynote talk, Susan Bates of the National Center for Atmospheric Research (NCAR), pointed out that “what we get from Blue Waters isn’t just getting our science done or jobs run on the machine, but it’s really made connections and collaborations that I think are going to last far beyond the machine.” Bates noted that one of the biggest contributions Blue Waters has made to this area of research is the capability to run “incredibly complex models” at high resolution and the “amazing animations” that result. She noted that Blue Waters provides a massively parallel system, one of the largest storage and bandwidth computing facilities and excellent sharing services, and the Blue Waters staff understand the needs of the projects in these domains and facilitate getting the science done. Without Blue Waters, she estimates that she and other scientists in these domains would have accomplished only about 10 percent of what they did in these past 5 years.

“Having a resource like Blue Waters in order to get a full set of simulations, many different scenarios, under different climate regimes and having those ensembles is really important. The simulations we’re running on Blue Waters are using a particular version of our model, the Community Earth System Model that is maintained by NCAR, and they are [run with] a version of the model that has not been previously well documented, and certainly not been

run at the high resolutions and with the number of simulations that we’ve been able to achieve with Blue Waters. Because of Blue Waters we now have a very large set of high resolution climate simulations with the Community Earth System Model that we’ll be able to use and mine that data, for years to come.”

Bates’ choices for the most significant discoveries enabled by Blue Waters in atmospheric science include:

- how tropical cyclones cycles and events will change in the future (Susan Bates and Nan Rosenbloom, NCAR);
- how tropical cyclones, heat storage in the ocean, and climate are interrelated (Ryan Sriver and Hui Li, University of Illinois);
- future changes in extreme temperature (Donald Wuebbles, Blue Waters Professor, and Zachary Zobel, University of Illinois);
- quantifying entrainment in clouds and storms that enable global scale atmospheric currents (Sonia Lasher-Trapp, Blue Waters Professor, University of Illinois);
- understanding how long track violent tornadoes develop in supercell thunderstorms (Leigh Orf, University of Wisconsin); and
- develop methods to efficiently and effectively generate and deliver data and insights from the NASA Terra satellite (Larry Di Girolamo, Blue Waters Professor, University of Illinois).




Image from Juan Perilla.

BIOPHYSICS

“The real advances to solve some of the most challenging or pressing problems facing humanity are going to come at the intersection of observational and simulation science,” noted Rommie Amaro of the University of California, San Diego in her keynote presentation. Machines like Blue Waters allow researchers to conduct all-atom molecular dynamics simulations at unprecedented size and scale, providing views into the inner workings of cells at the molecular level. She and other researchers use Blue Waters like a computational microscope, a term coined by the late Klaus Schulten, which has led to some exciting discoveries that were only possible because of Blue Waters.

“Blue Waters has been super critical to our research because it’s one of the only places in the world where we can actually get access to the size and scale of computer architecture that we need in order to build and simulate these highly complex scenes of biological systems. Blue Waters is critically important for us because it’s one of the largest machines that we can get access to for the simulation of these systems, and it really is not possible to run simulations of that scale for any appreciable amount of time or at all on sort of the normal or the standard supercomputers. ... Blue Waters took us into and across key ‘capability gaps’; engaging all-atom and coarse-grained molecular

dynamics to give unseen views into the inner workings of cells at the molecular level.”

Amaro’s choices for the most significant biophysics discoveries enabled by Blue Waters include:

- a close look at a fine-grained structure of nucleic acids, including force field development and RNA (Tom Cheatham, University of Utah);
- closely examining the mechanisms by which potassium channels open and close (Benoit Roux, Eduardo Pirozo and Jing Li, all at University of Chicago);
- unveiling the functions of the Hepatitis B virus capsid (Juan Perilla, University of Delaware);
- large-scale, coarse-grained molecular simulations of the viral lifecycle of HIV-1, providing new insight into the HIV infection process (Gregory Voth, University of Chicago);
- molecular dynamic simulation of a fully enveloped flu virus, the largest biological system ever simulated—about 165 million atoms (Rommie Amaro, University of California, San Diego); and
- determining the precise chemical structure of the HIV capsid (Klaus Schulten, University of Illinois).

SPACE SCIENCE

AEROSPACE ENGINEERING

ASTROPHYSICS

COSMOLOGY

HELIOPHYSICS

- 20 *Disks around Merging Binary Black Holes: From GW150914 to Supermassive Black Holes*
- 22 *Multimessenger Astrophysics with the Blue Waters Supercomputer*
- 24 *Realistic Simulations of the Intergalactic Medium: The Need for Enzo-P/Cello*
- 26 *Core-Collapse Supernova Simulations: Simulating the Brightest Objects in the Sky and the Source of Life's Building Blocks*
- 28 *Deep Learning for Multimessenger Astrophysics: Real-Time Discovery at Scale*
- 30 *Tiny Galaxies Host the First Giant Black Holes: BlueTides Simulation Makes Contact with the First 700 Million Years of Cosmic History*
- 32 *Shedding Light on Inspiring Binary Black Holes with Magnetized Mini-Disks*
- 34 *Feeding Black Holes: Tilt with a Twist*
- 36 *3D Stellar Hydrodynamics of H-Ingestion Flashes and Shell Mergers*
- 38 *3D Nature of Collisionless Magnetic Reconnection at Earth's Magnetopause*
- 40 *Kinetic Simulations of Plasma Turbulence*
- 42 *Modeling Plasma Flows with Kinetic Approaches using Hybrid CPU-GPU Computing*
- 44 *Unified Modeling of Galaxy Populations in Clusters*
- 46 *Astrophysics on Graphical Processing Units*
- 48 *Petascale Simulations of Merging Black Holes and Neutron Stars*
- 49 *Chemistry and the Early Universe*
- 50 *Simulating Galaxy Formation across Cosmic Time*
- 52 *Elucidating the Alignment Mechanism for Black Hole Accretion Disks Subjected to Lense-Thirring Torques*
- 54 *Modeling Physical Processes in the Solar Wind and Local Interstellar Medium with a Multiscale Fluid-Kinetic Simulation Suite: Connecting Solar, Heliospheric, and Astrophysical Scales*
- 56 *Simulating Two-Fluid MHD Turbulence and Dynamos in Star-Forming Molecular Clouds*
- 58 *The Evolution of the Halo Mass Function in an Accelerated Universe*
- 60 *Inductive E Fields in Earth's Magnetosphere*
- 62 *Multiscale Space Weather Simulations*
- 64 *Cosmic Reionization on Computers*
- 66 *Using Blue Waters to Understand the Origins of Galaxies and the Nature of Dark Matter*
- 68 *Simulating the Co-Evolution of Galaxies and Their Supermassive Black Holes*

DISKS AROUND MERGING BINARY BLACK HOLES: FROM GW150914 TO SUPERMASSIVE BLACK HOLES

Allocation: Illinois/1,499 Knh

PI: Stuart Shapiro¹

Co-PIs: Milton Ruiz¹, Antonios Tsokaros¹

Collaborators: Abid Khan¹, Vasileios Paschalidis²

¹University of Illinois at Urbana-Champaign

²University of Arizona

EXECUTIVE SUMMARY

The coincident detection of gravitational waves (GWs) with electromagnetic (EM) signals from the coalescence of black hole binaries is a new observational challenge. Combining GW and EM observations offers a unique probe to understanding black hole cosmological evolution and accretion processes. We report results from general relativity simulations of circumbinary magnetized disks accreting onto nonspinning merging black holes. We survey different disk models to quantify the robustness of previous simulations on the initial disk model. Scaling our simulations to LIGO (Laser Interferometer Gravitational-Wave Observatory) GW150914 we find that such systems could explain possible gravitational wave and electromagnetic counterparts, such as the Fermi GBM hard X-ray signal reported 0.4 seconds after GW150915 ended. Scaling our simulations to supermassive binary black holes, we find that the observable flow properties such as accretion rate periodicities and the emergence of jets throughout inspiral, merger, and postmerger we reported in earlier studies display only modest dependence on the initial disk model.

RESEARCH CHALLENGE

Accreting black holes are central in explaining a range of high-energy astrophysical phenomena that we observe in our universe, such as X-ray binaries, active galactic nuclei (AGN), and quasars. Recently, substantial theoretical and observational effort has gone into understanding accretion onto binary black holes, and the emergent electromagnetic (EM) signatures these systems may generate, because they are anticipated to exist at the centers of distant AGNs and quasars [1,2]. The bulk of the research so far has focused on supermassive black hole binaries (SBHBHs), and about 150 candidate accreting SBHBHs have been identified in quasar surveys [3,4]. Depending on the physical properties of the above systems, such as the mass ratio and orbital period, some of these candidates may be the gravitational-wave (GW)-driven regime [5]. However, in addition to accreting SBHBHs, there may exist black hole binaries (BHBHs) of a few tens of solar masses that could be accreting matter from a circumbinary disk. This scenario has attracted a lot of attention recently because of the direct detection of GWs by the LIGO/VIRGO collaboration. As a crucial step toward solidifying the role of BHBHs as multimessenger systems, we report results from general relativity simulations of

circumbinary magnetized disks accreting onto merging BHBS, starting from relaxed disk initial data. We evolve the systems after binary-disk decoupling through inspiral and merger, and analyze the dependence on the binary mass ratio and initial disk model.

METHODS & CODES

Magnetohydrodynamic (MHD) numerical simulations in full general relativity (GR) require the solution of the Einstein 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 [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 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].

RESULTS & IMPACT

We performed MHD simulations of binary black holes with different mass ratios that accrete magnetized matter from a circumbinary accretion disk. We considered three initial disk models that differ in their scale heights, physical extent, and in their magnetic field content, in order to test whether previous properties of MHD accretion flows onto binary black holes are sensitive to the initial disk model [7]. We find that the presence of periodicities in the accretion rate, the emergence of jets, the time delay between merger, and the boost in the jet luminosity that we previously discovered [1,8–9] are all robust features and largely insensitive to the choice of initial disk model (see Fig.1).

As in our previous studies, we ignored the disk self-gravity and adopted a simplified Γ -law equation of state ($\Gamma=4/3$), which

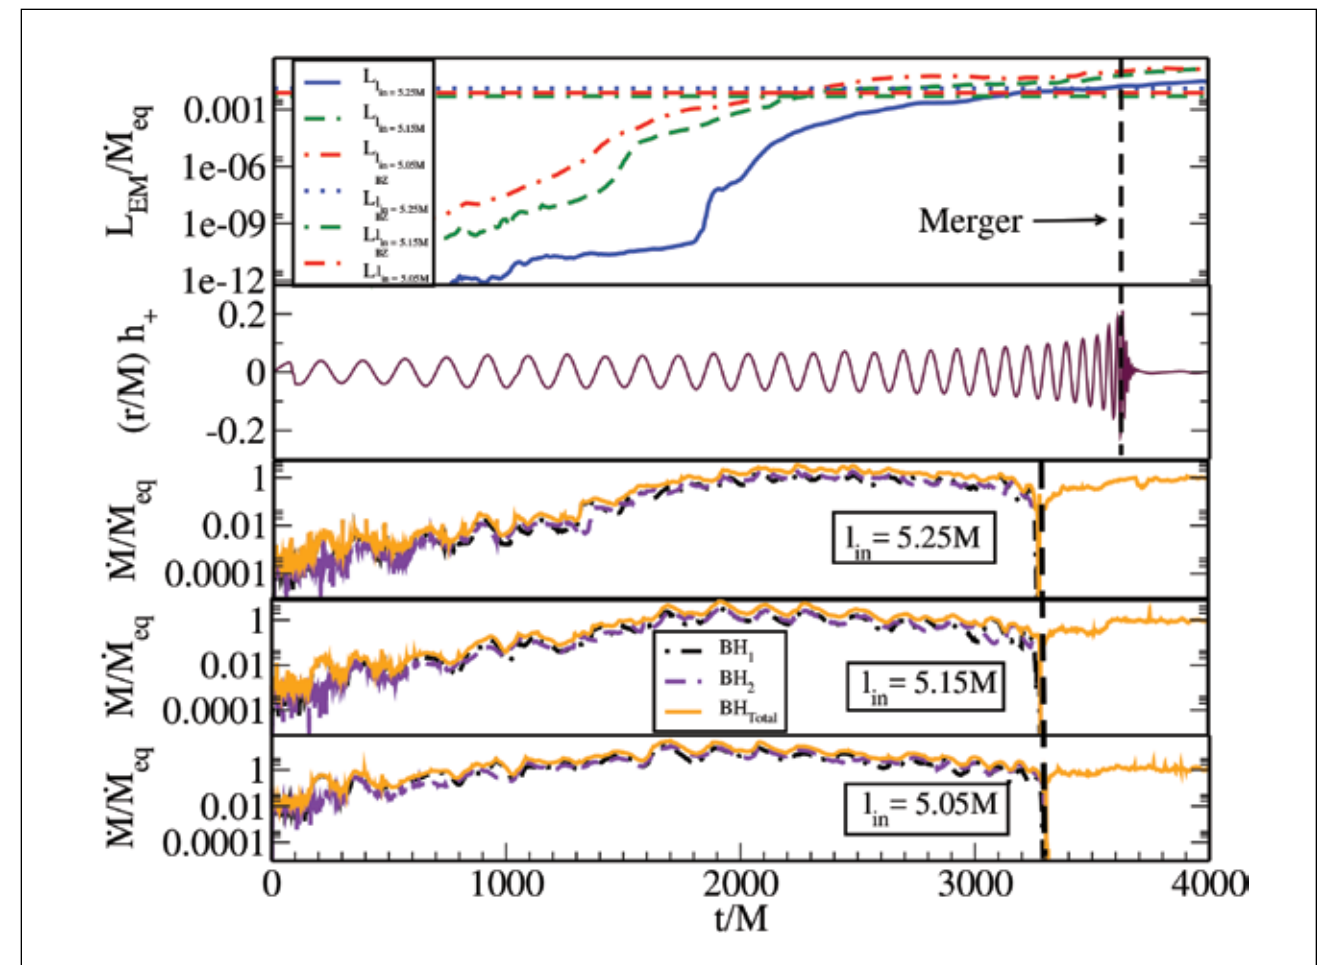


Figure 1: Poynting luminosity, accretion rate, and the strain of the gravitational waveform (“plus” polarization) vs. time for different disk models (see [7]). The displacements in the dashed vertical merger line in the luminosity and GW plots with respect to the merger lines in the accretion plots account for the light travel time.

allows us to scale the binary black hole mass and disk densities to arbitrary values. Thus, our results have implications both for LIGO black hole binaries and for supermassive black hole binaries at centers of luminous AGNs and quasars. Scaling our simulations to LIGO GW150914 we find that magnetized disk accretion onto binary black holes could explain both the GWs detected from this system and the EM counterpart GW150915-GBM reported by the Fermi GBM team 0.4 seconds after LIGO’s GW150915. When scaling to supermassive black hole binaries, we find that at late times flow properties, temperatures, and thermal frequencies are all robust, displaying only modest dependence on the disk model. Nevertheless, the range of disk thickness and ratio of magnetic-to-gas pressure in our survey is limited by what we can achieve with current computational resources and methods. As computational resources grow and numerical techniques advance we will be able to probe wider ranges of these parameters.

WHY BLUE WATERS

Blue Waters provides the required computational power to simulate these cosmic sources in a timely manner. By adding OpenMP support to our MPI-based code, scalability on multicore machines has improved greatly. With the Blue Waters next-generation interconnect and processors, our hybrid OpenMP/MPI code exhibits greater scalability and performance than on any other supercomputer we have used. Recently, we were able to build our code with the Blue Waters Intel compilers. This resulted in a significant boost of our code’s performance by about 30%, making Blue Waters unique for tackling the astrophysical problems we want to address.

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

DI

MULTIMESSENGER ASTROPHYSICS WITH THE BLUE WATERS SUPERCOMPUTER

Allocation: Innovation and Exploration/1,000 Knh

PI: Eliu Huerta¹

Collaborators: Roland Haas², Daniel S. Katz², Timothy Bouvet², Jeremy Enos², William T.C. Kramer², Hon Wai Leong², David Wheeler², Edgar Fajardo³, Stuart Anderson⁴, Peter Couvares⁴, Joshua Willis⁴

BW

MP

¹University of Illinois at Urbana-Champaign

²National Center for Supercomputing Applications

³University of California, San Diego

⁴California Institute of Technology

EXECUTIVE SUMMARY

We developed a novel computational framework that connects Blue Waters, the NSF-supported leadership-class supercomputer operated by NCSA, to the Laser Interferometer Gravitational-wave Observatory (LIGO) Data Grid via Open Science Grid technology. This work represents the first time Open Science Grid, containers, and Blue Waters were unified to tackle a scientific problem at scale. This new framework has been used during LIGO's second discovery campaign to run the most computationally demanding LIGO workflows on Blue Waters to accelerate discovery in the emergent field of gravitational wave astrophysics, and to validate the first gravitational wave detection of two colliding neutron stars with the LIGO and Virgo gravitational wave detectors. This discovery that marks the beginning of multimessenger astrophysics (MMA). Supporting LIGO data analysis workflows concurrently with highly parallel numerical relativity simulations is the most recent success and most complex example of successfully achieving convergence with the Blue Waters supercomputer.

RESEARCH CHALLENGE

Future gravitational wave discovery campaigns will be longer and will involve more detectors, thereby requiring additional compute resources. Furthermore, to support urgent needs such as the detection of gravitational waves that are accompanied by emission of light and neutrinos, LIGO will need to rely on

resources beyond those available for its normal processing to validate these discoveries. Other challenges have to do with NSF and HPC infrastructure. Computationally demanding workflows allow Blue Waters to increase cluster utilization and throughput. The computational tasks required for this work is good utilization of backfill for unused Blue Waters nodes. MMA require the interoperability of NSF cyberinfrastructure resources and show how large projects can benefit from making use of existing resources rather than having to build their own custom solutions for all possible needs.

METHODS & CODES

The Open Science Grid (OSG) provides federated compute resources for data-intensive research in a variety of science areas [2]. OSG targets typical high-throughput workloads consisting of spatially small, loosely coupled science jobs that are executed on any of the participating resources providing clusters [3]. We used this flexibility to target high-throughput computing workloads on Blue Waters.

LIGO currently uses the Pegasus Workflow Management System [4] as a layer on top of DAGMan (Directed Acyclic Graph Manager) to manage dependencies. DAGMan is provided by HTCondor to enforce dependencies among jobs in large workflows, and reliably restart workflows from point of failure.

PyCBC [5] is one of LIGO's most computationally intensive gravitational wave search pipelines and the only production pipeline that currently runs on OSG. It has been thoroughly tested by LIGO on a number of different clusters. However, all of them used a variant of the RedHat or Debian operating systems on their compute nodes, where specific versions of these operating systems were used to certify the software stack. This presented a challenge, as Blue Waters does not operate on these tested Linux variants, but rather on a lightweight Linux variant based on SUSE. To overcome this challenge, we adopted Shifter [6] as Blue Waters' container solution. Shifter accepts Docker image files and converts them into a disk image suitable for concurrent use by multiple Blue Waters compute nodes. Further, Shifter ensures that the




Figure 1: Use of Shifter to run LIGO workflows on Blue Waters.

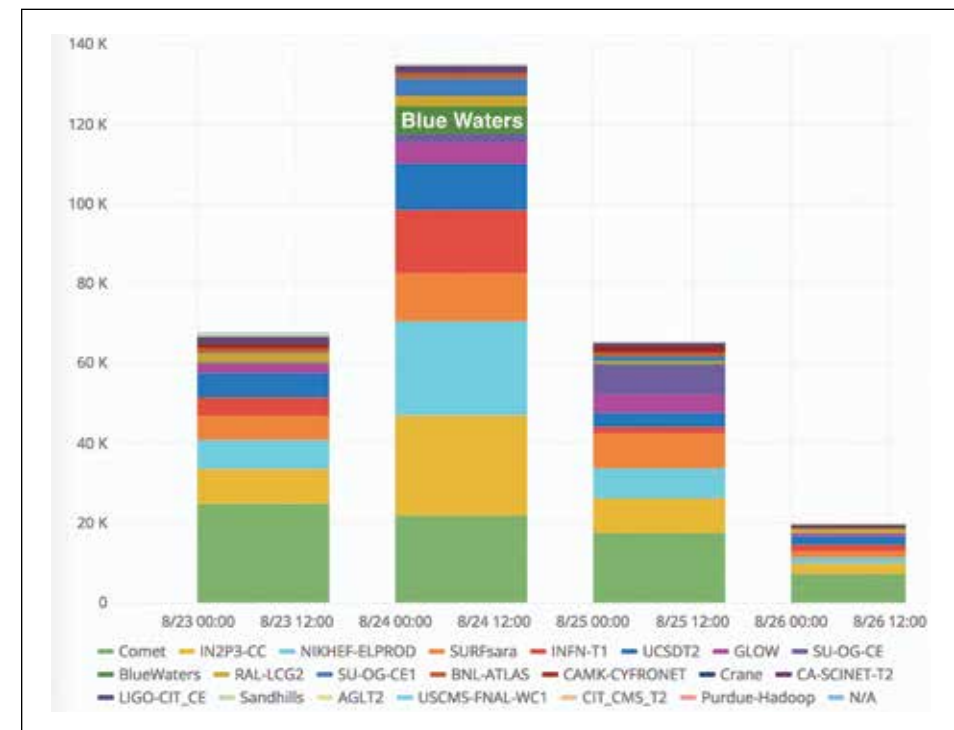


Figure 2: Open Science Grid compute resources used for large-scale gravitational wave data analysis. The chart shows the first time Blue Waters was used at scale as an Open Science Grid compute element, which corresponds to the gravitational wave discovery of two colliding neutron stars by the LIGO and Virgo detectors.

systemwide, parallel file systems are visible inside the container; that MPI can be used; and that Blue Waters' security policy is enforced on the container.

To validate our results, we first ran a small PyCBC workflow on OSG facilities using the data set utilized by GitHub Travis CI tests on LDG clusters. This data set and the results obtained from this analysis have been thoroughly cross-checked using LDG and OSG resources.

Having a baseline for comparison, we ran a PyCBC workflow on Blue Waters using the same validation data set and thoroughly checked that the results reported in both independent analyses were identical. Thereafter, we repeated the same exercise running 10 times larger PyCBC workflows both on OSG and Blue Waters, and confirmed that the results were consistent. Upon confirming that our computational infrastructure works in a stable manner and that we were able to accurately reproduce results obtained with OSG resources, we stress-tested this new framework with several production-scale workflows. The computational framework used for these studies is presented in Fig. 1.

RESULTS & IMPACT

This work marks the first time convergence was reached on the Blue Waters supercomputer and exhibited the flexibility and interoperability of NSF cyberinfrastructure to enable and accelerate scientific discovery. We have used this novel computational framework at scale to validate the gravitational wave detection of two colliding neutron stars with the LIGO and Virgo detectors [7], as shown in Fig. 2.

This framework can be readily used to run other scientific workflows on the Blue Waters supercomputer if they meet the following requirements: they are a good match to the OSG infrastructure, the software can be containerized, and a workflow manager can be used to monitor the workflow from end to end, i.e., Pegasus, Swift, etc. This is a minimal set of requirements that may be easily met by existing OSG users, who may already use portable, self-contained software that could be containerized.

WHY BLUE WATERS

The Blue Waters supercomputer is ideally suited to facilitate large-scale gravitational wave data analysis because the large number of independent jobs in these analyses can quickly be run using the reasonably large set of otherwise unoccupied nodes through backfill. Blue Waters staff assisted us throughout the development and exploitation of this framework at scale for gravitational wave discovery.

PUBLICATIONS & DATA SETS

Huerta, E.A., et al., BOSS-LDG: A Novel Computational Framework that Brings Together Blue Waters, Open Science Grid, Shifter and the LIGO Data Grid to Accelerate Gravitational Wave Discovery. *eScience*, (2017), DOI:10.1109/eScience.2017.47.

FS

BI

REALISTIC SIMULATIONS OF THE INTERGALACTIC MEDIUM: THE NEED FOR ENZO-P/CELLO

Allocation: NSF PRAC/125 Knh
PI: Michael Norman¹
Co-PI: James Bordner¹
Collaborator: David Tytler¹

¹University of California, San Diego

EXECUTIVE SUMMARY

In the continuing quest for more comprehensive models of the intergalactic medium, our research team is developing Enzo-P/Cello, the petascale fork of the widely used Enzo cosmological adaptive mesh refinement (AMR) code. Unlike Enzo, Enzo-P/Cello employs an array-of-octrees AMR mesh and is parallelized using Charm++. Here, we present the first scaling results of Enzo-P/Cello on Blue Waters showing near ideal scaling on a hydrodynamic AMR test problem to 262,000 floating point cores. This problem generates 1.7 trillion cells at four levels of refinement, making it one of the largest AMR tests ever performed.

RESEARCH CHALLENGE

Improved computational models of the intergalactic medium (IGM) are needed to extract information encoded in the high-resolution optical spectra of distant quasars. Information includes the physical state of the mostly primordial gas pervading the universe, but also the dark matter that shapes the gas into discrete intergalactic absorption line systems (the Lyman alpha forest). Standard computational models disagree with certain aspects of the observational data [1], suggesting there is some key ingredient missing in the models. Previously, we explored whether modeling the population of quasars that provide a photoionizing bath




Figure 1: Small section of the Alphabet Soup weak scaling test of the Enzo-P/Cello extreme-scale adaptive mesh refinement code. The image shows interacting blastwaves sourced by an array of high-pressure regions in the shape of letters of the alphabet.

of ultraviolet radiation as discrete point sources rather than a homogeneous background could improve agreement with observations. Using the Enzo hydrodynamic cosmology code [2] enhanced with multigroup flux-limited diffusion radiative transfer, we found this change did not improve the results. We are therefore investigating our next hypothesis: that dense gas bound to galaxies that is unresolved in the Enzo simulations supplies significant absorption of the quasar light and modifies the key observables in such a way to improve agreement with observations.

METHODS & CODES

Including galaxies in simulations of the IGM poses severe resolution requirements that can be addressed using adaptive mesh refinement (AMR) around galaxies. Enzo’s AMR capability is not sufficiently scalable to permit a full-frontal assault on this problem. For this reason, we have developed a successor to the Enzo code called Enzo-P (for petascale) built on an entirely new highly scalable AMR framework called Cello. The combined code—Enzo-P/Cello [4]—uses the Illinois-developed Charm++ parallel object framework for parallelization. We have implemented the already proven scalable Forest-of-Octrees AMR algorithm [3] on top of Charm++ and have obtained excellent parallel scaling results on Blue Waters as a prelude to our target application problem.

RESULTS & IMPACT

We have achieved ideal weak scaling to 262,000 cores on a hydrodynamic AMR test problem we call “Alphabet Soup” involving an array of blastwaves driven by high-pressure regions in the shape of letters of the alphabet. Each core is assigned a blastwave in the $N \times N \times N$ array, where N ranges from 1 to 64. The problem is evolved with AMR until the blastwaves interact. We measure parallel efficiency for both memory and execution speed, finding 100% for the former and 85% for the latter at 262,000 cores. The largest problem evolved 1.7 trillion cells representing over 52 million Cartesian blocks of 32^3 cells arranged in a 64^3 array of octrees each with four levels of AMR refinement.

Weak and strong parallel scaling tests of a cosmology test problem including dark matter particles, adiabatic gas dynamics, and self-gravity on uniform meshes as large as $4,096^3$ and 110,000 cores have been performed with excellent results. We are currently repeating these tests with AMR turned on, but those results were not ready in time for this report.

The significance of this work is that it provides a path to exascale for the entire Enzo community, which numbers over 100 active developers and users. The technical approach is highly innovative as well, leveraging prior developments in parallel AMR algorithms (Forest-of-Octrees) and task-based dynamic execution with Charm++.

WHY BLUE WATERS

The scale of the system and our ability to get large scaling tests approaching full system size to execute quickly and reliably accelerated our code development substantially.

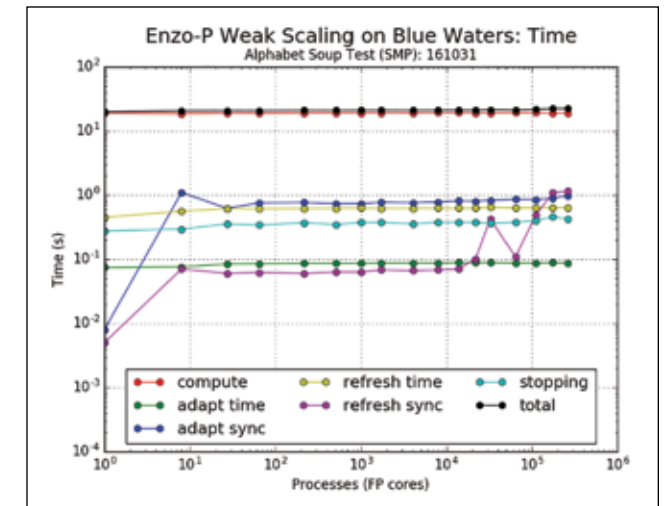


Figure 2: Weak scaling results for Enzo-P/Cello on the Alphabet Soup AMR hydrodynamics test problem. Plotted are components of the execution time per timestep per core versus floating point core count. The red line is time spent doing compressible hydrodynamics. The black line is total time including AMR and Charm++ overhead functions (other colored lines).

CORE-COLLAPSE SUPERNOVA SIMULATIONS: SIMULATING THE BRIGHTEST OBJECTS IN THE SKY AND THE SOURCE OF LIFE'S BUILDING BLOCKS

Allocation: Illinois/1,000 Knh
PI: Roland Haas¹
Co-PIs: Edward Seidel¹, Luke Roberts², Philipp Moesta³, Eliu Huerta⁴
Collaborators: Erik Schnetter⁵, Christian Ott

¹University of Illinois at Urbana-Champaign
²Michigan State University
³University of California, Berkeley
⁴National Center for Supercomputing Applications
⁵Perimeter Institute for Theoretical Physics

EXECUTIVE SUMMARY

Core-collapse supernovae explosions are cosmic engines that enrich the interstellar medium with the ashes of thermonuclear fusion inside massive stars. Their shock waves can shut off and trigger star formation and regulate the gas budget of galaxies. They are the birth sites of neutron stars and stellar mass black holes. In this project, we aimed to understand the influence of properties of the star just before the onset of explosion on the supernova and the signals that will be observed from the next supernova in the Milky Way.

RESEARCH CHALLENGE

Supernovae play a crucial role in the evolution of galaxies such as the Milky Way, and understanding the mechanism that makes them explode is of key interest to astronomers and astrophysicists.

METHODS & CODES

Our simulation code, Zelmani, is built on the Cactus framework and uses Berger–Olinger-type adaptive mesh refinement with full subcycling in time, HDF5 output for all data, hybrid OpenMP+MPI parallelization, and SIMD vectorization for the spacetime evolution. The magnetohydrodynamics use high-resolution, shock-capturing finite volume methods, and the neutrino transport code ZelmaniM1 currently employs an analytic closure.

RESULTS & IMPACT

We have published a first set of results in [1] where we present a study of the progenitor dependence of a three-dimensional neutrino mechanism of core-collapse supernovae. Our results suggest a complex, nonmonotonic dependence on progenitor parameters necessitating more detailed numerical studies to fully understand these effects and hinting at a complex interplay between multiple proposed explosion mechanisms.

WHY BLUE WATERS

Without Blue Waters, this research would have been impossible. Fully three-dimensional general-relativistic radiation-

magnetohydrodynamic simulations are simply too demanding of computational resources for any but a leadership-class facility to support them. Our simulations require the use of hundreds of compute nodes to provide sufficient memory for the simulation's state vector. The exceptional speed of Blue Waters' network lets us scale to the number of nodes required.

PUBLICATIONS & DATA SETS

Ott, C.D., et al., The Progenitor Dependence of Core-collapse Supernovae from Three-dimensional Simulations with Progenitor Models of 12–40 M_{\odot} . *Astrophys.J.*, 855 (2018). DOI:10.3847/2041-8213/aaa967.




Figure 1: Volume rendering of the strongly aspherical supernova shock (in blue) for different progenitor stars at the end of the simulations. Yellow and red colors indicate hotter gas behind the shock. Figure taken from [1].

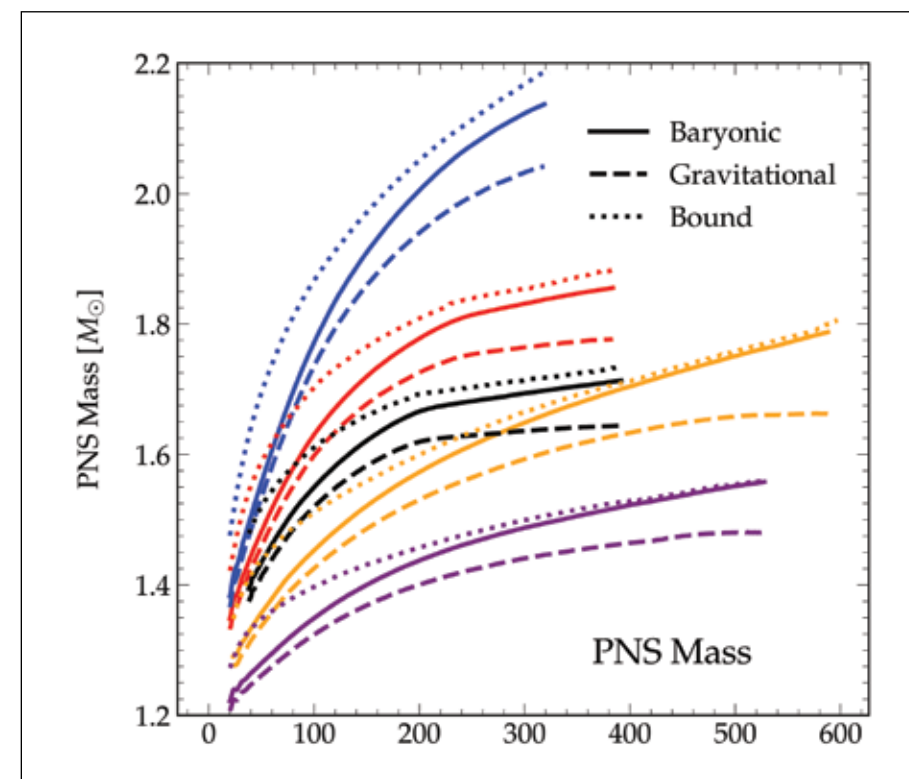


Figure 2: The mass of the neutron star forming at the center of the supernova for the different progenitor stars simulated in [1]. The highest mass simulation (in blue) is still accreting at 0.45 solar masses per second and will likely collapse to a black hole.

DEEP LEARNING FOR MULTIMESSENGER ASTROPHYSICS: REAL-TIME DISCOVERY AT SCALE

Allocation: Illinois/570 Knh

PI: Eliu Huerta¹

Collaborator: Daniel George²

¹University of Illinois at Urbana-Champaign

²National Center for Supercomputing Applications

EXECUTIVE SUMMARY

The observation of two colliding neutron stars in gravitational waves and light marks the beginning of multimessenger astrophysics. To accelerate discovery in this emergent field of science, we pioneered the use of deep learning for rapid detection and characterization of gravitational wave signals. We initially demonstrated this approach using simulated Laser Interferometer Gravitational-Wave Observatory (LIGO) noise. We have now shown for the first time that deep learning can detect and characterize gravitational wave signals in real (nonstationary and non-Gaussian) LIGO data, achieving similar sensitivities and lower errors compared to established LIGO detection algorithms. This new paradigm is far more computationally efficient and resilient to glitches, allowing faster-than-real-time processing of weak gravitational waves in real LIGO data with minimal computational resources and the detection of new classes of gravitational wave sources that may go unnoticed with existing detection algorithms. In addition, the new paradigm is ideally suited to enable real-time multimessenger discovery campaigns.

RESEARCH CHALLENGE

Matched-filtering searches, the most sensitive gravitational wave (GW) detection algorithms used by LIGO, currently target

a 3D parameter [1]. Extending these template-matching searches to target the 9D parameter space available to GW detectors is computationally prohibitive [2]. To address these limitations, we pioneered the use of GPU-accelerated deep learning algorithms [3]. Our technique, Deep Filtering, employs a system of two deep convolution neural networks (CNNs) that directly take time-series inputs for both classification and regression.

In our foundational article [3], we showed that CNNs can outperform traditional machine learning methods, reaching sensitivities comparable to matched-filtering for directly processing highly noisy time-series data streams to detect weak GW signals and estimate the parameters of their source in real time, using GW signals injected into simulated LIGO noise.

Deep Filtering demonstrated, for the first time, that machine learning can successfully detect and recover true parameters of real GW signals observed by LIGO, and achieve performance comparable to matched-filtering methods while being several orders of magnitude faster and far more resilient to transient noise artifacts, such as glitches. Furthermore, we showed that after a single training process, Deep Filtering can automatically generalize to noise having new Power Spectral Densities from different LIGO events without retraining.




Figure 1: The curve shows the sensitivity of detecting GW signals injected in real LIGO noise from our test set using Deep Filtering and with matched-filtering with the same template bank used for training. These results imply that deep learning is capable of detecting signals significantly weaker than the background noise.

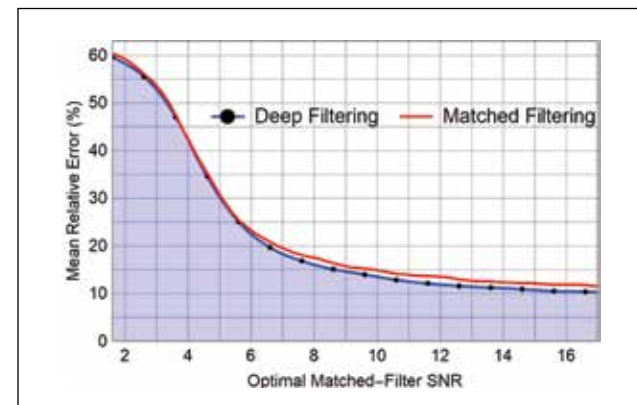


Figure 2: Mean percentage absolute error of estimating masses on testing signals at each SNR, injected in real LIGO noise from events not used for training, compared to matched-filtering using the same template bank that was used for training. Deep learning can interpolate to test set signals with intermediate parameter values.

METHODS & CODES

The training set contained about 2,500 waveform templates, generated with the open source EOB model [4], with black hole component masses between 5 and 75 solar masses sampled in steps of 1 solar mass. The input duration was fixed at 1 second, with a sampling rate of 8,192 Hz. The testing dataset also contained approximately 2,500 templates with intermediate component masses.

We produced copies of each signal by shifting the location of their peaks randomly within the final 0.2 seconds to make the CNNs more resilient to time translations. We obtained real LIGO data from the LIGO Open Science Center (LOSC) around the first three GW events; namely, GW150914, LVT151012, and GW151226. Each event contained 4,096 seconds of real data from each detector. We used noise sampled from GW151226 and LVT151012 for training and validation of our model and noise from GW150914 for testing.

We superimposed different realizations of noise randomly sampled from the training set of real LIGO noise from the two events GW151226 and LVT151012 and injected signals over multiple iterations, thus amplifying the size of the training datasets. The power of the noise was adjusted according to the desired optimal matched-filter signal-to-noise ratio (SNR) for each training round. The inputs were then whitened with the average Power Spectral Densities of the real noise measured at that time period.

We also scaled and mixed different samples of LIGO noise together to artificially produce more training data, and we also added various levels of Gaussian noise to augment the training process. However, the testing results were measured using only pure LIGO noise not used in training with true GW signals or with signals injected from the unaltered test sets.

We used the Wolfram Language neural network functionality, built using the open-source MXNet framework, that uses the cuDNN library for accelerating the training on NVIDIA GPUs. The learning algorithm was ADAM, and other details were the same as before [3]. While training, we used the curriculum learning strategy in our first article [3] to improve the performance and reduce training times of the CNNs while retaining performance at very high SNRs.

RESULTS & IMPACT

This research has shown for the very first time (Figs. 1 and 2) that CNNs can be used for both detection and parameter estimation of GW signals in raw LIGO data [5]. The intrinsic scalability of deep learning can enable fast, automated GW searches covering millions or billions of templates over the full range of parameter-space that is beyond the reach of existing algorithms. Extending Deep Filtering to predict any number of parameters such as spins, eccentricities, etc., or additional classes of signals or noise is as simple as adding an additional neuron for each new parameter, or class, to the final layer and training with noisy waveforms with the corresponding labels. Furthermore, the input dimensions of the

CNNs can be enlarged to take time-series inputs from multiple detectors, thus allowing coherent searches and measurements of parameters such as sky locations.

The average time taken for evaluating each of our CNNs per second of data is approximately 85 milliseconds and 540 microseconds using a single CPU core and GPU, respectively, thus enabling analysis even faster than in real time.

WHY BLUE WATERS

Blue Waters played a critical role in creating the numerical relativity waveforms used to train and test deep learning algorithms. In recent developments, Blue Waters has provided the required scale and computational power to construct deep neural networks using distributed learning involving over 1,024 GPUs.

PUBLICATIONS & DATA SETS

George, D., and E.A. Huerta, Deep learning for real-time gravitational wave detection and parameter estimation: Results with Advanced LIGO data. *Physics Letters B*, 778 (2018), pp. 64–70.

TINY GALAXIES HOST THE FIRST GIANT BLACK HOLES: BLUETIDES SIMULATION MAKES CONTACT WITH THE FIRST 700 MILLION YEARS OF COSMIC HISTORY

Allocation: NSF PRAC/1,856 Knh
PI: Tiziana Di Matteo¹
Co-PIs: Yu Feng², Rupert Croft¹
Collaborators: Ananth Tenneti¹, Steven Wilkins³

¹Carnegie Mellon University
²University of California, Berkeley
³University of Sussex

EXECUTIVE SUMMARY

Quasars, powered by supermassive black holes, are the most luminous objects known. As a result, they enable studies of the universe at the earliest cosmic epochs. A record-holding black hole with a mass 800 million times that of the Sun has recently been discovered. It formed when the universe was only 690 million years old—just 5% of its current age. The BlueTides (BT) simulation is the largest state-of-the-art hydrodynamical cosmological simulation that incorporates the processes of structure formation at cosmological scales with all the physics at much smaller galactical scales. Thus, it captures our most realistic understanding of black holes and their connection to galaxy formation in the early universe. These rare black holes grow at a super high rate, emitting enough energy to blow out gas that would eventually make stars in the galaxies. By growing so fast,

the black hole stunts the growth of its host galaxy. We predict that the JWST (James Webb Space Telescope, the successor of the Hubble) should be able to peek at the primordial environment of these extreme objects.

RESEARCH CHALLENGE

The first billion years is a pivotal time for cosmic structure formation. This epoch is also known as the cosmic dawn: the galaxies and black holes that formed then shaped and influenced all future generations of stars and black holes. Understanding and detecting the first galaxies and black holes is therefore one of the main observational and theoretical challenges in galaxy formation. The fundamental challenge to understanding this epoch of the Universe is that extremely large volumes need to be simulated as the first objects are rare, while at the same time extremely high resolution is required as the first galaxies and quasars are expected to be small and compact. With our BT simulation on Blue Waters (BW), this has now become possible, and it is the first and only cosmological simulation of structure formation that has run on the full BW machine.

A complete simulation of the universe at the epochs we are studying requires a small enough particle mass to model the first galaxies. It also requires an enormous volume, of the order of 1 cubic gigaparsec (1 Gpc³ is 3×10^{28} cubic light years) to capture the rarest and brightest objects. The first requirement is therefore equivalent to a high particle density, and the second to a large volume.

In the last year, we have extended the BT-I simulation to cover the evolution of the first billion years of cosmic history (second phase: BT-II). The goal is to significantly increase the scientific impact of this calculation to the community. The effort has paid off: BT is the only simulation that makes contact with the recently discovered supermassive black hole and can now make predictions regarding its formation, history, and future observations by next-generation telescopes.




Figure 1: Images showing the distribution of gas in the region centered on the most massive black hole of a size corresponding to the JWST field of view. The background image shows the distribution of gas (density color coded by temperature). The insets show the host galaxy of the brightest quasar as JWST should see it.

METHODS & CODES

The run was made possible through our new code, Massively Parallel (MP)-Gadget. MP-Gadget is a cosmological hydrodynamic simulation code. Recent radical updates to the code efficiency and the Smooth Particle Hydrodynamics formulation and star formation modeling mean that we have met the challenge of simulating the next-generation space telescope fields and are effectively using the full BW 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

The most distant known quasar (from 690 million years after the Big Bang) was recently discovered by [1]. We explore the host galaxy of this quasar in the large-volume cosmological hydrodynamic simulation BT, which in Phase II has reached these redshifts. The brightest quasar in BT has a luminosity of 10^{13} times that of the Sun and a black hole mass of 700 million solar masses, comparable to the observed quasar (Fig. 1).

The quasar resides in a rare halo of 10^{12} solar masses, and has a host galaxy comparable to the Milky Way's mass. We derive quasar and galaxy spectral energy distributions in the mid- and near-infrared range and make predictions in JWST bands. We predict a significant amount of dust is present in the galaxy. We present mock JWST images of the galaxy: The host galaxy is detectable in NIRCcam filters, but it is extremely compact (10 times smaller than our Milky Way). JWST's exquisite sensitivity, resolution, and wide wavelength coverage will be essential (and, we hope, sufficient) to constrain the stellar mass of these tiny host galaxies that are currently undetectable with the Hubble Space Telescope.

Many theoretical models predict that quasar-driven outflows account for the observed quenching of star formation in massive galaxies. There is growing observational evidence for quasar-launched massive outflows even in the very early universe. This is referred to as "quasar feedback." We have studied the feedback around the highest redshift quasar in the BT simulation and predict that there are significant outflows around this quasar. The gas is blown out from the galaxy; the quasar stops growing and so does its host galaxy.

The outflow gas contains a cold, dense molecular component that originates from the inner region of the halo, within a few kiloparsecs of the central black hole. This would be observable in CO emission at radio wavelengths. The velocities of the outflow gas reach thousands of km/s, within which the molecular component has mass averaged outward radial velocity of 1,300 km/s, consistent with observations. The averaged outflow rate has an enormous value, about 200–300 solar masses per year, or one hundred times greater than the current outflow from our own Milky Way galaxy.

We predict that the outflows we have seen in the simulated galaxy halo are likely to be present in the observed quasar. In addition, the presence of such significant quasar-driven outflows may help explain the low star formation rate in the host galaxy.




Figure 2: Conical gas outflows around the most massive black hole are powered by the quasar, which energizes and blows gas out from the host galaxy and also quenches star formation. The insert shows a histogram of the gas velocities around the black hole. The simulations predict this outflow should be observable around the first massive black holes.

WHY BLUE WATERS

In broad terms, the size of our problem meant that BW was the only system that could solve it, both in terms of memory and compute cores (we use essentially all system memory and cores). There were many essential aspects for this success. For example, scheduling allowed the system to be drained efficiently so our run could begin and make regular progress. Project staff returned aspects of the file system and scheduling. In addition, they identified obscure issues with the system, working with Cray in a fashion that we could never have accomplished. Thus, we achieved some of the highest I/O speeds seen on BW, which was essential to our research. The PAID program, linking computer scientists and domain scientists, was forward-thinking, progressive, and immensely useful.

PUBLICATIONS & DATA SETS

- Tenneti, A., et al., A tiny host galaxy for the first giant black hole: $z = 7.5$ quasar in BlueTides. *MNRAS*, submitted (2018).
- Ni, Y., et al., Gas outflows from the $z = 7.54$ quasar: predictions from the BlueTides Simulation. *MNRAS*, submitted (2018).
- Huang, K.-W., et al., BLUETIDES simulation: establishing black hole-galaxy relations at high-redshift. *MNRAS*, (2018), DOI:10.1093/mnras/sty1329.
- Bhomwick, A., and T. Di Matteo, The clustering of $z > 7$ galaxies: predictions from the BLUETIDES simulation. *MNRAS*, 474:4 (2018), pp. 5393–5405.
- Wilkins, S., et al., Dust-obscured star-forming galaxies in the early universe. *MNRAS*, 473:4 (2018), pp. 5363–5369.
- Tenneti, A., et al., The descendants of the first quasars in the BlueTides simulation. *MNRAS*, 474:1 (2018), pp. 597–603.

SHEDDING LIGHT ON INSPIRALING BINARY BLACK HOLES WITH MAGNETIZED MINI-DISKS

Allocation: NSF PRAC/1,500 Knh

PI: Scott Noble¹

Co-PIs: Manuela Campanelli², Julian Krolik³

Collaborators: Stephane d'Ascoli^{2,4}, Mark Avara², Dennis Bowen², Vassilios Mewes²

¹University of Tulsa

²Rochester Institute of Technology

³Johns Hopkins University

⁴École Normale Supérieure

EXECUTIVE SUMMARY

Observing electromagnetic and gravitational waves from supermassive binary black holes and their environments promises to provide important new information about both strong-field gravity and galaxy evolution. Since little theoretical understanding about the details of these accreting binary black hole systems exists, our project's aim is to continually advance the realism and rigor of simulations of these systems. The problem is complicated because dynamical general relativity, plasma physics, and radiation physics all must be calculated together over vast spatial and temporal scales. Our predictions will be critical to the success of electromagnetic searches and source characterization leading up to the launch of the Laser Interferometer Space Antenna (LISA). This past year, we finished performing the first magnetohydrodynamics simulations of black holes with their own mini-disks surrounded by a circumbinary disk. We have recently finished creating the first detailed electromagnetic predictions consistent with simulations using postproduction radiation transport.

RESEARCH CHALLENGE

Realistic accretion disk simulations are particularly challenging as they involve a multitude of physical processes interacting over large dynamic ranges in space and time. In actual systems, gas is collected at scales a million times larger than the black holes themselves, yet many cells per black hole width must be used to capture the relativistic plasma dynamics in their vicinity. Reliable angular momentum transport of gas through the disk requires solving the magnetohydrodynamics equations of motion at sufficiently high resolution to adequately resolve the responsible internal magnetic stresses. Consistency between the gas's thermodynamics and radiation model is desirable to produce self-consistent predictions of the light produced by the modeled systems, which is the ultimate goal of our program. Then, transporting the produced light to a distant observer requires us to calculate how light moves in curved, time-dependent spacetime to the black hole binary, and how it scatters and is absorbed by intervening gas; i.e., solving the general relativistic geodesic and radiative transfer equations.

All of this sophistication is built so we can confidently predict what electromagnetic counterparts may exist to the extremely bright gravitational wave sources LISA will see over cosmological distances. The first truly multimessenger event last year, GW/GRB 170817, provided a glimpse of the rich and varied information we can discover when gravitational waves accompany more conventional astronomical messengers such as light. Accreting supermassive binary black holes at the centers of galaxies are the only promising scenario for seeing light and gravitational waves from binary black holes, as stellar-mass systems in the local universe are unlikely to have any appreciable mass left over from their earlier stellar life. Since LISA is expected to launch in more than a decade, we can begin to search for these systems with our predictions in hand. Our simulations may discover features that are unique to binaries and, thus, inform the search for them.

METHODS & CODES

Our team uses the flux-conservative General Relativistic Magnetohydrodynamic (GRMHD) code called HARM3d. It is written in a covariant way such that arbitrary spacetime metrics and coordinates of systems may be used without the need to modify core routines. We exploit this property for the sake of time-dependent fixed mesh refinement to resolve the large range in spatial scales present in our systems of interest.




Figure 1: Volume rendering of the accretion flow's rest-mass density from the magnetized mini-disk simulation. (Credit: M. Van Moer, NCSA)

RESULTS & IMPACT

Our 3D magnetohydrodynamics simulation of mini-disk and circumbinary disk interactions advanced the field in a number of ways (Fig. 1). First, the inclusion of magnetic stresses provides greater realism to how gas is brought in from large distances and how it dissipates its orbital energy. The resolutions and run-times required to adequately describe this process are only possible on Blue Waters. Second, we discovered a new phenomenon in which the irregular circumbinary flow can modulate the rate of accretion onto the mini-disks. Because our black holes are relatively close, in that they are approaching the nonlinear merger phase of their evolution, the mini-disks drain faster than expected from standard Newtonian accretion disk theory. The fast depletion rate of the mini-disks leads them to be depleted, and then refilled, as they pass by the over-density feature in the circumbinary disk, suggesting a means to generate periodic electromagnetic emission.

In order to predict the electromagnetic emission from this simulation, we performed first-of-a-kind radiative transfer calculations in time-dependent general relativity using the simulation's data as an emitting source. A range of viewing angles and observation frequencies were surveyed for all time slices of data to explore the energy, time, and angle dependence of the emission. Our calculation resulted in the first electromagnetic spectrum of accreting supermassive black holes in the inspiral regime. Simulated images of the accreting black hole binary are shown in Fig. 2.

We are now exploring how this changes when the disk is tilted with respect to the binary's orbit, since the gas fed to the system need not always be aligned. We have completed a first survey and are analyzing the results now.

WHY BLUE WATERS

The 3D GRMHD mini-disk simulation ran for three orbital periods and used 18 million floating-point-core-hours or 1.2 million node-hours on Blue Waters. The simulation used $600 \times 160 \times 640$ or approximately 60 million cells on about 3 million timesteps using 600 nodes or 19,200 Blue Waters cores. Our latest tilted circumbinary disk simulations use up to 500 Blue Waters nodes at a time, with more than half the number of cells as the mini-disk simulation. NCSA Blue Waters staff, David King and Jing Li, were helpful in arranging reservations for our runs.

We have further benefited from Mark Van Moer's efforts over the past three years to produce state-of-the-art visualizations of our simulations. Recently, these included volume renderings of the gas density (Fig. 1) and a method to advect the seed points used to integrate the magnetic field lines using the fluid's velocity so that the lines are consistent across several frames in the animation.

PUBLICATIONS & DATA SETS

Bowen, D., et al., Quasi-periodic Behavior of Mini-disks in Binary Black Holes Approaching Merger. *Astrophysical Journal*, 853 (2018), DOI:10.3847/2041-8213/aaa756.




Figure 2: General relativistic ray-traced image of emission from magnetized gas accreting onto a binary system of black holes. The intensity of light is represented by the intensity of the hue. (Credit: S. Noble, University of Tulsa and NASA-GSFC)

DI

GA

TN

BW

CI

FS

BI

FEEDING BLACK HOLES: TILT WITH A TWIST

Allocation: NSF PRAC/5,220 Knh

PI: Alexander Tchekhovskoy¹

Co-PI: Nicholas Stone²

Collaborators: Matthew Liska³, Casper Hesp³, Adam Ingram³, Michiel van der Klis³, Sera Markoff³, Koushik Chatterjee³, Zack Andalman⁴

¹Northwestern University

²Columbia University

³University of Amsterdam

⁴Evanston Township High School

EXECUTIVE SUMMARY

For the past 40 years, the scientific community has been wondering what happens to a black hole feeding on the disk of gas that is tilted relative to the black hole midplane. The analytic prediction of Bardeen and Petterson in 1975, [1] that such a disk aligns to conform to the black hole midplane, has not been seen previously in numerical simulations that took into account the crucial effects of general relativity, magnetic fields, and turbulent motions in the disk. Using the power of Blue Waters' GPUs, our team for the first time demonstrated that such disks indeed align with the black hole midplane and, surprisingly, break up into individual segments with important astrophysical implications.

RESEARCH CHALLENGE

Tilted accretion is common in astrophysical systems. In fact, we expect that nearly all black hole accretion disks are tilted at some level relative to the black hole rotational equator. This is because the gas that approaches the black hole from large distances has no idea which way the black hole is spinning. However, studies of such tilted accretion disks are extremely challenging, especially in the crucial regime of luminous, radiatively efficient accretion that powers bright quasars. Such accretion disks are razor-thin and difficult to resolve numerically, requiring high resolutions and adaptive grids to follow the body of the disk as it moves through the computational grid. In a seminal paper, Bardeen and Petterson [1] predicted that purely general relativistic effects of black hole spin would cause the inner regions of such disks to align with the black hole equatorial plane. However, this model used a simplified model of the disk. In nature, the accretion disks are magnetized and turbulent, yet no general relativistic simulation of such a tilted magnetized turbulent disk has shown the long-sought alignment.

METHODS & CODES

Using our new code H-AMR [2] (pronounced "hammer"), which includes adaptive mesh refinement, local adaptive timestepping, and efficiently runs on GPUs, we were able to overcome these challenges. H-AMR performs 10 times faster on a GPU than on similar vintage 16-core CPU. H-AMR is parallelized via MPI with domain decomposition and scales well to thousands of GPUs, achieving weak scaling efficiency of 85 per cent on 4,096 GPUs

on the Blue Waters supercomputer. The performance of the code allowed us to study tilted disks at higher resolutions and over longer durations than was previously possible.

RESULTS & IMPACT

More than 40 years after it was first proposed by Bardeen and Petterson, we have finally demonstrated the existence of the Bardeen–Petterson alignment of geometrically thin disks. In fact, at larger tilt angles, the outer misaligned part of the disk breaks off from the inner aligned part.

WHY BLUE WATERS

Blue Waters access has been instrumental to our ability to obtain these groundbreaking results, which require not only enormous amounts of computing power, but also fast interconnect to make use of hundreds of XK nodes. Further, Mark Van Moer helped us enormously with 3D visualization.

PUBLICATIONS & DATA SETS

Liska, M., et al., General relativistic simulations of moderately misaligned thin accretion flows. *Monthly Notices of the Royal Astronomical Society*, in preparation (2018).

Liska, M., et al, Formation of precessing jets by tilted black hole discs in 3D general relativistic MHD simulations. *Monthly Notices of the Royal Astronomical Society Letters*, 474 (2017), DOI:10.1093/mnrasl/slx174.




Figure 1: The inner part of a disk of half-thickness $h/r = 0.03$ tilted by 45 degrees aligns with the black hole midplane and breaks off from the outer misaligned part of the disk. This is the first demonstration of the Bardeen–Petterson effect in a general relativistic magnetohydrodynamic (GRMHD) numerical simulation.




Figure 2: If h/r decreases from 0.03 to 0.015, the size of the aligned part increases from 5 to 10 gravitational radii, indicating that the alignment effect becomes stronger for thinner disks. This is the simulation of the thinnest black hole accretion disk in general relativistic magnetohydrodynamic (GRMHD) numerical simulations.

3D STELLAR HYDRODYNAMICS OF H-INGESTION FLASHES AND SHELL MERGERS

Allocation: NSF PRAC/4,000 Knh

PI: Paul R. Woodward¹

Collaborator: Falk Herwig²

¹University of Minnesota

²University of Victoria

EXECUTIVE SUMMARY

Our teams at the University of Minnesota, led by PI Paul Woodward, and at the University of Victoria, led by collaborator Falk Herwig, are simulating brief events in the interiors of evolved stars that result in the 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, affects the ultimate expulsion of heavier elements into the surrounding interstellar gas, either by 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 detailed 3D simulations of the entirety of the stellar interior, for which Blue Waters is ideally suited.

RESEARCH CHALLENGE

The teams are simulating brief but important events in the lives of stars that can greatly affect the heavier nuclei that the stars produce. We began by concentrating on hydrogen ingestion flashes, in which unburned hydrogen-rich fuel is brought into the convection zone above the helium-burning shell. Recently, we have turned our attention to massive stars, where in addition to H-ingestion flashes there is the possibility of the convection zones above different nuclear-burning shells actually merging. These events can have important consequences for the nucleosynthesis of heavy elements in these stars. For massive stars, such events can significantly affect the conditions just before the star explodes, and in this way, influence the injection of heavier elements from these stars into the surrounding interstellar medium.

METHODS & CODES

We are exploiting the Piecewise-Parabolic Method (PPM) gas dynamics method 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. This new code also solves nonlinear dynamical equations in which all the largest terms arising from the unperturbed, hydrostatic state of the star have been explicitly cancelled out. This makes

the calculation accurate enough that the simulation is able to use 32-bit precision, which results in a doubling of the code's speed on Blue Waters. The new code also produces superior results on coarse grids, which increases its cost-effectiveness still further. In order to accurately describe the dynamics of fluid mixing at the convection zone boundary, we simulate many large-eddy overturning times in the convection zone above a nuclear-burning shell. The simulations are scaled to 2,176 nodes on grids of 768³ cells and to 7,344 nodes for grids of 1,536³ cells. It can perform up to 48 timestep updates per second on Blue Waters. Thus, the millions of timesteps needed to simulate an ingestion event accurately are eminently practical on Blue Waters with the new code.

This year the team turned its attention 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 with the goal of alternating between 1D and 3D simulation in order to span the time necessary in leading up to a shell merger. The 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. We have been working up an attack on this shell-merger problem by first addressing ingestion of hydrogen-rich gas into the convective core of a massive star while it is still on the main sequence. Here, it is important to include the effect of rotation.

Early results from recent simulations on Blue Waters are shown in Figs. 1 and 2. We have described the techniques we use to achieve the extreme scalability of our codes in [1]. In [2], we describe the technique we developed with partial support from the Blue Waters PAID program to enable GPU acceleration. Our new code's exclusive use of 32-bit arithmetic is essential in enabling it to exploit the power of GPU accelerators by dramatically reducing the code's cache footprint.

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. Global convection modes play an important role in these situations, making simulation difficult and costly [3–5]. 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 [6,7].

WHY BLUE WATERS

We have carried out our simulations on Blue Waters because of its special ability to run at a sufficiently large scale so that our large computations can each be completed in less than one week [1]. 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 & DATA SETS

Woodward, P.R., F. Herwig, and T. Wetherbee, Simulating Stellar Hydrodynamics at Extreme Scale. *Computing in Science & Engineering*. In press (2018).

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.

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.

Herwig, F., et al., Global Non-Spherical Oscillations in 3-D 4 π Simulations of the H-Ingestion Flash. *Astrophys. J. Letters*, 792 (2014), p. L3.

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.

Ritter, C., et al., Convective-reactive nucleosynthesis of K, Sc, Cl, and p-process isotopes in O-C shell mergers. *Monthly Notices Royal Astron. Soc.*, 474 (2018), pp. L1–L6.

Herwig, F., et al., Cyberhubs: Virtual Research Environments for Astronomy. *Astrophys. J. Suppl.*, 236 (2018).




Figure 1: Snapshots of the distribution of ingested gas from above the convection zone generated by core hydrogen burning in a model 25 solar mass star. The star at the right rotates more rapidly. We look along the rotation axis at the far hemisphere in each case.




Figure 2: For the same model star as in Fig. 1, but with less rotation, we see a prominent dipole-like convection flow. Left, entrained gas, as in Fig. 1. Right, radial component of velocity (outward is red and yellow, inward is blue and white). Flow goes directly through the center of the star.

3D NATURE OF COLLISIONLESS MAGNETIC RECONNECTION AT EARTH'S MAGNETOPAUSE

Allocation: NSF PRAC/1,800 Knh
PI: Yi-Hsin Liu¹
Co-PIs: Michael Hesse², William Daughton³
Collaborators: Tak Chu Li¹, Ari Le³, Shan Wang⁴

¹Dartmouth College
²University of Bergen
³Los Alamos National Laboratory
⁴University of Maryland, College Park

EXECUTIVE SUMMARY

Earth's magnetosphere (a region formed from its magnetic fields) 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 the mechanisms, magnetic reconnection, which refers to the breaking and reconnecting of oppositely directed magnetic field lines in a plasma, is arguably the most active process at the center of many spectacular events in our solar system.

Magnetic reconnection not only allows the transport of solar wind plasmas into Earth's magnetosphere, but also releases magnetic energy and changes the magnetic topology. At Earth's magnetopause, magnetic reconnection proceeds between the shocked solar wind plasmas and the magnetosphere plasmas. Many 3D properties of magnetic reconnection in such an asymmetric geometry remain unclear. In this work, we used 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 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 and magnetosphere coupling, and in 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 (the null line where magnetic reconnection occurs), which points in the direction that maximizes the speed characterizing the reconnection outflow, in such an asymmetric current sheet?" The solution of this problem remains unclear with our current understanding of magnetic reconnection; thus, we aim to study the 3D nature of reconnection. It will be a crucial step in the quest to predict the location and rate of flux transfer at Earth's magnetopause, improving the forecast of space weather.

METHODS & CODES

This project employs the particle-in-cell code, VPIC [1], which solves the relativistic Vlasov–Maxwell system of equations using an explicit charge-conserving approach. Charged particles are advanced using the leapfrog method 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.




Figure 1: In (a), the total current density on a 2D plane where $y = 0$. In (b), the x - y cut of the current density across the location of the intense current near the x -line. Similarly, in (c), the reconnected field; in (d), the electron outflow; in (e), the ion outflow.




Figure 2: The magnetic field exponentiation factor σ from a plane of seed points in each 3D simulation. The ridges of large values highlight the quasi-separatrix layer, which is an approximate x -line, and pair of separatrices. The red contour in each panel marks where the electron mix measure is 0.99.

RESULTS & IMPACT

The orientation and stability of the reconnection x -line in asymmetric geometry was studied in 3D systems. We initiated reconnection at the center of a large simulation domain to minimize the boundary effect. The resulting x -line had sufficient freedom to develop along an optimal orientation, and it remained laminar. Companion 2D simulations indicated that this x -line orientation maximizes the reconnection rate, which has an important implication. We then designed 3D simulations with one dimension being short to fix the x -line orientation but long enough to allow the growth of the fastest growing oblique tearing modes. This numerical experiment suggested that reconnection tends to radiate secondary oblique tearing modes if it is externally (globally) forced to proceed along an orientation not favored by the local physics. The development of oblique structure easily leads to turbulence inside small periodic systems. This result could help interpret the local geometry of reconnection events observed by the Magnetospheric Multiscale Mission (MMS) and perhaps help determine an appropriate LMN coordinate. The question we are exploring is also relevant to the upcoming ESA–CAS joint mission, the Solar wind Magnetosphere Ionosphere Link Explorer (SMILE), which will study the development of reconnection lines at Earth's magnetopause using X-ray and UV imagers.

In separate work, plasma parameters were selected to model MMS magnetopause diffusion region crossings with a varying guide field strength. In each case, strong drift-wave fluctuations were observed in the lower-hybrid frequency range at the steep density gradient across the magnetospheric separatrix. These fluctuations give rise to cross-field electron particle transport. In addition, this turbulent mixing led to significantly enhanced electron parallel heating in comparison to 2D simulations. Comparing different methods of quantifying the anomalous dissipation revealed complications. Nevertheless, the anomalous

dissipation from short wavelength drift fluctuations appeared weak for each case, and the reconnection rates observed in 3D were nearly the same as in 2D models. The 3D simulations exhibited a number of interesting and new features that are consistent with recent MMS observations.

WHY BLUE WATERS

Because the x -line has a dimensional range 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 the magnetohydrodynamics scale. A representative 3D run in this project traced the motion of two trillion charged particles under the interaction of self-generated electromagnetic fields, which are evaluated on six billion grids. The output data can easily have a size of hundreds of terabytes 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.

PUBLICATIONS & DATA SETS

Yi-Hsin Liu, Y.-H., et al., Orientation and Stability of Asymmetric Magnetic Reconnection X-line. *Journal of Geophysical Research*, 123:6 (2018), DOI:10.1029/2018JA025410.

Ari Le, W.D., et al., Drift Turbulence, Particle Transport, and Anomalous Dissipation at the Reconnecting Magnetopause. *Physics of Plasmas*, 25 (2018), DOI:10.1063/1.5027086.

DI

KINETIC SIMULATIONS OF PLASMA TURBULENCE

Allocation: NSF PRAC/3,850 Knh

PI: Vadim Roytershteyn¹

Co-PIs: Yuri Omelchuk¹, Kai Germaschewski²

Collaborators: John Podesta¹, Stansilav Boldyrev³, Gian Luca Delzanno⁴, William Matthaeus⁵

¹Space Science Institute

²University of New Hampshire

³University of Wisconsin

⁴Los Alamos National Laboratory

⁵University of Delaware

EXECUTIVE SUMMARY

Our 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 the majority of situations encountered in space and in 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 a distinct population of the same species (e.g., the thermal part of the distribution or the energetic particles). An understanding of these issues may help advance our knowledge of systems where turbulence operates. For example, both the solar wind and the solar corona are famously known to be significantly hotter than could be explained with simple models, with local heating by turbulence often proposed as one of the likeliest explanations.

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 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.

METHODS & CODES

The most complete description of the plasmas of interest is provided by so-called Vlasov–Maxwell equations, a six-dimensional system of partial differential equations. In order to solve these equations, we use two complementary approaches. One is the well-known particle-in-cell (PIC) technique, which represents plasma as a collection of computational particles, while equations describing electromagnetic fields are solved on a computational grid. A typical large-scale simulation can simultaneously track upwards of a trillion particles in order to obtain a reliable statistical description. This requires a petascale computational resource such

as Blue Waters. In this work, we use VPIC: a general-purpose, high-performance plasma simulation code.

While PIC simulations have been quite successful in describing many microscopic physical phenomena, they also have a number of well-recognized limitations. For example, particle methods tend to have low accuracy or, more precisely, slow convergence rates associated with a finite number of particles per cell. To overcome these limitations, we investigated another approach based on fully spectral decomposition of the plasma species distribution function in the phase space. We use the code SpectralPlasmaSolver (SPS), developed in collaboration with Los Alamos National Laboratory. SPS uses dual Fourier–Hermite bases, is fully implicit, and possesses exact conservation laws for long-term, accurate simulations.

RESULTS & IMPACT

During the last year, we focused the investigation on characterizing kinetic plasma turbulence in the regimes where the ratio of plasma internal energy to the magnetic energy density (the so-called plasma parameter β) is small. Such regimes are encountered, for example, in the solar corona, the Earth's magnetosheath, interplanetary coronal mass ejections, regions downstream of collisionless shocks, hot accretion flows, and elsewhere.

In the summer of 2018, NASA will launch a revolutionary spacecraft mission—the Parker Solar Probe (PSP)—a "NASA mission to touch the Sun." PSP will pass as close as 9.8 solar radii from the Sun's surface, providing measurements that are expected to revolutionize our understanding of the solar corona and the solar wind, and will eventually lead to better models of space weather. There exists an urgent need to develop theoretical models and computational tools that correctly treat plasma turbulence in the regimes to be encountered by PSP.

Meanwhile, the majority of prior investigations of kinetic processes in turbulence focused on the regimes with β order unity, which is typical of the solar wind at the distances corresponding to the Earth's orbit. Recent theoretical analyses suggest that kinetic plasma turbulence is strongly modified under the parameters expected to be encountered close to the Sun, relative to solar wind turbulence at the Earth's orbit [1,2]. This change in the

nature of fluctuations affects many of the practically important processes, such as the mechanisms of energy dissipation. Using a combination of 3D SPS and 2D VPIC simulations, we have been able to confirm the existence of this new regime of plasma turbulence. Statistical properties of the turbulence measured in the simulations agree well with the theoretical predictions. These results are summarized in a publication submitted to the *Astrophysical Journal* [3].

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 a high-performance computing resource like Blue Waters and cannot be conducted on cloud resources.

PUBLICATIONS & DATA SETS

Roytershteyn, V., and G.L. Delzanno, Spectral Approach to Simulations of Kinetic Phenomena in Plasmas Based on Hermite Decomposition in the Velocity Space. *Frontiers in Astronomy and Space Sciences*, submitted (2018).

Roytershteyn, V., et al., Numerical Study of Inertial Kinetic–Alfvén Turbulence. *Astrophysical Journal*, submitted (2018).




Figure 1: Visualization of small-scale structure in kinetic plasma turbulence. The simulation was conducted using a newly developed version of spectral, fully implicit code SPS [4].

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

Allocation: Illinois/224 Knh
 PI: Deborah Levin¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

The main objective of our Blue Waters proposal has been to characterize the backflow contamination environment due to plasma created by electric-propulsion (EP) plumes and 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 have built on our earlier work where we developed an object-oriented C++ Direct Simulation Monte Carlo (DSMC) code that uses 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 this computational work has been to take advantage of our recent unique advances in GPU multithread parallelization applied to tree-based computational strategies. Blue Waters has been especially crucial to this modeling since we were able to use up to 256 GPUs per run for the plasma plume simulations on the XK nodes. During this work, we implemented the first DSMC calculations coupled to the particle-in-cell method with a particle approach for electrons as well.

RESEARCH CHALLENGE

The main objective of our work is to characterize the backflow contamination environment due to the plasma created by EP plumes, and their interaction with the spacecraft environment




Figure 1: Non-Maxwellian Electron velocity distribution functions of a shifted source plume as a function of the evolution time of an electric propulsion xenon ion plume. The asymmetry of the main xenon beam and the electron neutralizer means that the electrons cannot be described by an electron temperature.

and neutralizer sources, using state-of-the-art high-performance GPU-based petascale computations. The improved predictability of key surfaces of solar cell arrays and spacecraft charging in the backflow environment of chemical and EP thrusters is crucial, particularly for small satellites and CubeSats where there is close placement of thrusters and solar cell arrays. To model the spacecraft environment, multiple sources and processes that must be considered include both neutral and charged species in the backflow region of an onboard EP thruster as well as highly reactive ion thrusters and ambient neutral species that occur over multiple time and length scales. The indirect environmental exposure of spacecraft material such as the micron-sized glass and aluminized Mylar coatings of solar arrays can cause appreciable sputtering and erosion. Sputtering and erosion are hard to quantify and predict because the backflow ion fluxes are about five orders of magnitude less than those due to main ion beam impingement. In addition, the external hollow cathodes are a source of electrons that must be modeled kinetically if one is to truly understand how charge-exchange (CEX) ions affect the erosion of solar cell panels.

METHODS & CODES

Our plasma modeling is an outgrowth of our DSMC code, CHAOS (Cuda-based Hybrid Approach for Octree Simulations), that was developed under a previous Blue Waters effort to study neutral flows through porous media. The approach is unique in that in order to compute volume of cut-leaf nodes it utilizes the Morton Z-curve octree structure, a volume-of-fluids (VOF) method, and ray-tracing, which is very efficient on GPUs. Since the space conditions are such that the local mean free path for collisions is about three orders of magnitude larger than the local Debye length ($\sim 10^{-6}$ m), we have implemented two linearized Morton-ordered forests of octrees (FOT) so that these grids can be adapted to meet these two diverse numerical criteria. We have shown in [1] that when a leaf node may only be one level larger than its smallest-faced neighbor (the “2:1 criterion”) we are able to obtain 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.

In addition, we have implemented new boundary conditions that enable us to model a stable, steady plume as the beam-front crosses the computational domain. Further, we have obtained a




Figure 2: CEX ion streamlines (white) superimposed on an ion thruster, ion velocity field. Vertical strips are notional solar cell arrays. Ion impact fluxes differ for two electron models (Boltzmann vs. polytropic).

physically reasonable result where the electron distribution function remains Maxwellian for a co-located ion and electron source with an electron temperature less than about 2 eV. This computational complexity could be avoided by modeling a much larger computational domain, but the computational cost increases to the point where the number of resources would be prohibitive. Finally, as is shown in [1], we obtained ideal scaling up to 128 GPUs. This work received a prestigious best paper award at the 45th International Conference on Plasma Science in 2018 [2].

RESULTS & IMPACT

Modeling of electron species directly requires very small timesteps to accommodate typical velocities of 200 km/s. Therefore, previous modeling of electrons as a separate species in electric-propulsion plumes assumed a much smaller ion-to-electron mass ratio than the true system and was limited to two-dimensional cases. With this research on Blue Waters we are able to model the actual xenon-to-electron mass ratio for three-dimensional geometries. Fig. 1 shows the evolution of the electron cross-stream velocity component at a location close to the thruster exit ($z = 0.005$ m, normalized by the electron plasma frequency, $\omega_e = 1.78 \times 10^8$ rad/s). The electrons were initialized with a temperature of 2 eV and an initial number density of $1.0 \times 10^{13}/\text{m}^3$. The figure shows that at $t_{\omega_e} = 400$, the spacing between the two peak locations decreases, and the normalized probability of

electrons with zero cross-stream velocity increases to 0.4. This supports the hypothesis that the ion beam is trapping the electrons, which, in turn, damps the electron oscillations. As the plume evolves to $t_{\omega_e} = 1000, 1500,$ and 3000 , the peak of the y-EVDF (electron velocity distribution function) lies at $v_e/v_{te0} = 0$, with a smaller peak at $v_e/v_{te0} = -8$ showing that most electrons do not oscillate, and the electrons emitted from the shifted source are immediately attracted toward the plume.

Fig. 2 shows an extension of these simulations, taking the fundamental output to a very practical level. Streamlines moving from right to left along the core of the plume (red region) produce thrust, whereas ions that stream back to the solar cell arrays (yellow vertical strip on the main plot) cause damage. Because the distribution of ions flowing in the backward direction toward the spacecraft depends on the electron model, only a fully kinetic approach can give a truly predictive answer.

WHY BLUE WATERS

Blue Waters has allowed us to test and develop our algorithms on a large number of GPUs for three-dimensional, fully kinetic plasma simulations. Compared to the present state-of-the-art plasma simulations, a uniform grid in 3D would require a factor of at least ten more cells than our use of AMR/octree. The use of GPUs vs. CPUs decreased the runtime by at least another factor of five.

UNIFIED MODELING OF GALAXY POPULATIONS IN CLUSTERS

Allocation: NSF PRAC/9,000 Knh

PI: Thomas Quinn¹

Co-PI: Fabio Governato¹

Collaborators: Michael Tremmel², Arif Babul³, Iryna Butsky¹

¹University of Washington

²Yale University

³University of Victoria

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 galaxy 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 tackling this challenging problem. Preliminary results show that models that have successfully reproduced the morphology and number densities of isolated 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 galaxy 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 the baryons are visible via X-ray and microwave radiation. In contrast to field galaxies, where feedback from supernovae and AGN puts gas into a mostly invisible circumgalactic medium (ICM), feedback from cluster galaxies will impact the state of the 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. ChaNGa 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 one-half million cores on Blue Waters [2].

ChaNGa includes a well-constrained model for star formation and feedback, and improved implementations of supermassive black hole formation, growth, mergers, and feedback [3,4]. In our previous Blue Waters allocation, 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 (via the Sunyaev-Zeldovich effect) to understand the relation between these observables and the underlying gas properties. Finally, the overall mass distribution will be used to better understand how these clusters gravitationally lens background galaxies.

RESULTS & IMPACT

We have completed the simulation of several smaller galaxy clusters in preparation for our flagship simulation. Even the completed smaller simulations are advancing the state of the art in the simulation of galaxy clusters. Of particular significance is that models based on energy injected from supernovae and active galactic nuclei in field galaxies are able to naturally explain the properties of clusters and the galaxies within them. In this way, the models become predictive of the growth of galaxies and the black holes that power the active galactic nuclei.

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 of order 10^5 solar masses to accurately follow star formation and galaxy morphology. Likewise, we need to model a galaxy cluster of order 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.

PUBLICATIONS & DATA SETS

Tremmel, M., et al., Dancing to CHANGA: a self-consistent prediction for close SMBH pair formation time-scales following galaxy mergers. *MNRAS*, 475:4 (2018), DOI:10.1093/mnras/sty139.

Tremmel, M., et al., Wandering Supermassive Black Holes in Milky-Way-mass Halos. *Ap. J.*, 857:2 (2018), DOI:10.3847/2041-8213/aabc0a.

Lentz, E., et al., A New Signal Model for Axion Cavity Searches from N-body Simulations. *Ap. J.*, 845:2, DOI:10.3847/1538-4357/aa80dd.

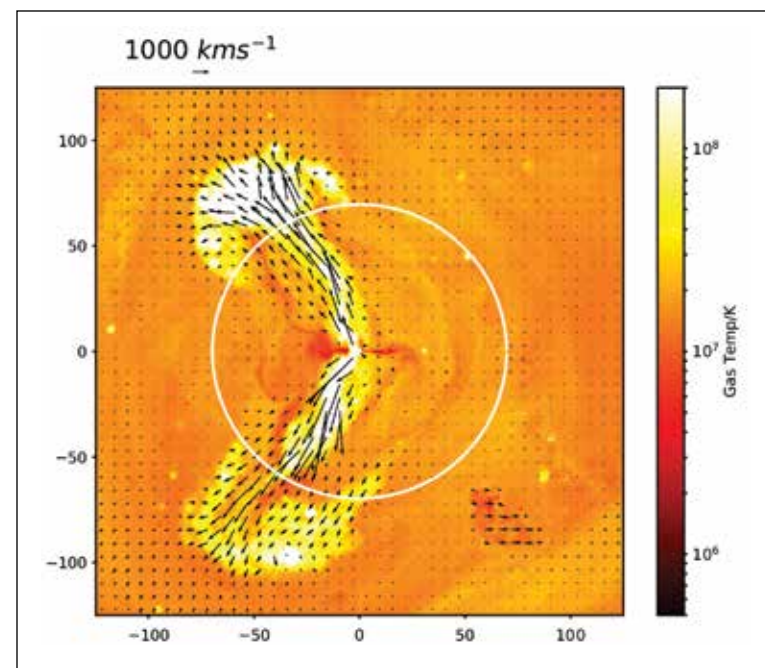


Figure 1: Outflows from the active galactic nucleus (AGN) of the central cluster galaxy. The columnation of the flow is emergent from the simulation: the energy from the AGN is injected isotropically. This outflow counteracts the cooling of the intracluster gas and suppresses star formation.

ASTROPHYSICS ON GRAPHICAL PROCESSING UNITS

Allocation: Illinois/230 Knh

PI: Hsi-Yu Schive¹

Co-PI: Matthew Turk¹

Collaborators: John ZuHone², Nathan Goldbaum³, Massimo Gaspari⁴, Chin-Yu Chen¹, Ui-Han Zhang⁵

¹University of Illinois at Urbana-Champaign

²Harvard-Smithsonian Center for Astrophysics

³National Center for Supercomputing Applications

⁴Princeton University

⁵National Taiwan University

EXECUTIVE SUMMARY

GAMER is a GPU-accelerated adaptive mesh refinement (AMR) code for astrophysics. In this project, we have incorporated a rich set of new physics and functionality into the code, including particles, magnetic field, chemistry and radiative processes, star formation, adaptive time-stepping, memory pool for efficient memory reuse, bitwise reproducibility, and optimization on load balancing. To demonstrate the capability of GAMER, we directly compared it with the widely adopted code Enzo [2] on a realistic astrophysical application, i.e., star formation in an isolated disk galaxy. This comparison showed that GAMER not only produces physical results that are very consistent with Enzo but also runs significantly faster, by an order of magnitude. In addition, we demonstrated good weak scaling with a parallel efficiency of 60%–70% and an overall performance of 5×10^{10} cells/sec using 4,096 GPUs and 65,536 CPU cores simultaneously for both hydrodynamic and magnetohydrodynamic applications.

RESEARCH CHALLENGE

Efficient utilization of the computing power of heterogeneous CPU/GPU supercomputers like Blue Waters remains extremely challenging for astrophysical AMR codes, especially when targeting realistic astrophysical applications requiring a rich set of physics and when running with a large number of nodes. The main challenge lies in the fact that the performance bottleneck, in general, does not reside in a single computational task. Instead, multiple tasks can be comparably expensive, for example, hydrodynamic solver, Poisson solver, chemistry and radiative processes, particle update, and AMR data structure reconstruction. The improvement in overall performance using GPUs is thus usually very limited if only certain specific tasks are ported to GPUs. On the other hand, not all tasks are GPU-friendly. Moreover, the GPU memory is still much smaller than the CPU's main memory. Therefore, it is necessary to efficiently utilize both CPU and GPU resources and carefully distinguish GPU-friendly and GPU-unfriendly tasks. Last, but not least, load imbalance and MPI communication may also become the performance bottleneck, especially if the computation time is significantly reduced by GPU utilization.

METHODS & CODES

According to the above, in this project we have significantly revised the GAMER code to incorporate a rich set of new physics, optimized the overall performance and parallel scalability, and directly compared it with other AMR codes on realistic astrophysical applications. The major features of GAMER include:

- *Multiphysics.* GAMER supports a rich set of physics, including hydrodynamics, magnetohydrodynamics, self-gravity, dark matter particles, chemistry and radiative processes, and star formation.
- *Adaptive mesh refinement* for resolving a large dynamic range.
- *Hybrid OpenMP/MPI/GPU.* GAMER uses GPUs as PDE solvers and CPUs to manipulate the AMR structure. It adopts OpenMP for intra-node parallelization in CPUs, MPI for inter-node communication, and CUDA as the GPU programming interface.
- *Asynchronous operations.* GAMER utilizes overlapping CPU computation, GPU computation, and CPU-GPU communication to boost the performance further.
- *Hilbert space-filling curve* for load balancing.
- *Memory pool* for efficient reuse of CPU memory.
- *Bitwise reproducibility* for scientific reproducibility.

GAMER uses the publicly available library GRACKLE [1] for the chemistry and radiative processes.

RESULTS & IMPACT

Fig. 1 shows the weak scaling of GAMER for the AMR-enabled hydrodynamic and magnetohydrodynamic (MHD) simulations on Blue Waters using up to 4,096 XK nodes. For the hydrodynamic test, we simulated a three-dimensional Kelvin–Helmholtz instability using the corner transport upwind scheme, piecewise parabolic reconstruction, and Roe's Riemann solver. For the MHD test, we simulated the Arnold–Beltrami–Childress (ABC) flow using the same hydrodynamic scheme but with the constraint transport technique to reinforce the divergence-free constraint on the magnetic field. GAMER achieves a parallel efficiency of 60%–70% and an overall performance of 5×10^{10} cells/sec. Note that when using, for example, 4,096 XK nodes, GAMER utilizes




Figure 1: Weak scaling of GAMER for the AMR-enabled hydrodynamic and MHD simulations using up to 4,096 XK nodes. Thanks to the hybrid CPU/GPU parallelization, GAMER can use both 4,096 GPUs and 65,536 CPU cores simultaneously and achieve a parallel efficiency of 60%–70% and an overall performance of 5×10^{10} cells/sec.




Figure 2: Performance speedup of GAMER over Enzo on the strong scaling of isolated disk galaxy simulations. GAMER outperforms Enzo by factors of four to eight. More importantly, the speedup increases when using more than 32 XK nodes, suggesting that GAMER also exhibits better parallel scalability.

both the 4,096 GPUs and the 65,536 CPU cores thanks to the hybrid OpenMP/MPI/GPU parallelization.

Fig. 2 shows the performance speedup of GAMER over Enzo for strong scaling using 1–128 XK nodes. Note that instead of using a simple test problem for the comparison, we conducted isolated disk galaxy simulations that combined gas, stars, and dark matter and incorporated hydrodynamics, self-gravity, particle dynamics, chemistry, radiative cooling/heating, and star formation. This comparison thus provides very convincing demonstrations of the accuracy and performance of GAMER. First, we found that the physical results obtained by the two codes are in very good agreement. Further, GAMER was measured as being four to eight times faster than Enzo. More importantly, this speedup ratio is roughly a constant when using a smaller number of nodes and increases when using more than 32 nodes, suggesting that GAMER not only runs faster, but also scales better. This significant performance improvement will allow us to study the dynamical evolution of the interstellar medium in unprecedented detail.

We have also compared GAMER with the AMR code FLASH [3] on cluster merger simulations and showed that GAMER is two orders of magnitude faster. Last, but not least, we have demonstrated that, for the uniform-resolution simulations, GAMER can achieve a maximum resolution of $10,240^3$ and $8,192^3$ for hydrodynamics and MHD, respectively, when using 4,096 XK nodes.

WHY BLUE WATERS

Blue Waters provides a unique opportunity for us to test and optimize GAMER on an extreme scale using thousands of GPUs and tens of thousands of CPU cores simultaneously. Moreover, the Blue Waters staff is very efficient in helping solve technical issues such as fine-tuning the hybrid OpenMP/MPI model.

Without their help, it would have been very difficult to achieve an optimal throughput.

PUBLICATIONS & DATA SETS

Zhang, U-H., H-Y, Schive, and T. Chiueh, Magnetohydrodynamics with GAMER. *The Astrophysical Journal Supplement Series*, 236 (2018), DOI:10.3847/1538-4365/aac49e.

Schive, H-Y., et al., GAMER-2: a GPU-accelerated adaptive mesh refinement code—accuracy, performance, and scalability. Submitted to *Monthly Notices of the Royal Astronomical Society*, arXiv: 1712.07070 (2017).

PETASCALE SIMULATIONS OF MERGING BLACK HOLES AND NEUTRON STARS

Allocation: NSF PRAC/8,200 Knh

PI: Saul A. Teukolsky¹

Co-PIs: Larry Kidder¹, Mark Scheel²

Collaborators: Francois Foucart³, Matt Duez⁴, Harald Pfeiffer⁵, Geoffrey Lovelace⁶, Scott Field⁷, Erik Schnetter⁸, Peter Diener⁹

¹Cornell University

²California Institute of Technology

³University of New Hampshire

⁴Washington State University

⁵Max Planck Institute for Gravitational Physics (Potsdam)

⁶California State University, Fullerton

⁷University of Massachusetts Dartmouth

⁸Perimeter Institute for Theoretical Physics

⁹Louisiana State University

EXECUTIVE SUMMARY

The primary purpose of the project is the numerical solution of Einstein's equations of general relativity. The goal is to track the coalescence and merger of binary black hole systems and to calculate the emitted gravitational waves (GWs). Another goal is to carry out a similar project for binary systems containing a black hole and a neutron star or two neutron stars. The work is aimed at providing theoretical predictions that can be compared with the signals measured by the National Science Foundation's LIGO (Laser Interferometer Gravitational-Wave Observatory) GW detector.

RESEARCH CHALLENGE

The primary scientific objective of our project is to theoretically underpin and improve the ability of LIGO to extract the rich information that the observed GWs carry. Gravitational waves provide a new window on the universe that will enable us to test our understanding of fundamental physics as well as learn about the most extreme events in the cosmos.

METHODS & CODES

Most of the computations are done with the SpEC code (Spectral Einstein Code) [1] developed by the collaboration. The numerical methods it uses make it the fastest and most accurate code for treating black holes [2]. We are also developing a new code, SpECTRE, with innovative methods to treat neutron star systems.

RESULTS & IMPACT

We will release a new version of our public catalog of gravitational waveforms for use by all researchers, largely through simulations on Blue Waters. The new version will increase the size of the catalog from 174 waveforms to well over a thousand. These waveforms have already been used to produce a very accurate waveform model that LIGO can use in its data analysis.

WHY BLUE WATERS

Our numerical code runs most efficiently on 50 to 70 processors for each waveform. Blue Waters' nodes are perfectly sized for us to use one or two nodes per waveform and explore hundreds of different parameter values to develop our catalog.




Figure 1: Two spinning black holes in orbit around each other, calculated by full solution of Einstein's equations of general relativity. In Newtonian gravity, the orbits would stay in a plane. Because the holes are spinning, the orbits precess (tilt).

CHEMISTRY AND THE EARLY UNIVERSE

Allocation: Director Discretionary/150 Knh

PI: Matthew Turk¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

This project focuses on constructing and executing models of the formation of the stellar nurseries of the first stars in the universe, simulating their hydrodynamic and chemical properties to high resolution. In doing so, we hope to constrain the masses of these stars and better understand their end products and descendants. In support of this, we have constructed a cross-domain software package for the efficient solution of chemical species abundances in hydrodynamic simulations. We also implemented a solver for the gravitational potential as defined in spherical coordinates in three dimensions.

RESEARCH CHALLENGE

Understanding the initial mass function of the first stars will guide our understanding of the sources of chemical elements in the modern universe, as well as help to refine our understanding of the progenitors of gravitational wave events.

METHODS & CODES

We used the GAMER-2 simulation code, the Dengo rate construction package, and the Grackle software package.

RESULTS & IMPACT

These developments will help to contextualize observations from the James Webb Space Telescope; making these software packages available and accessible will enable scientific inquiry in a number of related domains.

WHY BLUE WATERS

Blue Waters provided the necessary environment to ensure that our solver would scale to the size and capacity we require for high-resolution studies.

DI

SIMULATING GALAXY FORMATION ACROSS COSMIC TIME

Allocation: NSF PRAC/4,000 Knh

PI: Brian O’Shea¹

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

Collaborators: Cameron J. Hummels⁴, Britton D. Smith⁵, Molly Peeples⁶, Jason Tumlinson⁶, Lauren Corlines⁶

¹Michigan State University

²Florida State University

³Georgia Institute of Technology

⁴California Institute of Technology

⁵University of California at San Diego

⁶Space Telescope Science Institute

BW

MP

CI

BI

EXECUTIVE SUMMARY

This research project addresses several critical questions about galaxy formation and evolution via a suite of physics-rich, high-dynamic-range adaptive mesh refinement simulations of cosmological structure formation. The two main goals 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 elucidate 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 our 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. Each of these telescopes is 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 extremely high dynamic range in space and time, complex physics (including radiation transport and nonequilibrium gas chemistry), and large simulation volumes.




Figure 1: This figure shows a projection of (left to right) gas density, temperature, and metal content for two simulations: one with an enhanced spatial refinement in the circumgalactic gas (top row), and one using standard refinement methods (bottom row).

METHODS & CODES

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

RESULTS & IMPACT

The two main results thus far involve the growth of supermassive black hole candidates in the early universe, and the cycling of gas into and out of galaxies. The former result involves analysis and resimulation of work done in an earlier PRAC allocation, and shows that, while stellar-mass black holes are not capable of growing into billion-solar-mass objects by the time that they can be observed (a billion years after the Big Bang), it is possible for massive gas clouds to directly collapse into much more massive objects that can easily seed black holes [3,4].

Our second important result involves the cycling of gas into and out of galaxies. Simulations 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. We found that increasing the resolution by more than an order of magnitude beyond previous state-of-the-art calculations resulted in the appearance of both spatial and chemical features that are seen in observations but not in previous models. This work is revolutionizing our understanding of the interface between the stellar component of galaxies and the diffuse corona of gas that surrounds them.

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 in order to model the many early galaxies that merged together to create a Milky Way-like galaxy at the present day. Further, 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 data set sizes and large computational resources, as well as an extremely high bandwidth and 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 & DATA SETS

Smith, B., et al., The Growth of Black Holes from Population III Remnants in the Renaissance Simulations. *The Astrophysical Journal*, accepted (2018), arXiv:1804.06477.

Wise, J., et al., Ubiquitous massive black hole formation in rapidly growing pre-galactic gas clouds. *Nature Astronomy*, under review (2018).

ELUCIDATING THE ALIGNMENT MECHANISM FOR BLACK HOLE ACCRETION DISKS SUBJECTED TO LENSE–THIRRING TORQUES

Allocation: GLCPC/320 Knh

PI: John Hawley¹

Co-PI: Julian Krolik²

¹University of Virginia

²Johns Hopkins University

EXECUTIVE SUMMARY

The research team uses Blue Waters to study astrophysical accretion onto a spinning black hole in which there is a misalignment between the orbital axis of the incoming gas and the black hole rotation axis. Astrophysicists have long expected that an initially misaligned orbiting accretion disk would align with the black hole's spin axis at some location near the hole. A detailed understanding of this alignment process has, however, been limited due to the assumption of a phenomenological viscosity to describe the internal dissipation necessary for alignment. Our simulations capture the physical internal stress due to magnetohydrodynamic turbulence with no reliance on phenomenological viscosity; such simulations are only possible with the high grid resolution made feasible by Blue Waters. The investigation will probe how a time-steady transition might be achieved between an inner disk region aligned with the equatorial plane of the central mass's spin and an outer region orbiting in a different plane.




Figure 1: A contour plot of log density after 20 orbits of evolution subject to the external torque. Overlaid is a line showing the equatorial plane for the spin axis of 24° showing how the inner disk has aligned with the black hole.

RESEARCH CHALLENGE

Accretion disks occur in a wide variety of astrophysical contexts, from planet formation to black holes, both in binary systems and in the cores of active galaxies. Whenever the disk's angular momentum is oblique to the angular momentum of the central object(s), as should often be the case, a torque causes rings within the disk to precess, twisting and warping it. Because the torque weakens rapidly with increasing radius, it has long been thought that some unspecified "friction" brings the inner portions of such disks into alignment with the equator of the central object while the outer parts remain in their original orientation. Identifying and quantifying that friction has been difficult.

Despite considerable theoretical effort, it is not known how to predict the radius at which a given astrophysical disk would become aligned with the equatorial axis of the black hole. Nearly all previous work on this topic has assumed that such a disk's internal stresses can be described by a parameterized isotropic viscosity. However, there is a well-established physical mechanism for internal stresses in accretion disks—correlated magnetohydrodynamic (MHD) turbulence, driven by the magnetorotational instability [1,2].

Treating MHD turbulence in a tilted disk requires very demanding numerical simulations. The central difficulty is that if a numerical simulation is to follow the MHD turbulence, it must have a timestep very short in comparison to an orbital timescale, whereas the precession timescale where the orientation transition may occur is almost certainly many orbital periods long. This project is to carry out such numerical simulations to investigate how twisted disks align when their mechanics are described only in terms of real forces, including MHD turbulence. The aim is to develop a model that will predict the location of the stationary alignment front in disks subjected to Lense–Thirring (relativistic) torques.

METHODS & CODES

This work uses a simplified disk model first studied in [3], consisting of an isothermal disk orbiting a point-mass in Newtonian gravity with a Keplerian angular velocity distribution. It includes only a lowest-order post-Newtonian term to represent the relativistic Lense–Thirring torque. This idealized model allows us to focus on the important physical processes governing alignment; these can be studied in isolation and in detail. In this project we

examine the influence of sound speed and black hole tilt. We evolve a series of models where the dependence of sound speed and tilt are studied in the sense of measuring a partial derivative: We change those terms without altering anything else. The use of an isothermal equation of state further isolates the effect of sound speed by ignoring the spatial and temporal changes expected in a more complex model.

The simulations were done with a Fortran 95 version of *Zeus*, an operator-split code that solves the equations of compressible MHD by time-explicit finite-differencing [4] using the CT (constrained transport) algorithm [5] to preserve the divergence-free property of the magnetic field. *Zeus* uses domain decomposition and MPI for parallelization.

RESULTS & IMPACT

The simulation results confirm that the influence of the sound speed can be encapsulated in a simple "lumped-parameter" model first proposed by [6]. In this model, alignment fronts propagate outward from the inner disk at a speed proportional to the local test-particle precession frequency. Meanwhile, transonic radial motions transport angular momentum both inward and outward at a rate that may be described roughly in terms of an orientation diffusion model with diffusion coefficient proportional to c_s^2/Ω for sound speed c_s and orbital frequency Ω . The competition between the two leads to, in isothermal disks, a stationary position for the alignment front at a radius proportional to $c_s^{-4/5}$. For alignment to happen at all, the disk must either be turbulent due to the magnetorotational instability in MHD, or, in HD, it must be cool enough for the bending waves driven by disk warp to be nonlinear at their launch point. The alignment process is largely independent of black hole tilt angles, at least for the range of angles studied here.

WHY BLUE WATERS

We have used Blue Waters to compute new thin disk simulations subject to Lense–Thirring torque with unprecedented resolution to improve the representation of MHD turbulence in the disk, and to make use of a larger black hole tilt angle to increase the torque effect and further explore the mechanisms behind, and scaling properties of, alignment. The unique high-performance capabilities of Blue Waters enabled key linchpin maximum-resolution simulations that support a wider effort involving a suite of less demanding simulations carried out on other systems.

MODELING PHYSICAL PROCESSES IN THE SOLAR WIND AND LOCAL INTERSTELLAR MEDIUM WITH A MULTISCALE FLUID-KINETIC SIMULATION SUITE: CONNECTING SOLAR, HELIOSPHERIC, AND ASTROPHYSICAL SCALES

Allocation: NSF PRAC/2,140 Knh
PI: Nikolai Pogorelov¹
Co-PI: Jacob Heerikhuisen¹
Collaborators: Sergey Borovikov¹, Tae Kim¹, Mehmet Sarp Yalim¹

¹University of Alabama in Huntsville

EXECUTIVE SUMMARY

We have investigated physical phenomena occurring when the solar wind (SW) interacts with the local interstellar medium (LISM), including: (1) propagation of coronal mass ejections (CMEs) through the solar wind flow governed by photospheric magnetograms; (2) data-driven simulations of transient phenomena affecting space weather near Earth and other planets; (3) SW propagation throughout the heliosphere and perturbations it creates in the LISM; (4) the effect of nonthermal pickup ions (PUIs) on spacecraft measurements; (5) magnetohydrodynamic (MHD) instabilities and magnetic reconnection near the heliopause; (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 help interpret and predict IBEX, New Horizons, Ulysses, and Voyager spacecraft 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 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 Interstellar Boundary Explorer (IBEX), New Horizons, Ulysses, Voyager, and a fleet of near-Earth spacecraft. We use vector magnetogram data and STEREO and SOHO observations to study the propagation of coronal mass ejections toward Earth, where they affect space weather. The Voyager 1 and 2 (V1 and V2) spacecraft crossed the heliospheric termination shock (TS) in December 2004 and August 2007, respectively, and after 45 years of historic discoveries, V1 is sampling the 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 give 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 (ENAs) in different energy bands [3]. Since most ENAs are created during charge exchange between hot PUIs and LISM neutral atoms, they bear an imprint of plasma properties of




Figure 1: Comparison of CME shapes: (left panel) SOHO LASCO/C3 coronagraph difference image; (right panel) temperature contours obtained from the simulation.




Figure 2: The distribution of magnetic field strength in the plane determined by the Sun's rotation axis and the LISM velocity vector. Because of their large gyroradius, simulations of multi-TeV cosmic ray transport toward observation sites on Earth require large computational domains.

the region where they were 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 MHD equations coupled with the kinetic Boltzmann equation describing the transport of neutral atoms. In a less strict approach, we model the flow of atoms with a few systems of the Euler gas dynamic equations describing the different atom populations dependent on the domains of their origin.

We have developed both fluid dynamics and kinetic models for PUIs and turbulence generated by kinetic instabilities of their distribution function. All these are components of a Multi-Scale Fluid-Kinetic Simulation Suite (MS-FLUKSS), an adaptive mesh refinement code we have built on the Chombo framework. MS-FLUKSS is now capable of performing simulations with fourth-order adaptive mesh refinement (AMR), both in space and time on composite meshes, e.g., cubed spheres.

RESULTS & IMPACT

As a result of our Blue Waters allocation we have: (1) used coronagraph images from the SOHO and STEREO spacecraft to simulate CMEs, starting from the lower corona of the Sun (Fig. 1); (2) used the Wang-Sheeley-Arge coronal model driven by the Air Force Data Assimilative Photospheric Flux Transport (ADAPT) model to simulate SW properties at Earth and have demonstrated

advantages of our model over existing space weather prediction models; (3) refined our model for shocks propagating through the LISM, demonstrating good agreement with observational data, and predicting a few events to be observed by V1 in the future; (4) performed a time-dependent simulation of the heliosphere that produces a comet-like heliotail and showed that the 11-year solar cycle leads to the formation of ENA lobes with properties remarkably similar to those observed by IBEX; (5) analyzed the heliotail flow and quantity distribution in the heliospheric bow wave for different LISM conditions in a computational box of 12,000 by 10,000 by 8,000 astronomical units, and showed that the observed multi-TeV cosmic ray anisotropy may be explained by LISM magnetic field distortion by the heliosphere; (6) used simulations on adaptive meshes to further understand MHD instabilities and magnetic reconnection at the heliopause, and (7) reproduced observations of pickup ions by New Horizons. Our results have been published in six papers (one more paper is in press and three are in preparation) and reported at more than 10 scientific meetings. In addition, our research was highlighted by the American Astronomical Society and other web news outlets.

WHY BLUE WATERS

The Blue Waters system and project staff responded in a timely manner to our concerns and were very helpful in the development of job-scheduling strategies. The overall performance and reliability of Blue Waters has been outstanding.

PUBLICATIONS & DATA SETS

Kim, T., N. Pogorelov, and L. Burlaga, Modeling shocks detected by Voyager 1 in the local interstellar medium. *Astrophysical Journal Letters*, 843:2 (2017), p. L32.

Lamy, L., et al., The aurorae of Uranus past equinox. *Journal of Geophysical Research: Space Physics*, 122 (2017), pp. 3997–4008.

McComas, D.J., et al., Heliosphere responds to a large solar wind intensification: Decisive observations from IBEX. *Astrophysical Journal Letters*, 856 (108), p. L10.

Pogorelov, N.V., et al., Heliosheath processes and the structure of the heliopause: Modeling energetic particles, cosmic rays, and magnetic fields. *Space Science Reviews*, 212 (2017), pp. 193–248.

Pogorelov, N., et al., Three-dimensional features of the outer heliosphere due to coupling between the interstellar and heliospheric magnetic field. V. The bow wave, heliospheric boundary layer, instabilities, and magnetic reconnection. *Astrophysical Journal*, 845 (2017), DOI:10.3847/1538-4357/aa7d4f.

Zirnstein, E., J. Heerikhuisen, and M. Dayeh, The role of pickup ion dynamics outside of the heliopause in the limit of weak pitch angle scattering: Implications for the source of the IBEX ribbon. *Astrophysical Journal*, 855 (2018), DOI:10.3847/1538-4357/aaaf6d.

SIMULATING TWO-FLUID MHD TURBULENCE AND DYNAMOS IN STAR-FORMING MOLECULAR CLOUDS

Allocation: NSF PRAC/2,700 Knh
PI: Dinshaw S. Balsara¹
Co-PI: Alexandre Lazarian²
Collaborator: Blakesley Burkhart³

¹University of Notre Dame
²University of Wisconsin-Madison
³Harvard-Smithsonian Center for Astrophysics

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, and 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 and dynamo processes that regulate magnetic field evolution and star formation. This turbulence, and the nonideal physics to which it gives rise, is fundamentally different from single-fluid magnetohydrodynamic (MHD) turbulence. We aim to study these differences.

RESEARCH CHALLENGE

Our understanding of the star formation process has reached the point where advanced observational capabilities are required. Consequently, NASA has made multimillion-dollar investments in the High-resolution Airborne Wideband Camera-plus (HAWC+) instrument aboard the Stratospheric Observatory for Infrared Astronomy (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 multiyear-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 and Blue Waters PRAC resources ([4]; [7–9]; Xu, et al., 2018a and 2018b). The partially ionized fluid in a GMC is best modeled as a neutral fluid and an ionic fluid that interact with the magnetic field. Such simulations are extremely CPU-intensive and only simulations with very limited resolution (5,123 zones) had been done before this Blue Waters PRAC award. At those resolutions, we would have been unable to match the observations from HAWC+. With our new generation of simulations, with 1,024³ zones and upwards resolution, we are

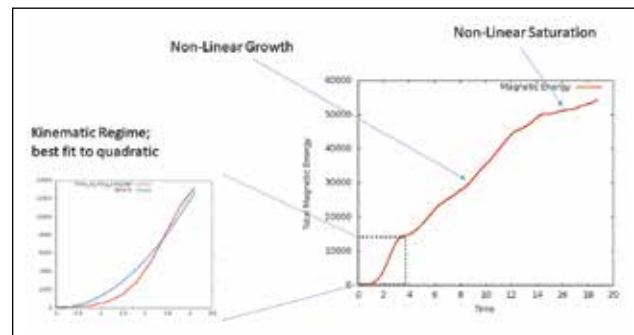


Figure 1: Fig. 1, from Xu, et al. (2018a), shows the evolution of magnetic energy with time in a two-fluid dynamo. The right panel shows magnetic field growth in a partially ionized dynamo. The inset on the left shows the initial growth of magnetic energy, with the red curve representing our highly resolved simulations and the blue curve showing the best fit to a quadratic.

able to obtain a well-defined inertial range in two-fluid and dynamo simulations. This ensures that NASA’s investment in HAWC+ is being matched by well-resolved simulations.

The most compelling motivation for understanding two-fluid, ambipolar diffusion-mediated turbulence comes from recent observations. Differences in the linewidths between neutral and ionized tracers have led to the suggestion that the dissipation of turbulence from ambipolar diffusion sets in on scales smaller than 0.0018 parsecs in M17 [2].

A breakthrough realization by Xu and Lazarian [7] claimed that magnetic fields would grow in a partially ionized plasma at rates that are very different from the growth in a fully ionized plasma. The growth of the magnetic field in a plasma is known as the dynamo problem. Their prediction was that while magnetic fields grow exponentially in a fully ionized plasma, they grow only quadratically in a partially ionized plasma. However, confirmation of the theory requires a highly resolved turbulent flow. Fig. 1 from the first listing under Publications and Data Sets shows the evolution of magnetic energy with time in a two-fluid dynamo. The right panel in Fig. 1 shows magnetic field growth in a partially ionized dynamo. The inset on the left shows the initial growth of magnetic energy with the red curve representing our highly resolved simulations and the blue curve showing the best-fit to a quadratic. The agreement with theory is very good and has




Figure 2: The study of the magnetic energy spectra is also significant. The peak in the magnetic energy spectrum is predicted to evolve with time at a set rate and the results conform to theory.

been documented. The study of the magnetic energy spectra is also significant and is shown in Fig. 2. The peak in the magnetic energy spectrum is predicted to evolve with time at a set rate and our results conform to theory.

METHODS & CODES

The core MHD algorithms in our RIEMANN code are based on higher-order Godunov schemes (conservative numerical schemes for solving partial differential equations). 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. We have described two-fluid methods [3], [5], and [6].

RESULTS & IMPACT

We have completed several large-scale simulations while others are still continuing on Blue Waters. The key points from our simulations have been documented above and are being readied for publication. This 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 the simulation code.

WHY BLUE WATERS

The simulations reported here are extremely CPU-intensive. The goal of this project is to use the petascale computing power of Blue Waters to push the resolution, accuracy, and fidelity of the simulations much higher so we can match theory with the observations coming from NASA-funded instruments.

PUBLICATIONS & DATA SETS

Xu, S., S. Garain, A. Lazarian, and D.S. Balsara, Turbulent dynamo in a weakly ionized medium—Study of growth rate. *Astrophysical Journal*, in preparation (2018a).

Xu, S., S. Garain, A. Lazarian, and D.S. Balsara, Turbulent dynamo in a weakly ionized medium—Spectral evolution. *Astrophysical Journal*, in preparation (2018b).

THE EVOLUTION OF THE HALO MASS FUNCTION IN AN ACCELERATED UNIVERSE

Allocation: Exploratory/50 Knh
PI: Matias Carrasco Kind¹
Co-PI: Robert Gramillano¹
Collaborators: Thomas Kirsh¹, Francisco Paz-Chinchon¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

We live in an accelerated expanding universe that is dominated by a cosmological constant, or the value of the energy density of the vacuum of space, also known as “dark energy.” This dark energy acts as a negative pressure pushing everything away and making the universe expand at faster rates. In the distant future, dark energy will completely dominate the evolution of large-scale structures. Based on a simple theoretical criterion, we can determine the limits of gravitationally bound, dark-matter halos in their critical shells for virialized haloes. (A system is virialized when the potential energy is twice the negative kinetic energy.) 70 billion years into the future, this predicts that the density inside the last bound shell has to be twice the critical density of the universe at that time. Using Blue Waters, we applied this criterion far into the future by using N-body simulations of dark matter particles, where all of the halos are fully virialized and “frozen” in the cosmic web. This simulation produced a halo mass function that reaches an asymptotic form, from which we can extract useful information to constrain the large-scale structures we observe today and put constraints on current cosmological models.

RESEARCH CHALLENGE

Recent observations suggest that we live in an expanding, zero-curvature universe dominated by a cosmological constant. This vacuum energy density, the so-called dark energy, appears to dominate the expansion rate of the universe today. In the present cosmological scenario, the universe transitioned from a matter-dominated stage to a dark-energy-dominated, accelerating stage at redshift $z \approx 0.7$. During the period of time when the universe was substantially dominated by matter density over dark-energy density, decelerating its expansion, the formation and evolution of structures were driven only by gravitational instability of matter

concentrations, producing structures that grew over time. At later times, when dark energy came to govern the dynamics of the universe, accelerating its expansion, the formation of structures slowed down. At those times, massive structures much denser than the dark energy will not be much affected by this acceleration and will remain gravitationally bound. As they evolve into the future, these structures will be fully virialized: They will be spherical, and they will stop accreting mass so they will be “frozen” in the future cosmic web.

By using simulations of the future evolution of dark matter particles, it is possible to place constraints on the bounds of superstructures by following the mass evolution of the biggest clusters found at redshift $z = 0$ (present day) in the simulation. However, due to resolution limitations in previous studies, a complete understanding of the evolution of mid-size halos based, among other properties, on local environments in a universe with a cosmological constant still remains incomplete. We wanted to understand how the mass of these structures evolves with time and how we can use that information to put constraints on today’s cosmological models of structure formation.

METHODS & CODES

We carried out several dark matter particle N-body simulations from very far in the past, when the universe was only 2% of its current size, to the future, where the universe will be 100 times its current size, in order to understand the evolution of dark matter haloes. Since our simulations included only dark matter particles, we evolved them in time using the parallel tree N-body/smoothed particle hydrodynamics (SPH) code GADGET-2 [1], which is the main software used. Only the tree code was used for this simulation. These simulations started at redshift $z = 40$ and




Figure 1: The particles that eventually in the future will collapse into a massive, virialized, and spherical halo. From left to right: $a=1$ (present), $a=1.5$, $a=8$ and $a=100$ (future). We see that the differences increase with the scale factor, and the structures at $a=100$ appear physically isolated.




Figure 2: A slice of one of the simulations for the same 50 megaparsec/hour (Mpc/h) region at the present ($a=0$, left) and in the distant future ($a=100$, right). Halos are more spherical and well defined in the future; everything is virialized, and no more mass accretion is occurring.

evolved until the scale factor reached 100, which is about five times the current age of the universe and with a size 100 times larger.

Our main bottleneck is the gravitational force computation that GADGET-2 completes for each timestep, but since our simulation is well within the XE compute node’s peak performance limits, this was not an issue. All of the runtime libraries needed for the GADGET-2 simulation are already included in Blue Waters’ software and packages list: hdf5 (for writing the output files), mpich2 (for communication), and fftw (for tree force computation).

RESULTS & IMPACT

Using Blue Waters, we were able to run several simulations at different scales; given the exploratory nature of our allocation, we are still processing and analyzing the data generated.

We can see from Fig. 1 that even the largest haloes in the future with close to 80,000 particles are very spherical and virialized, where only the dark energy is dominating the evolution and no more mass is being accreted to these structures. Fig. 2 shows a small section of one of the simulations where we can see the same effect but at larger scales; we can still see the web-like structure, but all the haloes are more concentrated after evolving into the future. The left panel shows the same structures as seen today; the right panel shows them in the future when the universe is 100 times larger (note that the units are scaled for the visualization).

We have created over 100 snapshots per simulation and are planning to study the evolution of these structures in much more detail. We plan to extend previous work on simulated future evolution of structures to smaller halos as well as to use different cosmological parameters. Initial results show that the halos stop accreting mass when the universe is twice today’s size. After that point, all structures will be defined and will reach a stable, virialized configuration.

WHY BLUE WATERS

N-body dark matter simulations at scale require the use of a high-performance computing resource that is multicore, scalable, stable, and has high bandwidth with low-latency interconnect communication, especially when running large simulations on the order of 1 billion particles. Parallel I/O access to data is also a requirement. Blue Waters provides all of the above in addition to a very helpful dashboard and outstanding help and technical support, which makes it a perfect platform for running such simulations.

INDUCTIVE E FIELDS IN EARTH'S MAGNETOSPHERE

Allocation: Illinois/60 Knh
PI: Raluca Ilie¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

The terrestrial magnetosphere has the capability to rapidly accelerate charged particles up to high energies over relatively short times and distances, leading to an increase in near-Earth currents. Since the energy of charged particles can be increased only by means of electric fields, knowledge of the relative importance of the potential versus inductive electric fields at intensifying the hot ion population is required to fully understand the dynamics of the inner magnetosphere. However, the contribution of

potential and inductive electric field-driven convection resulting in the development of the storm-time ring current has remained an unresolved question in geospace research. Understanding the implications of the induced electric fields requires a continuous global representation of the electromagnetic fields. This involves a combination of fluid and kinetic approaches that includes all relevant species (ions and electrons), and self-consistent three-dimensional magnetic, convective, and induced electric fields, as well as the relevant loss mechanisms.




Figure 1: The figure shows the percent of the inductive field (color) from the total electric field at a specific time during a geomagnetic storm. The blue sphere represents the inner boundary of the simulation, which is set at 2.5 Earth radii.

RESEARCH CHALLENGE

The overarching goal of this research is to determine what role the electric field induced by the time-varying magnetic field plays in the overall particle energization and to achieve cross-scale knowledge of the particle transport across the magnetosphere system. Thus, we will gain insight into the relationship that controls the global and local dynamical changes in the inner magnetosphere particle population. This new knowledge will allow us to work toward improvements of physics-based space weather prediction models.

METHODS & CODES

Due to the large variability in both temporal and spatial scales of various processes that control different domains inside the terrestrial magnetosphere, physics-based space weather modeling becomes a challenging problem. Therefore, software frameworks are suitable to model multiphysics systems since they offer the possibility of treating each region of the space with an appropriate numerical model. For this project, we use the Space Weather Modeling Framework (SWMF), which is capable of simulating physical processes and coupled domains from the solar surface to the upper atmosphere of the Earth. The SWMF is a robust, high-performance numerical tool for heliophysical simulations. It resolves the disparate spatial and temporal scales with independent spatial and temporal discretizations in the various integrated numerical models. Each domain/component may be represented with one of several alternative physics models, and the coupling of these modules makes the SWMF a unique and powerful tool in simulating the dynamic evolution and interactive physical dependencies within and between coupled systems of the space environment.

RESULTS & IMPACT

Assessing the relative contribution of potential versus inductive electric fields at the energization of the hot ion population in the terrestrial magnetosphere is only possible by thorough examination of the time-varying magnetic field and current systems using global modeling of the entire system. Numerical experiments using our method of separation of the electric field into inductive and potential components, based on Helmholtz vector decomposition of the motional electric field, reveal that the inductive component of the electric field is comparable, and even higher at times than the potential component. This suggests that the electric field induced by the time-varying magnetic field plays a crucial role in the overall particle energization in the inner magnetosphere.

In addition, due to the localized nature of the inductive electric field, knowledge of the relative contribution of potential versus inductive electric fields at intensifying the hot ion population provided new insight into the connection between the macroscale dynamics and microscale processes that govern this region. Further, it solidified our comprehension of the physical processes

controlling the magnetosphere dynamics. The results highlight the importance of accounting for inductive electric fields in space weather prediction models, a component long ignored in the description of near-Earth plasma dynamics. The implications of these findings are immediate as space weather prediction is critical to a forewarning of solar events that could generate severe space weather at Earth.

WHY BLUE WATERS

In order to obtain a physically consistent, realistic, and accurate understanding of plasma transport we use an array of numerical approaches for several disturbed condition intervals, comparing the results against each other and against satellite and ground-based observations. This involves modeling very disparate spatial and temporal scales and, therefore, requires a multiscale modeling approach combined with large computational resources. Blue Waters provided the much-needed platform to run the simulations. The vector decomposition and the calculation of the electric field components by source requires high-accuracy, high-resolution simulations across a very large domain, which is not possible on smaller clusters.

DI

MULTISCALE SPACE WEATHER SIMULATIONS

Allocation: NSF PRAC/3,000 Knh

PI: Gabor Toth¹

Co-PI: Ward Manchester IV¹

Collaborators: Bart van der Holst¹, Yuxi Chen¹, Hongyang Zhou¹

¹University of Michigan

TN

EXECUTIVE SUMMARY

Our efforts aim to achieve a breakthrough advance in the understanding of space weather. The most destructive forms of space weather are caused by 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

Our research solves fundamental problems in plasma physics, solar physics, and magnetospheric physics that relate to magnetic field energization and reconnection. Consequently, the results are significant to solar and plasma scientists, as well as magnetosphere and space weather scientists.

METHODS & CODES

Our approach combines the efficiency of global fluid-type models with the physical capabilities of computationally 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 and 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 12×24 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 about coronal or photospheric conditions.

RESULTS & IMPACT

We have used this unique opportunity to simulate space weather events using the MHD with Embedded Particle-in-Cell (MHD-EPIC) model, where the reconnection is handled by a kinetic PIC code. With this approach, we focused on modeling the fundamental process of reconnection and its impact on global dynamics. Currently, the MHD-EPIC model is the first three-dimensional global study of the complex reconnection process using a high-fidelity kinetic model for the magnetic reconnection. We also made breakthrough advances in simulating flux emergence at an active-region scale in spherical geometry—simulations that will produce eruptions, most notably flares and coronal mass ejections. This research addresses the most salient question of space weather: how the buildup of magnetic energy results in solar eruptions.

WHY BLUE WATERS

Our project uses the Blue Waters petascale computing resource 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 an entire 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-EPIC model and gain a better understanding of the intricate interaction between small kinetic and global scales that result in magnetospheric storms.

PUBLICATIONS & DATA SETS

Toth, G., et al., Scaling the ion inertial length and its implications for modeling reconnection in global simulations. *Journal of Geophysical Research*, 122 (2017), p. 10336.

Chen, Y., et al., Global three-dimensional simulation of Earth's dayside reconnection using a two-way coupled magnetohydrodynamics with embedded particle-in-cell model: initial results. *Journal of Geophysical Research*, 122 (2017), p. 10318.

Manchester, W., et al., Coupled Simulations of Magnetic Flux Emergence. In preparation (2018).




Figure 1: MHD-EPIC simulation of the Earth's magnetosphere. Left: color image represents pressure with projected field lines. The MHD simulation extends to 256Rs; the red box indicates the embedded PIC model ($4R_e \times 6R_e \times 6R_e$). Right: simulated flux transfer event where a magnetic flux rope is formed by magnetic reconnection. Field line colors represent plasma pressure.

MP

COSMIC REIONIZATION ON COMPUTERS

Allocation: NSF PRAC/5,850 Knh
PI: Nickolay Gnedin¹

¹University of Chicago

EXECUTIVE SUMMARY

Cosmic reionization—the process of ionization of the bulk of cosmic gas by ultraviolet radiation from the first galaxies and quasars—is the last frontier of modern cosmology. The Cosmic Reionization on Computers (CROC) project aims to, over the course of several years, produce numerical simulations of reionization. These will self-consistently model all relevant physics from radiative transfer to gas dynamics and star formation in simulation volumes of over 100 comoving megaparsecs (Mpc), necessary to model a representative sample of high-mass galaxies, and with spatial resolution approaching 100 parsecs (pc) in physical units, necessary to reliably model star formation in galaxies. CROC simulations, therefore, cover the full range of spatial, temporal, and mass scales important for studying reionization. The largest CROC simulations can only be completed on Blue Waters-class machines.

RESEARCH CHALLENGE

Study of cosmic reionization has been highlighted by the last Decadal Survey as one of the most promising areas of astrophysical research. Because of the observational constraints on reionization, theoretical modeling, including numerical simulations, plays a relatively larger part in reionization studies than in many other fields of modern astrophysics. While the first simulations of reionization were attempted nearly 20 years ago, major breakthroughs in this field are only possible with modern petascale supercomputing platforms.

Taking advantage of this technological progress, we have initiated a CROC project. This aims, over the course of several years, to produce numerical simulations of reionization that model self-consistently all relevant physics, ranging from radiative transfer to gas dynamics and star formation, in simulation volumes of over 100 comoving Mpc, which is necessary to model a representative sample of high-mass galaxies, and with spatial resolution approaching 100 pc in physical units, which is necessary to reliably model star formation in galaxies. Our simulations, therefore, cover the full range of spatial, temporal, and mass scales important for studying reionization.

The primary motivation for focusing on reionization now is the expected major advance in observational capabilities: The James Webb Space Telescope (the next flagship NASA mission) is expected to be launched in mid-2020, and studying galaxies responsible for cosmic reionization is one of its primary goals. Studies of the intergalactic gas will be propelled forward by the deployment of 30-meter telescopes, several of which will become

operational in the first half of the next decade. Other novel observational tools will follow in the second half of the next decade.

METHODS & CODES

In order to reach the required dynamic range, we rely on the Adaptive Mesh Refinement (AMR) technique. Our simulations are run with the Adaptive Refinement Tree (ART) code, a publicly available cosmological simulation code developed and supported by our group. The code includes all necessary physical modules for simulating cosmic reionization (dynamics of dark matter and gas, atomic processes, interstellar chemistry, star formation and stellar feedback, radiative transfer of ionizing and UV radiation). ART is MPI+OpenMP parallel and scales perfectly (on this type of simulation) to about 50,000 cores, with parallel scaling remaining acceptable to about 100,000 cores.

RESULTS & IMPACT

CROC simulations are defining the state of the art in this field. By virtue of including all the relevant physics and extending to volumes that are required to properly capture the process of reionization, they are creating a physically plausible model of cosmic reionization that can be matched against any existing observational data. Such comparisons have been made by our group in the last several years in a series of papers. However, the most constraining observational data set that exists today is the distribution of optical depth in the spectra of distant quasars. Over 100 quasars during the reionization epoch have been discovered so far, and for almost 70 of them, high-resolution and high-quality spectra exist. These data probe the largest spatial scales relevant for reionization of about 70 comoving Mpc.

Our previous simulations in computational volumes of 60 Mpc were just a bit too small to make a fair comparison with the data. Simulations on Blue Waters reach scales of 120 Mpc and are the only existing numerical simulations of reionization suitable for the comparison with the optical depth data. Such comparison is one of the main science goals for our computational program, to be achieved by the end of the proposal period (summer of 2019).

However, comparison of theory (as realized by numerical simulations) with observational data is not the end in itself. A more important and far-reaching goal of our project is to refine and calibrate theoretical tools for the upcoming observations with JWST, 30-meter class telescopes, and other, more distant, observational advances. The success of the CROC project in matching all of the existing observations demonstrates that such a goal is close to being achieved.

WHY BLUE WATERS

Blue Waters is the only existing U.S. supercomputer where the CROC project can be efficiently completed. ART code does not support GPUs, so Titan is not suitable for our purposes. Prior to the Blue Waters allocation, the CROC project was supported by the U.S. Department of Energy's INCITE program with a more modest allocation on Mira at the Argonne Leadership Computing Facility. However, Mira is significantly slower than Blue Waters, is much less stable, and is being poorly administered, so using Blue Waters is more than 10 times more efficient for our purposes.

PUBLICATIONS & DATA SETS

Villanueva–Domingo, P., N.Y. Gnedin, and O. Mena, Warm Dark Matter and Cosmic Reionization. *The Astrophysical Journal*, 852:2 (2018), p. 139.




Figure 1: Propagation of ionized bubbles (translucent blue color) into the primordial neutral gas (nontransparent dark red). Ionization is driven by early galaxies (yellow dots). The cube is a small region of the whole computational volume, which is 20 times larger in each direction.

USING BLUE WATERS TO UNDERSTAND THE ORIGINS OF GALAXIES AND THE NATURE OF DARK MATTER

Allocation: NSF PRAC/5,060 Knh

PI: Philip Hopkins¹

Collaborators: Claude-Andre Faucher-Giguere², Dusan Keres³, Eliot Quataert⁴, Norman Murray⁵, Mike Boylan-Kolchin⁶, James Bullock⁷, Andrew Wetzel⁸, Robyn Sanderson⁸, Chris Hayward⁹, Robert Feldman¹⁰, Christine Corbett-Moran¹¹

¹California Institute of Technology

²Northwestern University

³University of California, San Diego

⁴University of California, Berkeley

⁵Canadian Institute for Theoretical Astrophysics

⁶University of Texas at Austin

⁷University of California, Irvine

⁸University of Pennsylvania

⁹Flatiron Institute

¹⁰ETH Zurich

¹¹NASA Jet Propulsion Laboratory

EXECUTIVE SUMMARY

Our research program uses Blue Waters to explore the origins of galaxies and stars and the nature of dark matter. At a fundamental level, the study of galaxies and stars seeks to answer the question: “How did we get from the Big Bang to the Milky Way?” This is an immensely challenging question involving the interplay among gravity, fluid dynamics, radiation and matter, and stars exploding as supernovae, giving rise to explosive outflows of material from galaxies that can reach across the observable universe. The physics is chaotic and wildly nonlinear, and the range of scales is tremendous (from one to 10 billion years). As such, our work requires massive numerical simulations that can follow all of these processes. By using these simulations, we have gained fundamental insights into why galaxies today look as they do and, in the process, strongly constrained the allowed nature of the dark matter.




Figure 1: Simulations of a Milky Way-like galaxy. Observed starlight (mock images) is presented with overlaid intensity maps showing mock carbon monoxide (molecular gas), X-rays, and dust emission. At the right is a mock Hubble image during a galaxy collision, where violent bursts of star formation are triggered.

RESEARCH CHALLENGE

Our program seeks to understand the origin and nature of galaxies, using massively parallel simulations that follow the birth and evolution of galaxies and stars from the very early universe to the present day. Our simulations model the origins, evolution, internal structure, and observable properties of galaxies ranging in size from the smallest observed “dwarf” galaxies with just a few thousand stars to the Milky Way and Andromeda (the “Local Group”).

Deep and fundamental questions remain unsolved in this area, including: “How did we get from the Big Bang to the Milky Way? Why did the universe form so few stars compared to its total mass? Why did stars form where and when they did? How can we use galaxies to probe the fundamental nature of dark matter?” At the heart of these issues lies the fact that stars, once they form, are not passive actors within a galaxy. They shine and emit tremendous amounts of energy in the form of light (radiation), stellar winds, and (sometimes) supernovae explosions. This energy can blow material out of the galaxy entirely and completely alter the evolutionary history of galaxies. But these processes remain poorly understood, in large part because they: (1) couple very small and very large scales in the universe, thus requiring simulations with enormous dynamic range to model them; and (2) involve a diverse range of physics including (but not limited to) gravity, fluid dynamics, magnetic fields, conduction and viscosity, radiation-matter interactions, interstellar chemistry, and stellar evolution. The simulations we have conducted incorporated all of these processes into the highest-resolution simulations yet run to address these questions for the first time at the level of detail needed to make observable predictions. Billions of dollars are being invested in new telescopes and instruments to explore these questions experimentally. Simulations will be critical tools

to make detailed predictions and leverage these transformative observations.

METHODS & CODES

We have run a large suite of cosmological, high-resolution simulations including detailed treatments of the physics of the interstellar medium, star formation, feedback in radiation and supernovae, magnetic fields, and cosmic rays. The simulations use the Feedback In Realistic Environments (FIRE) physics methods in the GIZMO code, a new massively parallel multimethod, hybrid Lagrangian-Eulerian finite-element, high-order, radiation-magnetohydrodynamics code (unique in numerical methods and physics-supported).

RESULTS & IMPACT

Our cosmological simulations target galaxies from the faintest dwarfs through the Milky Way at the ultra-high-resolution and realism required to leverage the next generation of observations. The petascale resources of Blue Waters allow us to resolve each galaxy with ~1 billion particles and follow them self-consistently over their entire history in live cosmological settings. When the interstellar medium is resolved into dense molecular clouds, massive stars naturally form and then inject large quantities of energy and momentum into the surrounding medium via “stellar feedback.” This feedback is critical to produce realistic galaxies and generate the powerful galactic winds observed, radically altering the baryon cycle between galaxies and the circumgalactic medium.

These simulations model the physics of galaxy formation with unprecedented realism, uniquely incorporating not only all of the important stellar feedback mechanisms, (radiation pressure, photo-heating, stellar winds, supernovae, and cosmic rays) but also magnetic fields, physical (anisotropic) Braginskii conduction and viscosity, passive scalar (metal) diffusion, and explicit, multiwavelength radiation hydrodynamics. This represents the culmination of several years of work supported by NSF, and has been critical to enable the science of the FIRE project, a collaboration of theorists across 13 different major institutions. The program has revealed fundamental new insights into how stars alter their galactic environments and has changed our observational inferences about the nature of dark matter in those galaxies. The simulations are also being used to support an outreach component involving high school students and teachers, undergraduate students, and a large science team. In addition, this work has been used to make predictions specifically for next-generation observatories including (but not limited to): the James Webb Space Telescope, the Large Synoptic Survey Telescope, Gaia, and the Hubble Space Telescope, in order to test theories of galaxy and star formation, the origin of the heavy elements in the universe, the reionization history of the early universe, the effects of fundamental plasma physics in the circum- and intergalactic medium, and the nature of cold dark matter.

WHY BLUE WATERS

Blue Waters is critical for this research because the enormous computational challenges enumerated above require >100 million CPU-hours on tens of thousands of processors and requiring tens of terabytes of active memory to store and evolve the immensely complex physical systems, and the simulations produce petabytes of data products. No other current facility enables this research.

PUBLICATIONS & DATA SETS

Lamberts, A., Predicting the binary black hole population of the Milky Way with cosmological simulations. *Monthly Notices of the Royal Astronomical Society*, in press (2018), arXiv:1801.03099.

Corbett Moran, C., M. Grudic, and P. Hopkins, The effects of metallicity and cooling physics on fragmentation: implications on direct-collapse black hole formation. *Monthly Notices of the Royal Astronomical Society*, submitted (2018), arXiv:1803.06430.

Hopkins, P., and M. Grudic, Numerical Problems in Coupling Photon Momentum (Radiation Pressure) to Gas. *Monthly Notices of the Royal Astronomical Society*, submitted (2018), arXiv:1803.07573.

Sanderson, R., et al., Reconciling Observed and Simulated Stellar Halo Masses. *Monthly Notices of the Royal Astronomical Society*, in press (2018), arXiv:1712.05808

Garrison-Kimmel, S., et al., The origin of the diverse morphologies and kinematics of Milky Way-mass galaxies in the FIRE-2 simulations. *Monthly Notices of the Royal Astronomical Society*, in press (2018), arXiv:1712.03966.

Hopkins, P., A New Public Release of the GIZMO Code. *arXiv:1712.01294* (2017).



Figure 2: Mock Hubble map of a simulated galaxy, as seen from a Sun-like star. Filamentary molecular cloud complexes and young star clusters are visible within the “Milky Way” (galactic disk). Our combination of physics and resolution allow us to model galactic structure with unprecedented realism.

SIMULATING THE CO-EVOLUTION OF GALAXIES AND THEIR SUPERMASSIVE BLACK HOLES

Allocation: GLCPC/910 Knh

PI: Claude-André Faucher-Giguère¹

Co-PIs: Daniel Anglés-Alcázar², Philip Hopkins³, Eliot Quataert⁴, Dušan Kereš⁵, Robert Feldmann⁶

Collaborator: Blakesley Burkhart³

¹Northwestern University

²Flatiron Institute

³California Institute of Technology

⁴University of California, Berkeley

⁵University of California, San Diego

⁶University of Zurich

EXECUTIVE SUMMARY

We performed high-resolution simulations following the formation and evolution of galaxies from the Big Bang to the present time. We focused on galaxies more massive than the Milky Way, which represent the frontier for high-resolution simulations of galaxy formation and that host the most massive black holes in the universe. We are using the simulations to study how galaxies and the supermassive black holes they host at their centers co-evolve. These new simulations, which pushed the dynamic range of galaxy formation simulations by about an order of magnitude relative to the previous state-of-the-art, will have legacy value in enabling a wide range of other scientific investigations, including studies of the gaseous halos of massive galaxies.

RESEARCH CHALLENGE

Observations indicate a close connection between the growth of supermassive black holes (SMBHs) and the evolution of their host galaxies. For example, the masses of SMBHs correlate tightly

with the stellar bulges of galaxies. The cosmic history of SMBH growth also roughly tracks the cosmic star formation history. Furthermore, SMBHs appear important in explaining the observed properties of the most massive galaxies, which are observed to have stopped forming stars long ago. However, the physical mechanisms driving this connection remain largely unknown. This is because the huge dynamic range of physical scales involved has so far made it extremely challenging to simultaneously simulate all the relevant dynamics. Key questions include: What is the origin of galaxy–black hole scaling relations? How is active galactic nucleus (AGN) activity triggered? What is the impact of AGN feedback on galaxy evolution? What are the effects of stellar feedback on SMBH growth?

METHODS & CODES

To produce our galaxy formation simulations, we used the GIZMO N-body+hydrodynamics code. GIZMO is a publicly available and highly scalable code that implements state-of-the-

art gravity and hydrodynamic solvers. Our simulations also build on a set of modules for galaxy formation physics developed as part of the FIRE (Feedback In Realistic Environments) project.

RESULTS & IMPACT

Scientifically, the new simulations pushed the dynamic range of galaxy formation simulations by about an order of magnitude and will enable a wide range of scientific investigations beyond supermassive black holes. Pushing our simulations to the scale of this allocation also enabled us to test and optimize the scaling of cosmological zoom-in simulations in a new regime. Since a version of GIZMO is publicly available, this optimization will benefit the broader high-performance computing community.

WHY BLUE WATERS

The Blue Waters computational capabilities are essential to simulate at high resolution the massive galaxies in which most supermassive black hole growth occurs, as these simulations require the largest number of resolution elements.

PUBLICATIONS & DATA SETS

Anglés-Alcázar, D., et al., Black Holes on FIRE: Stellar Feedback Limits Early Feeding of Galactic Nuclei. *MNRAS Letters*, 472 (2017), p. L109.




Figure 1: Multiscale rendering of the gas distribution (color scale indicates temperature) in one of the simulations of a massive galaxy with a central supermassive black hole (Anglés-Alcázar, et al., 2017). The leftmost panel shows the large-scale distribution of matter (the dashed circle indicates the size of the dark matter halo) and the rightmost panel focuses on the galactic nucleus.

GEOSCIENCE

CLIMATE

ENVIRONMENT

GEOLOGY

SEISMOLOGY

WEATHER

- 72** *Simulating the Most Devastating Thunderstorms: Big Tornadoes and Big Data*
- 74** *Using a 3D Particle-Resolved Aerosol Model to Quantify and Reduce Uncertainties in Aerosol–Atmosphere Interactions*
- 76** *Analyzing Tropical Cyclone–Climate Interactions Using the High-Resolution Community Earth System Model*
- 78** *The Response of Tropical Cyclone Activity to Global Warming in the Community Earth System Model*
- 80** *Atmospheric Rivers: Potential Drought Busters*
- 82** *Forecasting Global Crop Productivity through Integrating Novel Satellite Data and Process-Based Models*
- 84** *“Breathing” Clouds and Storms: Inflow and Entrainment, Precipitation and Outflow*
- 86** *The Terra Data Fusion Project*
- 88** *Enhanced Digital Elevation Model for the Arctic*
- 90** *Image Processing to Build a Multi-Temporal Vegetation Elevation Ecosystem Model of the Great Lakes Basin*
- 92** *Particulate Matter Prediction and Source Attribution for U.S. Air Quality Management in a Changing World*
- 94** *Direct Numerical Simulations of the Relative Motion of High-Inertia Particles in Isotropic Turbulence*
- 96** *Modeling 4D Earth Evolution: From Continental Cratons to the Yellowstone Supervolcano*
- 98** *Simulating Alternative California Earthquake Histories using Rate and State Friction*
- 100** *Extreme Convective Storms under Climate Change*

SIMULATING THE MOST DEVASTATING THUNDERSTORMS: BIG TORNADOES AND BIG DATA

Allocation: NSF PRAC/3,630 Knh
PI: Leigh Orf¹
Co-PI: Catherine Finley²
Collaborators: Bruce Lee³, Robert Wilhelmson⁴

¹University of Wisconsin-Madison
²St. Louis University
³High Impact Weather Research and Consulting, LLC
⁴University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Tornadoes are among the most violent atmospheric phenomena, creating winds that can exceed 300 miles per hour. Supercells, long-lived rotating thunderstorms that often track across the Central/Midwest/Southeastern United States in the spring and fall, produce the strongest tornadoes. Currently, very little is known about what causes and supports the strongest tornadoes. In our work on Blue Waters, we have simulated several supercell thunderstorms that produce long-lived EF5 (the strongest category on the Enhanced Fujita scale) tornadoes. Data from these simulations, which will take years to fully analyze, have been saved as frequently as every model timestep, allowing for novel postprocessing and analysis approaches. Results thus far indicate features embedded within the storm's cold pool (including the newly identified streamwise vorticity current) help strengthen the storm's updraft very near the ground, supporting the formation and maintenance of long-lived tornadoes.

RESEARCH CHALLENGE

Tornadoes kill many people every year across the world, but the strongest, longest-lived tornadoes occur in the United States, where hundreds of fatalities and serious injuries can result from a single tornado such as in Joplin, Missouri, in 2011, which was also the costliest tornado in U.S. history. Currently, forecasters are unable to predict with any skill these kinds of tornadoes in advance. Our research aims to understand the inner workings of the most powerful tornado-producing thunderstorms in order to better forecast their behavior. The long-term goal of our work is to provide accurate forecasts of these types of storms (as well as less damaging storms) in order to provide ample time for the public to find shelter, and to reduce the high tornado-warning false alarm rate (about 70%) that currently plagues the National Weather Service in the United States.

METHODS & CODES

We used the CM1 cloud model developed at the National Center for Atmospheric Research, a model designed from the ground up to run efficiently on massively parallel distributed memory supercomputers. However, the PI rewrote the I/O driver completely in order to facilitate the frequent saving of

large amounts of model output data. This process resulted in a file system (LOFS) and associated tools to read back and convert data to widely supported formats such as netCDF.

RESULTS & IMPACT

The breakthrough supercell thunderstorm simulations we have conducted on Blue Waters are the first of their kind, containing devastating tornadoes lasting over an hour and a half. The PI has created hours of high-definition visualizations of model data, much of which has been shared on a YouTube channel dedicated to this research.

The identification of the streamwise vorticity current (SVC) in model data has had a big impact in the field of mesoscale




Figure 1: The streamwise vorticity current is comprised of cool air embedded within the storm's forward flank cold pool. It is associated with a drop in pressure 1-2 km above ground that results in a very strong low-level updraft that appears to initiate and maintain a long-lived EF5-strength tornado.




Figure 2: Volume rendered cloud and rain fields for a 15-meter simulation. At this time, two counter-rotating tornadoes descend from the wall cloud beneath the storm's mesocyclone. The storm's cold pool, crucial to the storm's morphology, is indicated at the surface by shades of blue.

meteorology. Efforts are underway to identify the SVC in field studies of supercells. The formation of the tornado, and its subsequent tracking for over an hour and a half, appears to be supported by the SVC and other features in the storm's cold pool along its forward flank. The LOFS file system has enabled us to save data as frequently as every model timestep, and we have saved 90 minutes of data every 1/6 second covering the formation, maintenance, and dissipation of the tornado in one of our simulations. Saving every model timestep to disk opens the door to novel ways of postprocessing and analysis that normally require the model to be running at full scale on a supercomputer. Our approach allows such analysis to be conducted "offline" on much more modest hardware. We are going to use GPU technology to trace millions of particles throughout the saved data and to visualize the data in novel ways. We have recently begun experimenting with creating temporally averaged data from saved 1-second data that reveals the scientifically important persistent forcing features while "smoothing out" the more transient features. The PI uploads "live" video presentations of his research to YouTube and these talks have collectively received tens of thousands of views. Other videos involving the PI and his research have received hundreds of thousands of views.

WHY BLUE WATERS

In order to simulate thunderstorms at resolutions where features such as tornadoes are properly resolved, a machine like Blue Waters is needed. Our most modest simulations (30-meter grid spacing) contain over 1.8 billion individual grid elements to resolve the storm, and simulations using 15-meter grid spacing required more than 18 billion elements. Beyond having the computational power to conduct a given simulation, dozens of simulations were conducted in order to get a feeling for the sensitivity of simulation results to initial conditions and model parameters. Finally, the amount of data produced by these simulations is astounding—each

simulation producing on the order of 50–100 TB, with the highest-resolution simulations creating closer to half a petabyte of data.

PUBLICATIONS & DATA SETS

- Shapiro, A., S. Rahimi, C.K. Potvin, and L. Orf, On the Use of Advection Correction in Trajectory Calculations. *J. Atmos. Sci.*, 72 (2015), pp. 4261–4280.
- 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.
- Li, S., et al., Spatiotemporal Wavelet Compression for Visualization of Scientific Simulation Data. *2017 IEEE International Conference on Cluster Computing (CLUSTER)*, pp. 216–227.
- Frank, L.R., V.L. Galinsky, L. Orf, and J. Wurman, Dynamic Multiscale Modes of Severe Storm Structure Detected in Mobile Doppler Radar Data by Entropy Field Decomposition. *J. Atmos. Sci.*, 75 (2017), pp. 709–730.
- Orf, L., The use of ZFP lossy compression in tornado-resolving thunderstorm simulations. *2017 AGU Fall Meeting* (New Orleans, La., December 11–15, 2017).

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

Allocation: Illinois/240 Knh

PI: Matthew West¹

Co-PI: Nicole Riemer¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

This research aims to reduce key uncertainties in quantifying the impact of atmospheric aerosol particles on Earth's climate. Aerosol particles can be brought into the atmosphere by a wide range of either human activities or natural sources. Aerosols profoundly impact the large-scale dynamics of the atmosphere because they interact with solar radiation—both by scattering and absorbing light, and by forming clouds. These impacts depend on the particles' sizes and their compositions. The uncertainties in quantifying these impacts originate from scale interactions and the high computational cost required for modeling them. To tackle this problem, we developed the particle-resolved 3D model WRF-PartMC, with the unique ability to track size and composition information on a per-particle level. Particle-resolved simulations require efficient numerical algorithms and a computational resource with the capabilities of Blue Waters. Together, they allow for ultra-high-detail simulations needed to quantify the impact of aerosol particles on weather and climate at the regional scale.

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 understanding the global radiation budget via the indirect and direct effects of aerosols [2]. Models provide important insights but experience a trade-off between the representation of physical detail and spatial resolution. Due to computational constraints, current 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 how aerosol particles impact regional and global climate—a topic of great societal relevance.

METHODS & CODES

To overcome the current limitations in representing aerosols and associated uncertainties, we coupled the particle-resolved model PartMC-MOSAIC [3] 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 model uses a 3D Eulerian grid for the atmospheric flow while explicitly resolving the evolution of individual aerosol particles per grid cell.

RESULTS & IMPACT

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 innovation of the WRF-PartMC model consists of representing many of these microscale processes explicitly on a per-particle level, which allows for an improved process-level simulation of the key interactions among aerosols, clouds, and radiation. WRF-PartMC is the only model of its kind, and this work is changing the field of aerosol science because it provides the first benchmark for more approximate models commonly used in the field. It also provides a basis for rigorous coarse-graining to develop physically robust parameterizations for use in larger-scale models. By simulating at a much higher level of detail, particle-resolved models can help close the gap in understanding the effects of modeling choices in global models. Regional-scale particle-resolved simulations allow the quantification of the spatial heterogeneity that determines the conditions where highly detailed aerosol composition is necessary. This next-generation model captures the complex aerosol composition that current-generation models are unable to simulate.

We present results from a particle-resolved aerosol simulation for a realistic, spatially resolved three-dimensional domain in California. Aerosol and trace gas emissions were taken from the 2011 National Emission Inventory [5]. The meteorology corresponded to June 17, 2010, which coincides with the Carbonaceous Aerosol and Radiative Effects Study (CARES) field campaign conducted during May–June 2010. We tracked approximately 50 billion computational particles in this simulation,

including their compositional changes due to gas-to-particle conversion, their coagulation events, and their transport by wind and turbulence. The simulation ran on 81 cores. Most of the compute time was spent on particle coagulation and dynamic gas-particle partitioning on a per-particle basis.

Fig. 1 shows the modeling domain and the spatial distribution of the aerosol number concentrations near the surface after 12 hours of simulation. High-number concentrations are present near major highways since traffic is a major particle source. While aerosol number concentration is a fundamental bulk quantity common to any chemical transport model, the particle-resolved aerosol representation provides unprecedented detail for particle composition and source tracking. Fig. 2 shows an example of the complex continuum of aerosol composition that exists within the single grid cell marked with a star in Fig. 1; 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, creating a high-dimensional composition space that only a particle-resolved model is able to resolve.

As the model tracks composition and source information of thousands of simulated particles per grid cell, individual particles may also be explored. For example, 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 sources such as agriculture and fossil fuel combustion. These capabilities will be useful in future studies for quantifying how much individual source categories contribute to pollution at a certain location.

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 process-level physical representation of box models. Modeling 3D domains with as many as 100 billion tracked particles creates 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 requires tens of thousands of cores, fast interconnections between those cores, and sufficient memory per process.

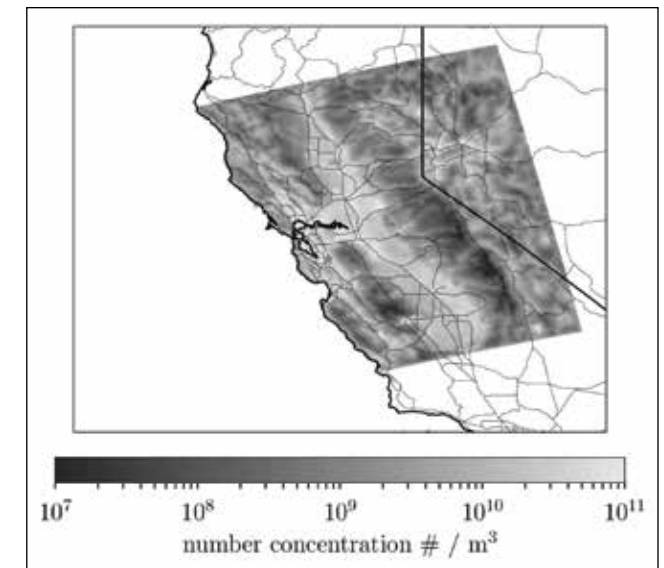


Figure 1: Horizontal distribution of particle number concentration located near the surface after 12 hours of simulation.




Figure 2: WRF-PartMC resolves population-level details. Two-dimensional number distribution as a function of particle dry diameter and black carbon mass fraction indicates the number concentration of particles within a range of diameters and a range of fractions of black carbon mass, for the grid cell denoted by a star in Fig. 1.

ANALYZING TROPICAL CYCLONE-CLIMATE INTERACTIONS USING THE HIGH-RESOLUTION COMMUNITY EARTH SYSTEM MODEL

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

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

This research aims to advance the understanding about the interactions between tropical cyclones (TCs) and climate using the high-resolution Community Earth System Model (CESM). The high-resolution CESM can generate TCs within the modeled climate, providing a valuable platform for TC-climate research. In this work, we performed a suite of CESM simulations using various model configurations including fully coupled, atmosphere-only, and ocean-only. We used the fully coupled simulation as the control and performed an atmosphere-only simulation to examine the sensitivity of modeled TCs to ocean coupling, concluding that ocean coupling is essential to capturing realistic TC intensity and intensification. We then investigated the impact of TCs on the global ocean using a set of ocean-only simulations with and without TC forcing. We found that TCs can influence global ocean temperature patterns, ocean energetics, and ocean heat transport. These results help provide insights into the model behavior and the physical nature of the connection between TCs and climate.

RESEARCH CHALLENGE

Tropical cyclones (TCs) are one of the world's most destructive natural hazards. The 2017 hurricane season caused over \$200 billion of damage. TC activity is closely linked to large-scale climate conditions [1], and the question of how TCs will change with the changing climate is drawing more and more attention. In addition, research has shown that TCs have the potential to actively contribute to the Earth's climate system through their interactions with the upper ocean [2,3]. Understanding these processes and feedbacks is of great importance to a better knowledge of the Earth's climate system and constraining uncertainty in climate projections.

The current generation of high-resolution Atmosphere General Circulation Models (AGCMs) with horizontal resolution finer than 25 km can generate realistic global TC activity and resolve the most intense TCs within the modeled climate [4,5], providing a unique platform for TC-climate research. The goals of this project are: (1) to characterize the model's self-generated TCs and analyze the sensitivity of the simulated TC characteristics to ocean coupling; and (2) to diagnose the impact of the model's self-generated TCs on the global ocean within the modeled climate.

METHODS & CODES

In this project, we analyzed the interactions between TCs and climate using various configurations of the high-resolution Community Earth System Model (CESM) [6]. We performed a 30-year fully coupled high-resolution simulation (CPL) in which the 0.25° (~25-km) resolution atmosphere component is coupled to the 1° (~110 km) ocean component. To analyze the sensitivity of the simulated TC activity to ocean coupling, we perform an atmosphere-only simulation (ATM) in which the atmospheric component is forced by sea surface temperature from the 30-year CPL. The impact of ocean coupling can then be diagnosed by comparing the simulation results between CPL and ATM. To analyze the impact of TCs on the global ocean, we performed two ocean-only simulations using the boundary conditions of the 30-year CPL in which TC effects at the ocean-atmosphere boundary were filtered out in one of the ocean simulations (OCN_FILT) while fully retained in the other (OCN_TC) in order to isolate the effect of the TCs on regional and global ocean variability across multiple time scales.




Figure 1: Thirty years of accumulated tropical cyclone (TC) tracks in the (a) observational best track during 1985–2014, (b) the coupled CESM simulation, and (c) the atmosphere-only simulation. (Right) Relationship between SST and maximum TC wind speed in the fully coupled simulation (blue) and the atmosphere-only simulation (salmon).

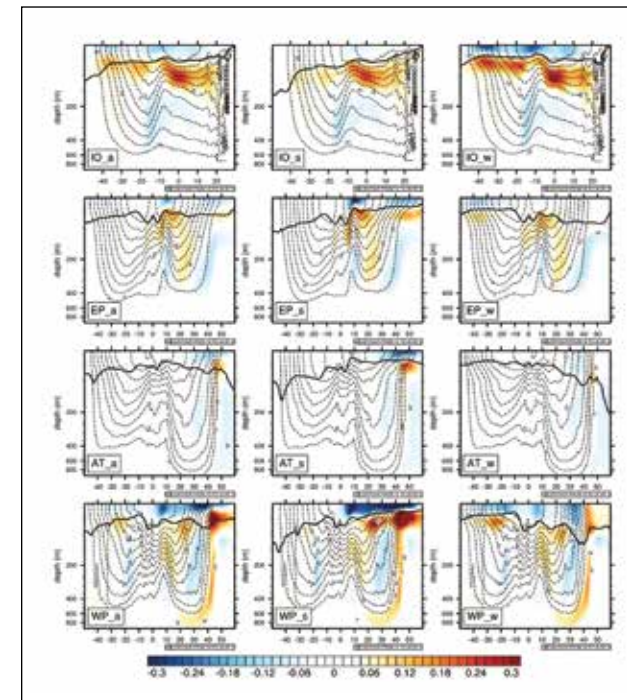


Figure 2: Meridional cross-section of long-term mean subsurface temperature differences between the OCN_TC run and OCN_FILT in the (top to bottom) Indian Ocean, eastern Pacific, Atlantic, and western Pacific for (left to right) annual, summer (June–November), and winter (December–May).

RESULTS & IMPACT

Results from the ATM simulation reveal that ocean coupling can influence the simulated annual TC number, spatial distribution, and storm intensity (Fig. 1). Despite using the same SST boundary forcing, ATM simulates more TCs globally than CPL, particularly for intense TCs stronger than Category 3. The differences are mainly attributable to the active air–sea coupling under TCs that accounts for the ocean's feedback: TC-induced SST cooling can restrain TC intensification, and this negative feedback mechanism can only be captured in a fully coupled model configuration.

For the ocean-only simulations, we found that the modeled TCs could induce changes in surface and subsurface ocean temperature patterns and seasonality (Fig. 2). On annual average, we observed SST cooling (up to 0.5°C) in major TC basins and SST warming (~0.1°C) at the higher latitudes. The surface warming occurs in the winter season, which is due to the reemergence of subsurface heat anomalies associated with the seasonal variations of the mixed-layer depth. This wintertime SST warming results in an anomalous SST gradient at the location critical for the midlatitude winter storm track. At the ocean subsurface, TC-induced heat anomalies generally travel along the isotherms following the ocean overturning cells, with a tendency to converge back to the near-equator upwelling zone. Moreover, we found that TCs can strengthen ocean gyre circulations and meridional overturning circulations, though the changes are small. The

maximum overturning in the subtropical cell was enhanced by 3% on global average and 1% in the Atlantic. In addition, we found that the modeled TCs can influence ocean meridional heat transport (MHT). In general, TCs amplify the background MHT patterns, featuring enhanced poleward heat transport and cross-equator northward transport. However, the contribution of TC-induced anomalies to the background total MHT is not significant, accounting for ~2.5% of the background peak transport.

These results help to better understand the relationship between TCs and climate within the state-of-the-art TC-resolving Earth System Model. Results reveal how the choice of model configuration can influence the model's self-generated TC activity and how the model's self-generated TCs can influence the model-simulated global ocean in various aspects and across different time-scales. This has provided insights into the effect of resolved transient extreme weather events on the simulated mean climate within the high-resolution Earth System Models.

WHY BLUE WATERS

TC-climate research falls at the interface between weather and climate modeling, requiring high-resolution grid spacing to resolve weather-scale TC features, as well as global-scale coverage and decades of integration time. Blue Waters has the unique capability 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 & 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.

Li, H., and R.L. Sriver, Tropical cyclone activity in the high-resolution community earth system model and the impact of ocean coupling. *Journal of Advances in Modeling Earth Systems*, 10 (2018), DOI:10.1002/2017MS001199.

Li, H., and R.L. Sriver, Impact of Tropical Cyclones on the Global Ocean: Results from Multi-decadal Global Ocean Simulations Isolating Tropical Cyclone Forcing. *Journal of Climate*, in review (2018).

THE RESPONSE OF TROPICAL CYCLONE ACTIVITY TO GLOBAL WARMING IN THE COMMUNITY EARTH SYSTEM MODEL

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

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

This research aims to advance our understanding about the interactions between tropical cyclones (TCs) and climate using the high-resolution Community Earth System Model (CESM) with 0.25° horizontal resolution in the atmosphere and 1° horizontal resolution in the ocean. The high-resolution CESM can generate TCs within the modeled climate, providing a valuable platform for TC climate research. This project builds on the PIs' recent Blue Waters allocation, analyzing tropical cyclones in high-resolution configurations of CESM under constant present-day climate conditions for different coupled and uncoupled configurations. Here, we expand these simulations into a more comprehensive climate change experiment, investigating the response of global TC activity to changing greenhouse gas forcing (e.g., atmospheric carbon dioxide concentrations) using idealized 4×CO₂ experiments. Results provide insights into the model behavior and potential response in global TC activity to changes in anthropogenic global warming.

RESEARCH CHALLENGE

Tropical cyclones are rare weather events, yet they consistently rank among the world's deadliest and costliest natural hazards.

The frequency, intensity, and spatial distribution of TCs can vary with climate. In addition, recent research suggests TCs actively contribute to the dynamics of Earth's climate system through complex ocean-atmosphere interactions that can alter ocean temperature patterns and influence circulations within the atmosphere and ocean. These cyclone-climate connections are poorly understood and largely missing from today's generation of Earth system models, yet they may be fundamentally important to understanding the mechanisms influencing climate variability and to improving projections of future climate change.

METHODS & CODES

The Community Earth System Model is a comprehensive global climate model that consists of atmosphere, land, ocean, and sea ice components that are connected via a central coupler that exchanges state information and fluxes between the components [1,2]. It is used by scientists and researchers worldwide from universities, national laboratories, and other institutions for climate, weather, atmospheric, oceanic, and land surface research. It represents the leading edge of communitywide efforts in global climate modeling, and it is considered a state-of-the-art Earth system model.

The atmosphere component of CESM is the Community Atmosphere Model (CAM5.3) with the spectral element (SE) dynamic core [3]. CAM5 has improved microphysics and cloud properties, and the prognostic modal aerosol package is activated in the current experiments. The SE version of the CAM model has a higher scalability than the finite volume version, making it suitable for high-resolution implementation on parallel computing clusters. The atmosphere is configured with 0.25° (~25 km) horizontal resolution and the standard 30 vertical layers. The land model is the Community Land Model version 4 and uses the same grid resolution as the atmosphere.

The dynamic ocean component of CESM is the Parallel Ocean Program version 2 (POP2), and the ice component is the Los Alamos Sea Ice Model. POP2 solves the three-dimensional primitive equations with the Boussinesq and hydrostatic approximations. We use the nominal 1° horizontal resolution on the displaced pole grid, which has a uniform zonal grid spacing of 1.125° and a varying meridional grid spacing that increases from 0.27° at the equator to a maximum grid spacing of 0.64° (northwestern Pacific). The model has 62 vertical levels, with layer thickness increasing from 10 m at the near-surface to 250 m at a depth of 6,000 m.

RESULTS & IMPACT

In this research project, we are building on our recent work analyzing the relationship between TCs and climate using a high-resolution, state-of-the-art CESM. We are introducing time-varying changes in CO₂ through an idealized ramped forcing experiment, in combination with the extension of the control simulations with constant pre-industrial forcing. We are analyzing potential sensitivities in cyclone activity and damage metrics to CO₂-induced changes in large-scale environmental factors, and we are assessing potential biases based on model-data comparisons for current climate conditions. These simulations are also being used to explore the potential for cyclone-induced climate feedbacks within the coupled Earth system.

These runs are orthogonal to our ongoing Blue Waters CESM simulations, in which we are testing the sensitivity of TCs to ocean-atmosphere coupling (Figs. 1 and 2). The preliminary model sensitivity results are very promising, and we are now conducting the climate change component of the experiment. The new simulations will provide enough years to assess statistically significant changes in regional-to-global TC activity under increased carbon dioxide scenarios in CESM.

The comprehensive suite of model experiments have provided us with thousands of simulated cyclone tracks for different climate conditions and coupling configurations using a dynamically consistent Earth system modeling framework, enabling robust assessment of cyclone activity and variability in response to changes in climate. A key outcome of these simulations is gridded fields of cyclone tracks, winds, and other useful cyclone metrics that can be used to assess potential changes in cyclone-related

damages under future global warming. These data products and model outputs are freely available upon request.

WHY BLUE WATERS

TC climate research falls at the interface between weather and climate modeling, requiring high-resolution grid spacing to resolve weather-scale TC features, as well as global-scale coverage and decades of integration time. Blue Waters provides the unique capabilities of handling 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 postprocessing and visualization of model results.

Understanding the physical relationship between tropical cyclones and climate, and assessing the societal and economic impacts of climate change, represents a grand challenge to the Earth system modeling community. Blue Waters provides the unique capability to solve this problem.

PUBLICATIONS & DATA SETS

Bock, D., H. Li, and R.L. Sriver, Simulation and visual representation of tropical cyclone-ocean interactions. *International Conference for High Performance Computing, Networking, Storage and Analysis (SC17)* (Denver, Colo., 2017).

Li, H., and R.L. Sriver, Tropical cyclone activity in the high-resolution community earth system model and the impact of ocean coupling. *Journal of Advances in Modeling Earth Systems*, 10 (2018), DOI:10.1002/2017MS001199.

Li, H., and R.L. Sriver, Impact of Tropical Cyclones on the Global Ocean: Results from Multi-decadal Global Ocean Simulations Isolating Tropical Cyclone Forcing. *Journal of Climate*, in review (2018).

Li, H., and R.L. Sriver, Impact of air-sea coupling on the simulated global tropical cyclone activity in the high-resolution Community Earth System Model (CESM). *Journal of Geophysical Research-Atmospheres*, in preparation (2018).




Figure 1: Box-whisker plot of global tropical cyclone number in different basins for observations (black), coupled CESM (steel blue), and atmosphere-only CESM simulations (salmon). Basin abbreviations: AT (Atlantic), EP (Eastern Pacific), NI (Northern Indian), WP (Western Pacific), and SH (Southern Hemisphere).




Figure 2: Binned frequency distribution of global annual average number of tropical cyclones for different maximum wind speeds from observations (grey), coupled CESM (steel blue), and atmosphere-only CESM (salmon).

ATMOSPHERIC RIVERS: POTENTIAL DROUGHT BUSTERS

Allocation: NSF PRAC/9,000 Knh
PI: Ryan Sriver¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Atmospheric rivers (ARs) are a subset of midlatitude storms that can have profound regional impact. This is particularly true in the western United States where ARs can make landfall and produce heavy rainfall events. Strong AR precipitation during the anomalously wet winter of 2016–2017 nearly eliminated the decade-long drought impacting most of California. In this project on Blue Waters, we use fully coupled high-resolution atmospheric and oceanic climate simulations to investigate how ARs on the West Coasts of the United States, United Kingdom, and Iberian Peninsula may change under a future high-emission warming scenario. Future change in AR behavior is linked to changes in the midlatitude jet. Our model results suggest more or stronger ARs on the California coast, the U.K., and Iberian Peninsula and fewer or weaker ARs in the Pacific Northwest. Precipitation associated with ARs is also projected to increase in intensity under global warming conditions.

RESEARCH CHALLENGE

An overarching objective of our research is to identify processes and mechanisms that characterize high-impact events and quantify how these events could change in the future. For this particular study, we focused on atmospheric rivers (ARs) that make landfall on the U.S. West Coast, the U.K., and Iberian Peninsula. Results indicate that AR properties are dependent on the location and strength of the subtropical and eddy-driven atmospheric jets [1] and that ARs using higher-resolution atmospheric models better match those found in observations [3]. In this study, we investigate

the representation of ARs and their potential future change using a climate model configuration with a high-resolution atmosphere (0.25°) uncoupled and coupled to both a medium- (1°) and high- (0.1°) resolution ocean.

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, and river transport, the CESM allows researchers to conduct fundamental research into the Earth's past, present, and future climate states. The experiments we conducted are at the highest resolution currently feasible within the CESM for century-scale climate studies in terms of both scientific fidelity and computational cost. They allow better quantification of long-term climate change projections at regional scales and enable a more accurate representation of important dynamical features. These features include realistic extreme precipitation events that compare better to observations in heavy precipitation categories as well as more realistic midlatitude and tropical storms. In this project, we investigated ARs—a subset of midlatitude storms. (The long, filamentous structures streaming from the tropics into the midlatitudes in both hemispheres in Fig. 1 are illustrations of ARs.) A tracking algorithm developed by Shields and Kiehl [2] was employed to obtain the number and properties of North Pacific (U.S. West Coast) and North Atlantic (United Kingdom and Iberian Peninsula) ARs.




Figure 1: A snapshot of total precipitable water in the atmosphere on August 30, 1990, from a high-resolution (0.25°) simulation of the CESM. This illustrates the location of moisture in the atmosphere. (Figure courtesy of Computational and Information Systems Lab/NCAR)




Figure 2: Snapshots of examples of ARs in the north Atlantic (bottom panels) and north Pacific (top panels) when using a standard 1° ocean configuration (right panels) and a high-resolution (0.1°) configuration (left panels). The high-resolution ocean provides significantly more detailed structure to the AR. (Figure courtesy of Christine Shields, NCAR)

RESULTS & IMPACT

We first compared the representation of ARs within the various model configurations using a previous study conducted by Shields and Kiehl [1,2] that examined simulations with a 0.5° atmosphere model and the simulations run on Blue Waters that use a 0.25° atmosphere. Results indicate that increasing resolution generally improves the AR count compared to observations. Increasing model resolution results in a decrease of AR counts along the California and U.K./Iberian Peninsula coasts whereas counts increase in the Pacific Northwest. We anticipate further improvement in the simulation using the 0.1° ocean, particularly in the North Atlantic where the latitudinal position of the underlying Gulf Stream leads to improvement in storm track location [4].

We next investigated changes in ARs that could occur by the mid-21st century under the future high-emission scenario (RCP8.5). Change in ARs are driven by changes in the subtropical and eddy-driven jets, with the U.S. West Coast and Iberian Peninsula ARs following the subtropical jet and the U.K. ARs following changes in both jets. In the North Pacific in the future scenario we anticipate a strengthening of the jet along its southern flank and a weakening along the northern flank, resulting in more or stronger ARs on the California Coast and fewer or weaker ARs in the Pacific Northwest. In the North Atlantic, strengthening of both jets results in more or stronger ARs for both the U.K. and the Iberian Peninsula.

AR research thus far has largely focused on the atmosphere [e.g., 1,2]. Our work investigated the ocean's role regarding AR development, intensity, and structure, and how this may change in future states of the earth system. High-resolution ocean models may resolve AR structure with much finer detail, in part by resolving the small-scale ocean eddies that sit below the AR transport path. More realistic ocean eddy structure could impact

AR structure and its role as a vehicle for moisture transport. Fig. 2 illustrates the ability of the higher-resolution ocean to imprint smaller-scale features onto the structure of ARs compared to the lower-resolution model.

WHY BLUE WATERS

In order to determine model representation of climate processes and to project future changes, multiple century-long simulations are needed, including a long, stable preindustrial control, followed by historical and future scenarios. For investigation of climate extremes such as ARs, high-resolution simulations are necessary. Furthermore, multimember ensembles are needed to quantify and reduce uncertainty. The simulations we have conducted on Blue Waters this past year consist of a 0.25° atmosphere/land coupled to both a 1° and 0.1° ocean. These high-resolution simulations 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 needs the help of the Blue Waters staff to achieve good throughput.

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

Allocation: Blue Waters Professor/240 Knh
PI: Kaiyu Guan¹
Co-PIs: Jian Peng¹, Bin Peng¹, Yunan Luo¹, Sibao Wang¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

The ultimate goal of our project is to improve the predictability of regional and global crop yields by integrating advanced remote-sensing and process-based modeling. We have developed the CLM-APSIM modeling framework, which combines the strengths of earth system models and agronomy crop models. We conducted parameter sensitivity analysis and spatially explicit parameter optimization using satellite-based constraints and ran a set of historical simulation and future projection experiments. For remote sensing, we developed a multisource satellite data fusion algorithm (STAIR, or SaTellite DAte IntegRation) that can generate daily cloud/gap-free high-resolution surface reflectance images. We have also developed a machine-learning framework that ingests a time series of high-resolution satellite data for field-level in-season crop-type identification. The combined advancements in crop modeling and remote sensing will allow us to develop a seasonal forecasting system for crop productivity in the U.S. Corn Belt.

RESEARCH CHALLENGE

Global food security is under continuing pressure from increased population and climate change [5–8]. Crop productivity forecasting offers a promising technical framework for food-related early warning and decision-making. Crop models, either statistical or process-based, are the most essential component in any crop productivity forecasting system. Compared to statistical models, process-based crop models are better tools for studying the impacts of historical and future climate on regional and global food production, for assessing the effectiveness of possible adaptations and their potential feedback to climate, and for attributing different pathways through which climate can impact crop yields. Nevertheless, current process-based crop models still have large uncertainties when considering the demands from a forecasting perspective. Moreover, an in-season field-level crop-type classification product is also missing and remains a challenge for seasonal forecasting of crop productivity.

METHODS & CODES

We are targeting better prediction performance for crop yield in the U.S. Corn Belt through combining advanced process-based modeling and remote-sensing observations. We first improved

the crop growth representation in one of the leading Earth System Models, the Community Earth System Model (CESM) [2]. We combined the strengths of the Community Land Model (CLM) [4], the land component of CESM, and the state-of-the-art agronomy crop model APSIM (Agricultural Production Systems sIMulator). We coupled the SALib Python package [1] with the CESM by using its multi-instance configuration for parameter sensitivity analysis, and coupled the PyDREAM Python package [3] with the CESM for Bayesian parameter calibration of the CLM-APSIM model.

We are also developing a remote-sensing monitoring system for U.S. croplands, which includes two major components: the satellite data preprocessing and fusion component, and the machine learning-based classifier component. The data preprocessing component is responsible for harmonizing the multisource satellite images, including data encoding, geo-projection, geo-registration, and quality control of the images. The multisource satellite data fusion component uses our newly developed STAIR algorithm to generate cloud/gap-free high-resolution (daily and <=30m) surface reflectance images. For the machine learning-based classifier component, we have implemented a machine-learning framework that can ingest time series high-resolution satellite images to provide in-season crop-type information. All of these workflows were developed on Blue Waters.

RESULTS & IMPACT

The CLM-APSIM combines the strengths of both the CLM version 4.5 and the APSIM models (Peng, et al., 2018; Fig. 1). An evaluation of results at the AmeriFlux sites located in the U.S. Corn Belt show that the CLM-APSIM model performs better than the original CLM4.5 in simulating phenology, surface fluxes, and especially biomass partition and maize yield. The CLM-APSIM model corrects a serious deficiency in CLM4.5 related to underestimation of aboveground biomass (i.e., overestimation of belowground biomass) and overestimation of the Harvest Index, which leads to a reasonable yield estimation with the wrong mechanisms. We are conducting parameter sensitivity analysis and spatially explicit parameter optimization using satellite-based constraints for the CLM-APSIM model. We are also running a set of historical simulation and future projection experiments aimed at disentangling the contribution of different mechanisms to high-temperature impact on crop yield.

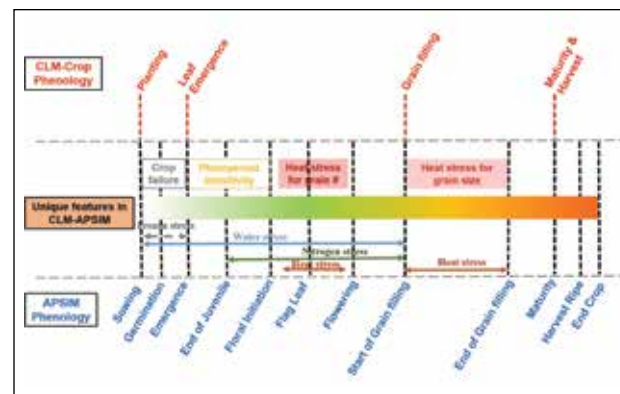


Figure 1: Conceptual diagram for phenological stages in the CLM, APSIM, and CLM-APSIM models. Unique features in the CLM-APSIM crop model are highlighted. Note that the stage duration in this diagram is not proportional to real stage length and is only presented for illustrative purpose.

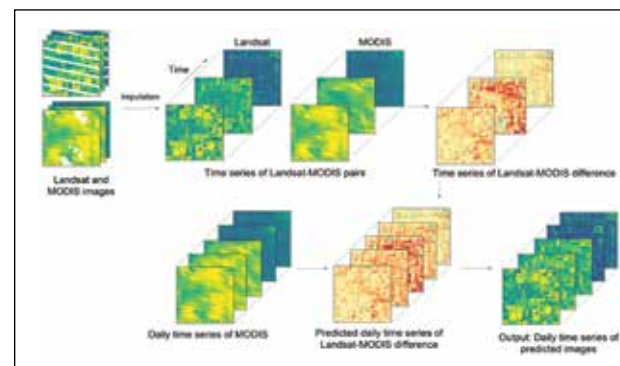


Figure 2: Schematic overview of the fusion method of SaTellite DAte IntegRation (STAIR).

The STAIR algorithm is a novel, fully automatic method for fusing multispectral satellite data to generate high-frequency, high-resolution and cloud/gap-free data (Luo, et al., 2018; Fig. 2). Building on the time series of multiple sources of satellite data, STAIR first imputes the missing-value pixels due to cloud cover or sensor mechanical issues in satellite images using an adaptive-average correction process. This takes into account different land covers and neighborhood information of missing-value pixels through an automatic segmentation. To fuse satellite images, it employs a local interpolation model to capture the most informative spatial information provided by the high-spatial-resolution data (e.g., Landsat) and then performs an adjustment step to incorporate the temporal patterns provided by the high-frequency data (e.g., MODIS). We tested our algorithm to fuse surface reflectance data of MODIS and Landsat for Champaign County, Illinois, and generated daily time series for all the growing seasons (April 1 to November 1) from 2000 to 2015 at 30-m resolution. Extensive experiments demonstrated that STAIR not only captures correct texture patterns but also predicts accurate reflectance values in the generated images, with a significant performance improvement over the classic STARFM algorithm.

This method is computationally efficient and ready to be scaled up to continental scales. It is also sufficiently generic to easily include various optical satellite data for fusion. We envision this novel algorithm can provide the effective means to leverage historical optical satellite data to build long-term daily, 30-m surface reflectance records (e.g., from 2000 to the present) at continental scales for various applications. In addition, it will produce operational near-real-time daily and high-resolution data for future earth observation applications.

The machine learning-based crop-type and area-mapping system has also been tested for Champaign County, Illinois (Cai, et al., 2018). Specifically, we used the U.S. Department of Agriculture’s (USDA) Common Land Units (CLUs) to aggregate spectral information for each field based on a time-series Landsat image data stack. This largely overcomes the cloud contamination issue while exploiting a machine-learning model based on Deep Neural Network (DNN) and high-performance computing for intelligent and scalable computation of classification processes. We used a total of 1,322 Landsat multitemporal scenes including all the spectral bands from 2000 to 2015. Our results show the inclusion of temporal phenology information and evenly distributed spatial training samples in the study domain improve classification performance. Benchmarked with USDA’s Crop Data Layer, our algorithm shows a relatively high overall accuracy of 96% for classifying corn and soybeans across all CLU fields in Champaign County from 2000 to 2015. Furthermore, our approach achieved 95% overall accuracy for classifying corn and soybeans by late July of each year.

WHY BLUE WATERS

Blue Waters is essential for our research since other resources, such as those available from XSEDE, are not suitable for our project considering the petabyte-level storage demand, data availability, and intensive I/O and computational demands.

PUBLICATIONS & DATA SETS

Peng, B., et al., Improving maize growth processes in the community land model: Implementation and evaluation. *Agricultural and Forest Meteorology*, 250–251 (2018), pp. 64–89.

Luo, Y., K. Guan, and J. Peng, STAIR: A generic and fully-automated method to fuse multiple sources of optical satellite data to generate a high-resolution, daily and cloud-/gap-free surface reflectance product. *Remote Sensing of Environment*, 214 (2018), pp. 87–99.

Cai, Y., et al., A high-performance and in-season classification system of field-level crop types using time-series Landsat data and a machine learning approach. *Remote Sensing of Environment*, 210 (2018), pp. 35–47.

Guan, K., et al., The shared and unique values of optical, fluorescence, thermal and microwave satellite data for estimating large-scale crop yields. *Remote Sensing of Environment*, 199 (2017), pp. 333–349.

“BREATHING” CLOUDS AND STORMS: INFLOW AND ENTRAINMENT, PRECIPITATION AND OUTFLOW

Allocation: Blue Waters Professor/275 Knh
PI: Sonia Lasher-Trapp¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

“Entrainment” describes how clouds/storms bring dry air from outside the cloud inward by their own turbulent motions. Its effects can limit storm development, longevity, and precipitation. Our understanding of entrainment has been limited by inadequate model resolution in past studies. Now, we are running high-resolution 3D storm simulations on Blue Waters to quantify the interactions among entrainment, precipitation, and the generation of new storms.

Our latest results suggest that: (1) Storms growing in an environment where the winds increase strongly with height entrain three to five times more dry air than storms growing in a less windy environment; (2) storms that develop closer together may precipitate less initially but can produce outflows (strong winds) capable of forcing new storms that precipitate much more; and (3) the strength of these outflows appears to be most dependent upon the amount of large ice particles, called “graupel,” that fall from the storm.

RESEARCH CHALLENGE

Deep convective clouds produce the majority of the earth’s precipitation, and yet it is difficult to predict if developing cumulus clouds will attain the depth and longevity required to produce heavy rainfall that can generate flooding in some cases. In addition, strong winds (outflows) from some storms are capable of forcing new storms, leading to thunderstorm outbreaks that can affect large geographical areas. “Entrainment” is the term for the process by which the turbulent motions within clouds bring dry air from outside the cloud inward. In time, entrainment can limit storm development and precipitation but is not always effective in doing so. Long-standing problems in meteorological models have been to understand why they tend to predict rain formation too early, and/or in excessive amounts, and why models often miss predicting outbreaks of storms. Our approach here is to investigate if these deficiencies in past models are related to poor representation of entrainment, and of the types of precipitation that produce the strongest storm outflows that may generate new storms. 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 used NCAR’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 utilized the NSSL microphysical scheme [2] within CM1 to model the details of precipitation formation. We conducted these simulations in both idealized and realistic environments. We evaluated entrainment with our own code [3] that calculates mass fluxes into the core of the storm as it evolves, and related it to the amount of precipitation the storm produces. We also used our own code that calculates the most important melting/evaporation factors that can strengthen storm outflows and potentially produce outbreaks of thunderstorms.

RESULTS & IMPACT

Our latest results are providing new quantitative information that can assist atmospheric scientists researching ways to represent thunderstorm entrainment, precipitation, and outflows in larger-scale weather and climate prediction models. They will also be of use to weather forecasters. Our results suggest that:

- Developing thunderstorms growing in an environment where the winds increase strongly with height (Fig. 1) initially entrain three to five times more dry air than storms developing in a less windy environment. These quantitative guidelines help establish how thunderstorm initiation can be suppressed in very windy environments.
- Storms that develop closer together may precipitate less initially due to their interference with the inflow of air into the neighboring storm bases that fuel them. However, if they do eventually precipitate, their outflows (strong winds) can meet and force new storms that precipitate much more.
- Calculations of the contribution of different kinds of precipitation to the strength of the storm outflows (Fig. 2) have revealed that the amount of large ice particles that fall from the storm appears to be most important rather than the evaporation of rain, as is sometimes assumed. This suggests that certain aspects of precipitation production in storms may be critical to thunderstorm outbreaks.

WHY BLUE WATERS

Our Blue Waters allocation has been essential for achieving the high resolution required within a given simulation to properly represent the smaller cloud motions that are important for entrainment and precipitation development over the larger spatial and temporal domains required for thunderstorms and their outflows. As a result of its huge number of nodes, its high speed, large memory, and its large storage capability for high-resolution model output and analysis, Blue Waters enables us to conduct our detailed calculations over millions of grid points. The hardware needed to run these kinds of simulations quickly supersedes the limits of most computers.




Figure 1: Precipitation mass in a thunderstorm developing in an environment where the wind speed increases greatly with height. Much of the mass from the older part of the cloud (blue colors on right side) was depleted by entrainment.

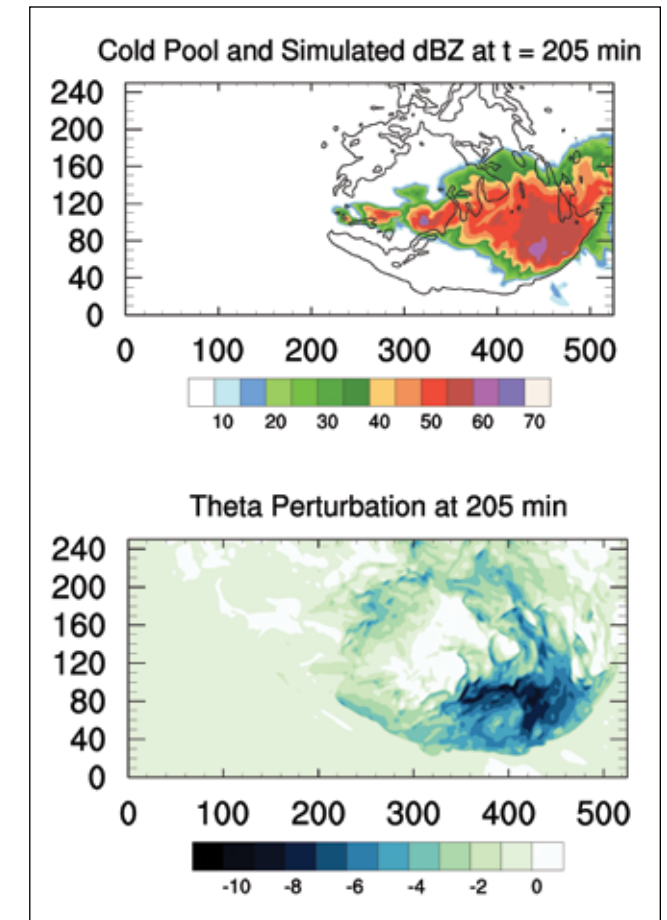


Figure 2: (Top) Simulated radar reflectivity for a system of thunderstorms, with cold outflow outlined in black. (Bottom) Associated temperature deficit (in degrees C) due to the cold precipitation outflow at this time. As this outflow propagates away from the storm, it can force new thunderstorms at its edge.

THE TERRA DATA FUSION PROJECT

Allocation: Blue Waters Professor/210 Knh

PI: Larry Di Girolamo¹

Co-PIs: John Towns², Shaowen Wang¹, Guangyu Zhao¹, Kent Yang³, Yan Liu¹

¹University of Illinois at Urbana-Champaign

²National Center for Supercomputing Applications

³The HDF Group

EXECUTIVE SUMMARY

We have led the Terra Data Fusion Project to success through collaborative efforts among NASA, the HDF Group, and NCSA. We generated and validated missionwide processed and calibrated (Level 1B [1]) fusion products (2.4 PB) on Blue Waters, which provides the necessary stepping stone for developing higher-level products and provides the framework for other flavors of fusion. For these data, we built an open-source tool that resamples Terra satellite data into a common grid adopted by any instrument or defined by any scalable map projection. The fusion data set has been further used to: (1) characterize ice crystal roughness of cirrus clouds, resulting in a better understanding of global ice cloud optical properties; (2) quantify regional biases in the retrieved cloud drop sizes of liquid water clouds induced from cloud heterogeneity; and (3) examine decadal changes in the Earth's radiance fields, revealing temporal and spatial variability at multiple scales.

RESEARCH CHALLENGE

The Terra satellite was launched in 1999 and continues to collect Earth science 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 data sets, 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 [2]. The challenge is particularly acute for Terra, given its growing data volume (>1 petabyte), the storage of different instrument data at different NASA centers, the different data file formats and projections, and inadequate cyberinfrastructure [3]. We recently initiated the Terra Data Fusion Project to tackle two long-standing problems: (1) How to efficiently generate and deliver Terra data fusion products; and (2) How to facilitate the use of Terra data fusion products by the community in generating new products and knowledge through national computing facilities, and disseminate these new products and knowledge through national data sharing services.

Solutions to these problems 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) reduce error and redundancy in the science community among those researchers using multiple instrument data sets; (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




Figure 1: Image frames captured from an animation clip that dynamically displays and projects the radiance imageries generated from one single Terra basic fusion granule (Orbit 3671) onto Earth as being orbited by Terra for all of the five Terra instruments. See it on YouTube at <https://youtu.be/C2uyjRGwwOs>.

NASA services. Thus far, we have accomplished key steps (1), (2), and (3), and we are working with NASA on step (4).

We also developed an open source tool to resample or reproject the radiance fields stored in the Terra basic fusion product into a common grid adopted either by a Terra instrument or defined by any map projections. We built two additional tools specifically to facilitate large data processing on Blue Waters: a Globus Python interface data transfer tool, which interacts with the Globus CLI, and a Python PBS workflow manager, which is a lightweight module with a friendly interface that provides a programmatic way to define complex job chains.

RESULTS & IMPACT

We successfully transferred mission-scale radiance data from all five Terra instruments to Blue Waters. We built a software tool to merge all Terra radiance granules into one Basic Fusion (BF) granule, which contains not only radiance measurements but also their uncertainties, geolocation, sun-view geometry, and observational time. The tool has produced missionwide basic fusion data (2.4 PB) on Blue Waters. We are currently working with NASA to transfer the entire data set to the Earthdata Cloud. We also produced missionwide metadata for the basic fusion product, which will be ingested into the NASA Common Metadata Repository such that the basic fusion product can be discovered and searched publicly. The overview of this project and progress report was given at the American Geophysical Union 2017 Fall Meeting [4].

In collaboration with the NCSA Advanced Visualization Lab, we developed a visualization tool that dynamically displays and projects the radiance imageries generated from one single BF granule onto a 3D Earth as being orbited by Terra for all of the five Terra instruments. We created and posted an animation clip on YouTube, with image frames from the clip given in Fig. 1. To the best of our knowledge, this is the first animation to simultaneously show five RGB imageries of all five Terra instruments on the same orbit. The tool not only helps us to validate and explore the BF data but will also have a profound educational influence among the broader scientific community.

We carried out science investigations using this Terra fusion data set primarily in two studies in the past year. We fused MISR and MODIS data to characterize biases in retrieved cloud drop effective radius inherent to MODIS-alone retrievals and further examined the underlying causes of the biases. 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. Our corrected cloud drop sizes are now in line with spot measurements from field campaigns. In the second study, we have been working closely with Prof. Ping Yang at Texas A&M University on a specialized MISR and MODIS fusion data set designed for retrieving ice cloud microphysical properties, including ice crystal roughness. Early results from his group were presented at several meeting and conference venues,

showing an altitude (hence temperature) and regional dependence on ice crystal structure.

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 workflow.

ENHANCED DIGITAL ELEVATION MODEL FOR THE ARCTIC

Allocation: NSF PRAC/9,696 Knh

PI: Paul Morin¹

Collaborators: Charles Nguyen², Claire Porter², Ian Howat³, M.J. Noh³, Michael Willis⁴, Brian Bates⁵, Cathleen Williamson⁵, Andrew Johnson⁵

¹University of Minnesota

²Polar Geospatial Center

³The Ohio State University

⁴University of Colorado Boulder

⁵National Geospatial-Intelligence Agency

EXECUTIVE SUMMARY

The Earth's land surface topography is arguably the most important single data set in the geosciences, geographical sciences, and civil engineering as it provided the basis or reference analysis and design. Accurate and timely elevation models are essential to determining the location of rivers and the extent of watersheds in hydrology, to permafrost collapse in areas, to the change in the shape of volcanoes in volcanology and to the effects of human activities. The Polar Geospatial Center (PGC) and its partners at The Ohio State University and University of Colorado Boulder adapted PGC's digital elevation model (DEM) production capabilities from small-area, on-demand creation to systematically process and mosaic the entire Arctic area from the sub-meter stereoscopic commercial imagery archive made

available by the National Geospatial Intelligence Agency. Such datasets make the Arctic one of the best-mapped regions on Earth and enable new aspects of Polar science. This dataset also allows for precise detection of change over time, enabling highly accurate measurements on this rapidly evolving landscape.

RESEARCH CHALLENGE

There is a lack of high-resolution, consistent, high-quality elevation data available for the Arctic and other regions. In 2000, the Shuttle Radar Topography Mission acquired Synthetic Aperture Radar data for the Earth that was processed into an elevation model with a 30-meter horizontal resolution. While invaluable for temperate regions, this mission could not provide coverage in Polar Regions because of the shuttle's orbital inclination. The National

Geospatial-Intelligence Agency (NGA), DigitalGlobe (DG), and PGC built a near-seamless archive of polar sub-meter stereo imagery that consists of almost 100,000 stereo pair strips from the Worldview-1, -2, and -3 satellites. Using photogrammetric algorithms, we are able to construct digital elevation models (DEMs) from the stereo pairs, enabling mapping of surface features at the 2-meter scale for the first time. These data are already being used by the Arctic research community to support activities that include transportation, national defense, land management, sustainable development, and scientific studies. Further, most areas include repeat coverage with a frequency of months or even days. Such temporal coverage can be used for change detection facilitating studies from land use to resource management and environmental change. The scale of this effort—to cover 1/6 of the earth's land at one time, is unprecedented.

METHODS & CODES

Our team has spent five years developing an efficient algorithm for constructing photogrammetric DEMs from satellite imagery with the objective of creating a fully automated pipeline capable of handling large amounts of data and very large areas. The Surface Extraction from TIN-based Search-space Minimization (SETSM) algorithm, initially designed to extract elevation data over ice sheets, has been refined and optimized to handle stereoscopic imagery over any land cover [1,2]. Unlike other DEM extraction algorithms, SETSM's structure eliminates the need for an existing (i.e., "seed") DEM for *a priori* constraints or any data-specific, user-defined search parameters, making it a truly automated algorithm. After an initial preprocessing step that corrects the source imagery for sensor-specific detector alignment artifacts, SETSM takes the two source images and derives increasingly detailed elevation models using its pyramid-based approach. The DEM extraction workflow currently runs on all cores of a single node for efficiency, and several thousand of these single-node tasks are bundled together using the Swift workflow management package in order to effectively submit jobs in hundreds to thousands of node batches. The method can also run across multiple nodes using MPI if needed.

RESULTS & IMPACT

Thus far, we have produced over 136,000 individual 2-m DEMs of the Arctic area. On average, the Arctic is covered six times with images taken over time, with some areas having over 100 unique DEMs for a given location. These data are also processed into continuous mosaics for over 99% of the 20,000,000-km² Arctic. All of these data have been or will shortly be released to the science community and the public through ArcticDEM.org. The Environmental Systems Research Institute (Esri) has developed web services and an interactive viewer for this and other DEM data. These data are now being used by scientists, national institutions, and regional and local governments for a broad range of scientific, civil engineering, and mapping applications. In addition, 18 scientific publications have used ArcticDEM data

since our first release in Aug 2016 and the summer of 2018, with more underway.

WHY BLUE WATERS

No other open research/academic computer has the capacity for this project. Over 17 million node hours (more than ½ billion core hours) were used to process the archive of stereoscopic imagery that was made available. Additionally, the staff at the Blue Waters Project were invaluable in adapting the leadership system to handle the single-node, high-throughput ArcticDEM workflow. With their help, the ArcticDEM project adopted a strategy that enabled ArcticDEM jobs to use primarily backfill nodes on a low priority, which increased overall system utilization, minimized impact on other projects and expanded the amount of computing that was achievable.




Figure 2: The North Slope of Alaska is the most poorly mapped and rapidly changing area of the United States. This image shows detailed topography over the transition from the Brooks Range in the south to permafrost tundra on the northern shores.




Figure 1: Mosaic tiles from Release 6 of ArcticDEM in color hillshade. The final ArcticDEM release will include updated DEMs and mosaics for the entire Arctic.

IMAGE PROCESSING TO BUILD A MULTI-TEMPORAL VEGETATION ELEVATION ECOSYSTEM MODEL OF THE GREAT LAKES BASIN

Allocation: GLCPC/540 Knh

PI: Jennifer Corcoran¹

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

Collaborators: Paul Morin¹, Joe Knight¹, Laura Bourgeau-Chavez⁴

¹University of Minnesota, Twin Cities

²U.S. Fish & Wildlife Service

³SharedGeo

⁴Michigan Tech Research Institute

EXECUTIVE SUMMARY

In 2017, the research team began acquiring, processing, and creating high-resolution, Multi-Temporal Vegetation Elevation Ecosystem Model (MTVEEM) maps across the entire Great Lakes Basin (GLB) where stereo, submeter, optical satellite imagery had been acquired over the last 15 years. This year, we have continued processing newly acquired imagery and adding processing steps to monitor changes in ecosystems over time, providing a new, essential data set that is otherwise not available to many fields of science and natural resource managers. The information we are generating will improve both our understanding of the complex processes at work in the GLB as well as our ability to forecast and mitigate adverse changes to the world's largest freshwater ecosystem. The amount of data to process and analyze for the GLB is well beyond the resources available from most academic, private, and government systems, and this computational need can only be addressed with a leading-edge petascale resource such as Blue Waters.

RESEARCH CHALLENGE

Ecosystem management requires knowledge of 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), is unprecedented for aerial or satellite surveys. Additionally, the binational management (Canada and United States) 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 data sets represent single time points and were not planned as continuous, basinwide acquisitions. Having high-resolution, multi-temporal information in three dimensions enables the kind of science that can address a multitude of critical questions that surround the ecosystems of the GLB.

The research questions remain: 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 from flooding, forest blowdown, fire, insects and disease outbreaks) and anthropogenic changes (such as harvest and land-cover change). Further, MTVEEM will improve habitat and biological modeling. Finally, MTVEEM will be used binationally to better visualize canopy change in forested habitats and freshwater wetland resources within the GLB.

METHODS & CODES

Stereo-mode acquisition through Digital Globe over the entire GLB started in 2016 and will continue through 2019 as clouds and other higher tasking priorities permit (Fig. 1). In 2017, we processed over 83,000 stereo pairs, where each job consisted of converting the input imagery into a standard format (GeoTIFF) and then calling the elevation extraction software (SETSM) [1]. We expect 50,000 additional satellite image stereo pairs for 2018. Each stereo pair task is run on a single node, submitted in batches of 2 to 100 tasks per job to the low-priority queue to maximize scheduler throughput. Complete processing of one stereo pair to 2m takes an average of 12 node hours (charged as six node hours due to using the low-priority queue), totaling 300,000 node hours. Additionally, we estimate it will take 150,000–200,000 node hours to process ortho-images and explore producing classifications based on the total number of image pairs we will have in hand.

RESULTS & IMPACT

As the data are processed in 2018, the resulting surface canopy models will be openly available initially through the University of Minnesota. Other partners' online distribution systems, such as NOAA's Digital Coast and the Great Lakes Observing System, will also be used. The final product, a seamless and registered surface elevation ecosystem model (MTVEEM) of the GLB will enable a large range of science activities at substantially higher resolution than currently available (current status shown in Fig. 1). These canopy maps and change detection products will provide positional accuracies of less than a couple meters with the added ground control points. We are assessing semidecadal changes in priority GLB areas where LiDAR-derived digital surface models



Figure 1: Great Lakes Basin digital surface model (DSM) production status as of March 2018. The source of the stereo imagery used to produce these DSMs span the archive, starting in about 2008 to early 2018. (2018 DigitalGlobe NextView License.)

from six to nine years ago are available. We are also beginning to look at intraseasonal differences by processing surface models from satellite stereo image pairs within a single growing season (see example of a blowdown event in Fig. 2). These preliminary results show great promise for providing valuable data to myriad coastal and terrestrial ecosystem science researchers and decision-makers across the entire GLB [2,3].

Our primary concentration will continue to be the GLB, with extended temporal and geographical footprints as efficiencies and capacities improve. Though we will continue to make significant progress in the GLB, this year's reallocation will also allow in-depth explorations of habitat types and terrain characteristics beyond our pilot study areas. This will allow us to adjust processes and perform quality control checks to enable scaling to much larger geographic regions, such as all of North America. With any remaining processing hours, the project may start half-meter MTVEEM processing, which has already been tested.




Figure 2: Near Duluth, Minn. False color-infrared images (left), corresponding digital surface models (DSM) (center), and difference map derived from a subtraction of the two DSM dates (right). (2018 DigitalGlobe NextView License.)

WHY BLUE WATERS

Stereo satellite imagery allows for the generation of highly accurate surface elevation models. We have already tasked stereo-mode acquisition through Digital Globe over the entire GLB. Each stereo pair is about 1.25 GB and the total number of pairs processed to date is about 83,000, soon to exceed 100,000. The amount of stereo imagery in a study area the size of the Great Lakes Basin and the computational burden to process each of these image pairs is well beyond those available from academic, private, and government systems. This is precisely why we need a leading-edge petascale resource such as Blue Waters.

PUBLICATIONS & DATA SETS

DigitalGlobe, WorldView-3 scene A 104001001E4FD500; B 104001001D127E00, Level Standard 2A, DigitalGlobe, Longmont, Colo., 06/08/2016.

DigitalGlobe, WorldView-3 scene A 1040010021C77D00; B 1040010020D20300, Level Standard 2A, DigitalGlobe, Longmont, Colo., 08/03/2016.

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

Allocation: Illinois/550 Knh
PI: Donald J. Wuebbles¹
Co-PI: Xin-Zhong Liang²

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

EXECUTIVE SUMMARY

This project looks into the changing climate and its impact on U.S. air quality, with a special focus on fine particulate matter and ozone, projecting their trends and quantifying key source attribution. We are using a state-of-the-science dynamic prediction system that couples a global climate–chemical transport model with regional climate and air-quality models over North America, to determine individual and combined impacts of global climate and emissions changes on U.S. air quality. These include uncertainty evaluations, from the present to 2050, under multiple climate and emission scenarios. The results from the global and regional model simulations for the past are evaluated with observational data to assess the capabilities of the model simulation and impacts of emissions change, climate change, and long-range transport on future U.S. air pollution.

RESEARCH CHALLENGE

The objectives of this study are to better understand how global changes in climate and emissions will affect U.S. air quality, focusing on particulate matter and ozone, to project their future trends, and quantify key source attributions. Thus, we aim to provide actionable information for U.S. environmental planners and decision makers to design effective dynamic management strategies, including local controls, domestic regulations, and

international policies, to sustain air quality improvements in a changing world. We are applying a state-of-the-science dynamic prediction system that couples global climate–chemical transport models with regional climate and air-quality models over North America 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. The aim is to quantify pollution sources and assign their attribution—natural vs. anthropogenic emissions, national versus international agents, natural variations versus climate changes—with associated probability and uncertainty.

We are conducting three primary experiments using the dynamic prediction system: (1) historical simulations for the period of 1990–2015 to establish the credibility of the system and refine process-level understanding of U.S. regional air quality; (2) projections for the period of 2030–2060 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. The advanced state of the prediction system will produce more complete scientific understanding of the challenges from global climate and emissions changes imposed on U.S. air quality management and a more reliable projection of future pollution sources and attribution changes.

METHODS & CODES

The proposed research uses 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.

We used CESM1.2 (Community Earth System Model) default emissions, which represent surface emissions of ~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 [1,2]. The results from the runs done in the current allocation have been presented in scientific conferences (e.g., Sanyal, et al., at the American Geophysical Union Annual meeting in December 2017).

The Blue Waters allocation supports the coupling and testing of the regional CWRf–CMAQ modeling system at NCSA. Historical runs (1990–2016) of U.S. air quality have been driven by the ECWMF ERI reanalysis and finished to test the model performance and the trends of U.S. air quality in the past decades. Runs for future periods are now in process.

RESULTS & IMPACT

The CESM model simulation was used to look into long-term Particulate Matter (PM) PM_{2.5} concentration globally from 1980–2005. Fig. 1 shows the annual average PM_{2.5} concentration, globally as well as in different regions such as the continental United States, Europe, India, and China. The PM_{2.5} concentration shows a decreasing trend globally and in the continental United States, Europe, and India whereas it shows an increasing trend in China.

To evaluate the model performance at simulating the major air pollutants, especially ozone and PM_{2.5}, we used EPA AQS ozone measurements and PM_{2.5} measurements from IMPROVE and CASTNET networks to evaluate the results (Fig. 2). The ozone comparison shows that CMAQ can capture the distribution of ozone pollution in the United States. The CWRf–CMAQ modeling system has substantial underestimations in urban and suburban areas, such as the central valley in California. For the PM_{2.5} pollution, both comparisons of CMAQ simulations and IMPROVE/CASTNET measurements suggest that the CWRf–CMAQ modeling system can successfully capture the PM_{2.5} pollution pattern, while some isolated sites have substantial discrepancies. This shows the capability of the dynamic CWRf–CMAQ system. We will extend the simulation period of CMAQ to 25 years (1990–2015) next year.

We hypothesize that the integration of the most advanced modeling system, most updated emissions treatment, multiscale processes representation, and various climate-emission scenarios assessment will improve the predictive capability of the model. This should result in more reliable projection of future likely changes in PM_{2.5}, O₃ and related pollutants as well as their global and regional sources. We will make a major contribution to a key goal of the U.S. Environmental Protection Agency’s Strategic Plan to address climate change. The advanced state-of-the-prediction system will produce more complete scientific understanding of the challenges from global climate and emissions changes imposed on U.S. air quality management and a more reliable projection of future pollution sources and attribution changes. The outcome will provide actionable information for U.S. regional and state agencies to design effective strategies to meet the air-quality standards and achieve sustainability in a changing world. Climate change will affect air quality. Policymakers need to understand what these effects are so they can plan for such effects in their management of air quality. These studies will guide policymakers in how to respond to climate change as they sort out air-quality management goals for the future.

WHY BLUE WATERS

The computational demand of high-resolution climate models used in this project is very extensive. In addition, we are using a fully coupled model of the Earth’s climate system with interactive chemistry, which is also computationally expensive even when not run at high resolution. Blue Waters, with its petascale computational facility, large number of nodes, and storage capability for the output from the high-resolution model simulation, is essential for our project. Blue Waters’ staff have been critical at 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.

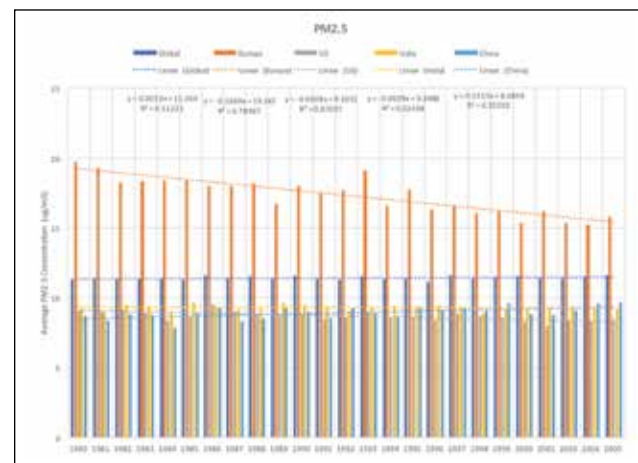


Figure 1: Annual average PM_{2.5} concentration globally and in the United States, Europe, India, and China.




Figure 2: Comparison of EPA Air Quality Standards observations and Community Multiscale Air Quality (CMAQ) simulations. Ozone values are from summer JJA MDA8 and PM_{2.5} measurements are the annual mean. Dots show AQS/CASTNET/IMPROVE observations and background shows CMAQ results (upper panel). Lower panel shows the ratio of observations to CMAQ results.

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

Allocation: NSF PRAC/2,660 Knh
 PI: Sarma Rani¹

¹University of Alabama, Huntsville

EXECUTIVE SUMMARY

The overall goal of our research is to investigate the role of turbulence in driving the relative velocities and relative positions of inertial particles in isotropic turbulence. During the second year of the PRAC grant, we hypothesized that the deterministic forcing used in the Direct Numerical Simulations (DNS) to achieve stationarity had the effect of artificially increasing the temporal coherence of eddies as the Reynolds number increased. This artifact, we surmised, would lead to higher Eulerian two-time correlations, and, thereby, diffusivities, particularly at large separations. In the third year, we investigated the effects of forcing on the Eulerian two-time correlations of fluid relative velocities in DNS of isotropic turbulence. Accordingly, we performed DNS of homogeneous, isotropic turbulence containing disperse but fixed particles for two grid sizes—128³ and 512³—while employing two forcing schemes, one deterministic and the other stochastic. The high-performance computing resources of the Blue Waters system were indispensable in performing the DNS runs needed to validate the hypothesis.

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 $St_\eta \sim 10\text{--}100$, 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) probability density function (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 timescale 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, we surmised that the forcing artificially increased the temporal coherence of the large-scale eddies, particularly as Re_λ was increased. The increased coherence led to higher magnitudes of the two-time correlations of fluid relative velocities (and thereby diffusivities) at separations that scaled with the integral length scale. The third-year study was motivated by the desire to quantitatively investigate the above hypothesis. Accordingly, we performed direct numerical simulations of forced isotropic turbulence laden with disperse but fixed particles. Two types of forcing were used to achieve statistical stationarity in the DNS runs.

METHODS & CODES

We performed DNS of forced isotropic turbulence 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 was 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, we adopted a pseudospectral approach wherein the nonlinear terms were first computed in physical space and then transformed into the spectral space. We used the P3DFFT library [2] to carry out the transforms between physical and spectral spaces.

RESULTS & IMPACT

In Fig. 1, we present the longitudinal component of the Eulerian two-time correlation of fluid relative velocities, i.e., $\langle \Delta u(r, x, 0) \Delta u(r, x, \tau) \rangle_{||}$, as a function of dimensionless time separation $\tau^* = \tau u_{rms}/L$ for $Re_\lambda = 210$. We compared the correlations obtained from the DF and SF1–SF5 simulations compared at four dimensionless separations, $r/L = 0.56, 1.12, 2.24$ and 3.36 , where L is the integral length scale and u_{rms} is the RMS fluctuating velocity. It is clear that at all four separations, the DF longitudinal correlation is




Figure 1: Longitudinal correlation $R_{||}(\tau, \tau) = \langle \Delta u(r, x, 0) \Delta u(r, x, \tau) \rangle_{||}$ as a function of dimensionless time separation $\tau^* = \tau u_{rms}/L$ for $Re_\lambda = 210$ at four dimensionless separations: (a) $r/L = 0.56$, (b) $r/L = 1.12$, (c) $r/L = 2.24$, and (d) $r/L = 3.36$. DF stands for deterministic forcing and SF for stochastic forcing.

higher than the SF1–SF5 correlations (except at small τ^*). At the three separations greater than L , we see that the DF correlation is significantly greater than the SF correlations.

WHY BLUE WATERS

Direct Numerical Simulations are the most accurate numerical approach to resolve all the temporal and length scales in a turbulent flow. However, DNS of particle-laden turbulent flows are computationally intensive, 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 microscale Reynolds number $Re_\lambda \sim 600$ with grid sizes $\sim 2,048^3$. This requires running our code on tens of thousands of cores. Also, each DNS run generates several terabytes of data. Due to these CPU time and storage requirements, the Blue Waters supercomputer is the ideal platform to achieve our objective. Further, 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 access to Blue Waters, where we ran the code on 20,000 cores.

PUBLICATIONS & 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. *American Physical Society 69th Annual Meeting of the Division of Fluid Dynamics* (November 20–22, 2016, Portland, Ore.).

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. *American Physical Society 69th Annual Meeting of the Division of Fluid Dynamics* (November 20–22, 2016, Portland, Ore.).

MODELING 4D EARTH EVOLUTION: FROM CONTINENTAL CRATONS TO THE YELLOWSTONE SUPERVOLCANO

Allocation: NSF PRAC/800 Knh
PI: Lijun Liu¹
Collaborator: Manuele Faccenda²

¹University of Illinois at Urbana-Champaign
²Università di Padova

EXECUTIVE SUMMARY

Many aspects of Earth’s dynamic evolution remain poorly understood. Our research focuses on two projects: 1) what causes the most stable portion of continents—cratons—to experience dramatic elevation change and internal deformation; and 2) what is the ultimate source of heat that fuels the Yellowstone supervolcano. By positioning all known and unknown tectonic elements in a common evolutionary system model, we discovered novel insights regarding both questions. We found that the cratonic lithosphere is not as stable as previously thought and its lower part would delaminate when perturbed by mantle plumes from below, causing rapid surface uplift and crustal erosion. Our answer to the supervolcano question is rather surprising: The majority of the hot mantle maintaining past Yellowstone volcanism came from below the Pacific Ocean and not from the deep mantle plume right below as traditionally believed. Both works were recently published in *Nature Geoscience*, with extensive media reports.

RESEARCH CHALLENGE

The dynamic behavior of the deep Earth such as mantle flow and heat migration near a subduction zone as well as the resulting surface responses such as volcanic eruption, topographic changes, and drainage system migration are physically and numerically complex. A quantitative understanding of these processes, especially their past evolution, is vital for explaining the various geological, geophysical, and geohazards observations. The challenge we geodynamicists face is how to accurately reproduce the various activities within the inaccessible interior of the solid Earth and how to connect them with the geological records.

METHODS & CODES

In order to quantitatively understand the current and past dynamic processes within the deep Earth, we adopt multiple data assimilation techniques into large-scale geodynamic modeling. For example, we employ seismic tomography images to represent the present-day mantle temperature and viscosity profiles. We assimilate the motion of tectonic plates that is 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 tested the scalability and I/O performance of this code on Blue Waters. CitcomS is community-based software; it has been designed and tested on many supercomputers.




Figure 1: Illustration of craton evolution in Africa and South America. (Upper) The stable cratonic lithosphere was destabilized by underlying mantle plumes leading to delamination of the denser root in the Cretaceous era. (Lower) The missing roots thermally grew back toward the present, while recording this history in the lithospheric fabrics.

RESULTS & IMPACT

In the first project about the evolution of the cratonic lithosphere, we quantitatively reproduced the subduction and mantle flow below South America since 100 million years ago [2]. Based on these reproductions, we found that the widespread fast-mantle seismic structures below the South Atlantic region are stalled pieces of foundered continental lithosphere [3]. Further examination of the topography, gravity, geology, and seismic properties of the lithosphere reveals that large volumes of the cratonic lithosphere were delaminated into the underlying mantle during the Cretaceous era. This caused the surface to uplift and shed enormous amounts of sediment offshore, leading to thinned crust and deformed lithosphere. Our work, therefore, revised the traditional view that the cratonic lithosphere is neutrally buoyant

and tectonically stable [4]. Part of this work was published in *Nature Geoscience*, and has enormous potential impact on other fields of geoscience, including seismology, geology, sedimentology, petrology, mineral physics, geomorphology, and the like.

In the second project about the heat source of the Yellowstone volcano, we developed a physical model that is consistent with all available geophysical and geological data constraints [5]. Our model reveals that most of the heat below the Snake River Plain and Yellowstone caldera originally came from under the Pacific Ocean [6]. This challenges the traditional hypothesis that the Yellowstone supervolcano has been fueled by a deep-mantle plume right below Wyoming. This result was also published in *Nature Geoscience*, with direct impact on many other fields.

WHY BLUE WATERS

The enormous amount of data processing and computation required for this work makes Blue Waters the best platform for our research.

PUBLICATIONS & DATA SETS

Hu, J., et al., Modification of the Western Gondwana craton by plume-lithosphere interaction. *Nature Geoscience*, (2018), DOI:10.1038/s41561-018-0064-1.

Zhou, Q., L. Liu, and J. Hu, Western U.S. volcanism due to intruding oceanic mantle driven by ancient Farallon slabs. *Nature Geoscience*, (2018), DOI:10.1038/s41561-017-0035-y.

Sun, W., et al., Post-ridge-acceleration of the Indian plate induced by slab rollback. *Solid Earth Sci.*, (2018), DOI:10.1016/j.sesci.2017.12.003.

Zhou, Q., and L. Liu, A Hybrid Forward-Adjoint Data Assimilation Approach for Reconstructing the Temporal Evolution of Mantle Dynamics. *Geochem. Geophys. Geosys.*, (2017), DOI:10.1002/2017GC007116.

Chen, L., F. Capitanio, L. Liu, and T. Gerya, Crustal Rheology Controls on the Tibetan Plateau Formation during India-Asia Convergence. *Nature Communications*, 8 (2017), DOI:10.1038/ncomms15992.

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 (2016), DOI:10.1002/2015GL067131.

Hu, J., L. Liu, A. Hermsillo, 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.

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.*, (2016), 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.*, (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 2: Schematic view of mantle dynamics below the western United States. The diverging motion of the hot intruding mantle, corresponding to the Newberry (NB) and Yellowstone (YS) tracks on the surface, is controlled by the sinking Farallon slab (blue) and the retreating Juan de Fuca slab (green).

SIMULATING ALTERNATIVE CALIFORNIA EARTHQUAKE HISTORIES USING RATE AND STATE FRICTION

Allocation: NSF PRAC/9,620 Knh

PI: John Vidale¹, Thomas Jordan (former PI)¹

Co-PIs: Yifeng Cui², Philip Maechling¹, Kim Olsen³, Ricardo Tabora⁴

Collaborators: Jacobo Bielak⁵, Md Monsurul Huda⁴, Naem Khoshnevis⁴, Andrea C. Riaño-Escandón⁵, Alexander Breuer², Scott Callaghan⁶, Jacquelyn Gilchrist¹, Christine A. Goulet⁶, R.W. Graves⁷, Dawei Mu², Dmitry Pekurovsky², Daniel Roten², Bruce Shaw⁸, Qian Yao³

¹University of Southern California

²San Diego Supercomputer Center

³San Diego State University

⁴University of Memphis

⁵Carnegie Mellon University

⁶Southern California Earthquake Center

⁷United States Geological Survey

⁸Columbia University

EXECUTIVE SUMMARY

The Southern California Earthquake Center (SCEC) is using Blue Waters to develop physics-based earthquake forecasts for broad impact users including the U.S. Geological Survey, California State emergency planning authorities, and civil engineering organizations. During the last year, we used the physics-based earthquake cycle simulator code known as RSQSim to produce several million-year California earthquake catalogs to investigate how fault complexities affect the probabilities of large, multifault ruptures and multievent sequences. This research is improving our understanding of earthquake processes, which will improve probabilistic seismic hazard analysis in the United States and benefit earthquake system science worldwide.

RESEARCH CHALLENGE

SCEC conducts an earthquake system science research program developing physics-based, predictive models of earthquake processes to improve probabilistic seismic hazards assessments. SCEC researchers use physics-based computational models, observation-based 3D earth structure models, and high-performance computer resources to improve probabilistic seismic hazard forecasts for California. Our research is leading to improved Probabilistic Seismic Hazard Analysis (PSHA) methods and results. Our computationally intensive CyberShake seismic hazard models estimate peak ground motions using 3D wave propagation simulations. Blue Waters enables SCEC seismic hazard modelers to benefit from significant advances in computing and storage capabilities. By improving PSHA and ground-motion simulations, this project contributes to improved seismic design and safety because PSHA ground motion estimates guide the long-term construction of a seismically safe built environment.

METHODS & CODES

SCEC's earthquake system science research requires an integrated collection of earth structure models and physics-

based simulation codes. SCEC has developed a seismic hazard software ecosystem that includes traditional probabilistic seismic hazard analysis software (OpenSHA), California velocity models (UCVM), finite difference wave propagation software (AWP-ODC), finite element wave propagation software (Hercules), and physics-based probabilistic seismic hazard software (CyberShake). The computational scale of our seismic hazard studies has led to close collaborations with the Blue Waters technical groups, who have helped us identify and remove technical roadblocks in our large-scale research calculations.

RESULTS & IMPACT

This year, the SCEC team used a physics-based code, RSQSim [1], to produce long-term (one million+ years) synthetic earthquake catalogs that comprise dates, times, locations, and magnitudes for earthquakes in the California region. RSQSim simulations impose stresses upon a representation of the California fault systems. These stresses are then distributed throughout a complex system of faults, generating cascades of earthquakes during the simulations. The RSQSim output is an earthquake rupture catalog that represents a sequence of earthquakes of various magnitudes expected to occur on the San Andreas Fault over a given time range. These catalogs represent a possible history of California earthquakes over the last million years.

After using RSQSim earthquake catalogs to estimate California seismic hazards, SCEC compared the RSQSim hazard estimates to seismic hazard estimates from traditional, empirically derived California Earthquake Rupture forecasts, called UCERF2 and UCERF3, which were designed and developed by SCEC and the U.S. Geological Survey. These California seismic hazard estimate comparisons (Fig. 1) show that the physics-based RSQSim model can replicate the seismic hazard estimates derived from the empirical models, but with far fewer statistical assumptions. This agreement gives researchers confidence that the seismic hazard

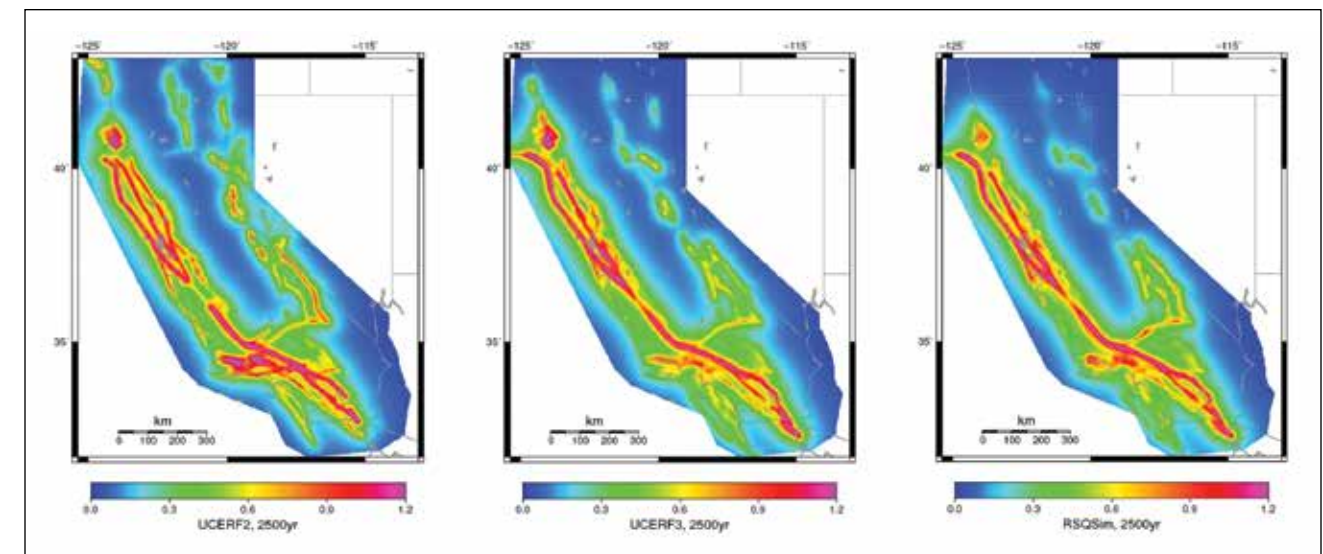


Figure 1: These maps show estimated shaking hazard produced using the RSQSim earthquake catalog as compared to the UCERF2 and UCERF3 hazard models. Maps show Peak Ground Acceleration with a probability of exceedance of 2% in 50 years. Units are in fractions of the acceleration of gravity, g. (a) UCERF2, (b) UCERF3, (c) RSQSim model.

models for California are consistent with what we know about earthquake physics.

WHY BLUE WATERS

SCEC's large-scale seismic hazard studies would not be possible without Blue Waters. Blue Waters is a highly productive research environment due to its variety and large number of computing nodes, the large filesystem, the fast I/O system and network, and the science-friendly system administration policies. Blue Waters makes it possible to perform the large-scale physics-based seismic hazard simulations that can reduce the uncertainties in long-term seismic hazard forecasts and improve risk assessments for critical facilities such as large dams, nuclear power plants, lifelines, and energy transportation networks.

PUBLICATIONS & DATA SETS

Restrepo, D., and R. Tabora, Multiaxial cyclic plasticity in accordance with 1D hyperbolic models and masing criteria. *Int. J. Numer. Anal. Meth. Geomech.*, submitted (2018).

Khoshnevis, N., and R. Tabora, Prioritizing ground motion validation metrics using semi-supervised and supervised learning. *Bull. Seismol. Soc. Am.*, in revision (2018).

Small, P., et al., The unified community velocity model software framework. *Seismol. Res. Lett.*, 88:6 (2017), pp. 1539–1552.

Savran, W.H., Developing stochastic models as inputs for high-frequency ground motion simulations. Doctoral dissertation, San Diego State University (2017).

Aslam, K., and R. Tabora, 3D earthquake ground motion simulations for the Christchurch area including the effects of the surface topography. *Abstr. QuakeCoRE Annu. Meet.* (Taupo, New Zealand, September 3–6, 2018).

Callaghan, S., et al., Using CyberShake 3D Ground Motion Simulation Workflows to Advance Central California PSHA. *Seismological Society of America Annual Meeting* (Denver, Colo., April 18–20, 2017).

Callaghan, S., et al., 10 years of CyberShake: Where are we now and where are we going with physics-based PSHA. *2017 Southern California Earthquake Center Annual Meeting* (Palm Springs, Calif., September 10–13, 2017).

Callaghan, S., et al., CyberShake: bringing physics-based PSHA to central California. *2017 Southern California Earthquake Center Annual Meeting*, poster 303, (Palm Springs, Calif., September 10–13, 2017).

Crouse, C., et al., Site-Specific MCER Response Spectra for Los Angeles Region based on 3-D Numerical Simulations and the NGA West2 Equations. *Proceedings of the 11th National Conference in Earthquake Engineering* (Los Angeles, Calif., June 25–29, 2018).

Gill, D.J., et al., UCVM: From supercomputers to laptops, querying and visualizing 3D seismic velocity models. *Abstr. SSA Annu. Meet.* (Denver, Colo., April 18–20, 2017).

Gilchrist, J.J., et al., Simulation Based Earthquake Forecasting with RSQSim. *2017 Seismological Society of America Annual Meeting* (Denver, Colo., April 18–20, 2017).

Gilchrist, J.J., et al., Conditional Probabilities of Large Earthquake Sequences in California from the Physics-based Rupture Simulator RSQSim. Poster presentation at *2017 Southern California Earthquake Center Annual Meeting* (Palm Springs, Calif., September 10–13, 2017).

Gilchrist, J.J., et al., Conditional Probabilities of Large Earthquake Sequences in California from the Physics-based Rupture Simulator

Continued on page 101

EXTREME CONVECTIVE STORMS UNDER CLIMATE CHANGE

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

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

One projected impact of human-induced climate change is an increase in the intensity of landfalling hurricanes. In our work, we have found that these more intense hurricanes also appear likely to be more hazardous in terms of inland flooding and tornado generation.

RESEARCH CHALLENGE

The future impact of thunderstorms and thunderstorm systems such as hurricanes under a globally warmed climate are still uncertain. Part of this uncertainty is related to the fact that thunderstorm hazards—tornadoes, hail, damaging straight-line winds, lightning, and localized flooding—have spatial scales that fall well below the effective resolution of typical climate models. Modeling approaches such as dynamical downscaling have begun to address this resolution issue. However, their applications thus far have generally been unconcerned with the basic question of whether significant events in the current climate will be more or less significant in the future. The answer to this question is important from the perspective of basic science but will also help to inform decision-makers such as emergency managers on how to prepare for future disasters.

METHODS & CODES

We used an event-based implementation of the pseudo-global warming (PGW) methodology. Modified atmospheric states drawn from global climate model (GCM) output were applied to constrain Weather Research and Forecasting (WRF) model simulations at high resolution. We supplemented these PGW simulations with idealized simulations using Cloud-Model 1 (CM1). Both WRF and CM1 are community codes.

RESULTS & IMPACT

Hurricane Ivan (2004) is the historical case of interest here, in part because of its relative intensity, but also because it generated a record-setting 118 tornadoes as well as considerable inland flooding. Thus, we were motivated to determine if such extreme tropical cyclone tornado (TCT) generation would be further enhanced in a future climate. Our basic approach was to compare a control simulation of Ivan to simulations in an atmosphere modified by PGW. The PGW simulations involved future climate conditions over the late (2080–2090) twenty-first century period under Representative Concentration Pathways 8.5, as extracted from three Coupled Model Intercomparison Project phase 5 GCMs (National Center for Atmospheric Research, Model for

Interdisciplinary Research on Climate, and Geophysical Fluid Dynamics Laboratory). Changes in tropical cyclone (TC) intensity and TCT generation for the PGW-modified Ivan were documented and analyzed.

Compared to the control, all three PGW simulations exhibited more intense TCs. The TCs under PGW also produced significantly more accumulated rainfall over the course of Ivan's inland progression. In addition, each of the PGW TCs generated more prelandfall TCTs than did the control simulation; more numerous and also more intense postlandfall TCTs resulted from PGW in some of the simulations. These and other experiments lend support to the hypothesis that an increase in sea surface temperature due to human-induced climate change will intensify landfalling TCs, which in turn will result in more numerous tornadoes.

In our forthcoming work, these PGW results are being used in a county-level event-based assessment of the risk of inland TC hazards, particularly TCTs.

WHY BLUE WATERS

The relatively small size of thunderstorms and the ranges of relevant scales within tropical cyclones, coupled with their episodic occurrence, necessitate a research approach that can account for temporal scales from minutes to decades and spatial scales of hundreds of meters to thousands of kilometers. In other words, we require very large geospatial domains that have fine gridpoint 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.

PUBLICATIONS & DATA SETS

Carroll-Smith, D., L.C. Dawson, and R.J. Trapp, High resolution real-data WRF modeling and verification of tropical cyclone tornadoes associated with Hurricane Ivan 2004. *Electronic Journal of Severe Storms Meteorology*, in revision (2018).

Marion, G.R., and R.J. Trapp, The dynamical coupling of convective updrafts, downdrafts, and cold pools in simulated supercell thunderstorms. *Journal of Geophysical Research-Atmospheres*, in review (2018).

Carroll-Smith, D., and R.J. Trapp, Hurricane Ivan (2004) under pseudo-global warming: An investigation of the influence of anthropogenic climate change on tropical cyclone inland hazards. *Journal of Applied Meteorology and Climatology*, in review (2018).

Continued from page 99

RSQSim. Poster presentation at 2017 American Geophysical Union Annual Meeting, (New Orleans, La., December 11–15, 2017).

Goulet, C.A., HPC Use for Earthquake Research. *HPC User Forum* (Tucson, Ariz., April 18–21, 2018).

Isbiloglu, Y., and R. Taborda, A preliminary study about the influence of building clusters on the variability of the ground motion during earthquakes. *Abstr. QuakeCoRE Annu. Meet.* (Taupo, New Zealand, September 3–6, 2017).

Jordan, T.H., et al., Integrating Physics-based Earthquake Cycle Simulator Models and High-Resolution Ground Motion Simulations into a Physics-based Probabilistic Seismic Hazard Model. *Blue Waters Symposium* (Bend, Ore., May 16–19, 2017).

Khoshnevis, N., and R. Taborda, An application of machine learning techniques to the evaluation of goodness-of-fit scores used in earthquake ground motion validation. *Proc. SCEC Annu. Meet.* (Palm Springs, Calif., September 10–13, 2017).

Khoshnevis, N., and R. Taborda, An overview of computational learning theory for use in ground motion simulation. *Proc. Machine Learning in Solid Earth Geoscience* (Santa Fe, N.M., February 20–23, 2018).

Khoshnevis, N., and R. Taborda, Application of constrained k-means clustering in ground motion simulation validation. *Proc. AGU Fall Meet.* (New Orleans, La., December 11–15, 2017).

Khoshnevis, N., and R. Taborda, Exploring the implementation of an equivalent linear method in 3D to approximate nonlinear response in regional ground motion simulation. *Proc. SSA Annu. Meet.* (Denver, Colo., April 18–20, 2017).

Khoshnevis, N., and R. Taborda, Towards an informed decision making in validation metrics. *Response History Analysis Validation Workshop, QuakeCoRE Annu. Meet.* (Taupo, New Zealand, September 3–6, 2017).

Maechling, P., et al., Current capabilities of the SCEC Unified Community Velocity Model (UCVM) software framework. *Proc. SCEC Annu. Meet.*, Poster No. CME-302 (Palm Springs, Calif., September 10–13, 2017).

Maechling, P.J., et al., The SCEC software ecosystem for enhancing earthquake system science. *Abstr. SSA Annu. Meet.* (Denver, Colo., April 18–20, 2017).

Taborda, R., and Y. Isbiloglu, Influence of buildings spacing in site-city interaction effects. *Proc. SCEC Annu. Meet.*, Poster No. EEII-240 (Palm Springs, Calif., September 10–13, 2017).

PHYSICS & ENGINEERING

AEROSPACE

FLUIDS

HIGH ENERGY PARTICLE PHYSICS

HIGH ENERGY PLASMA PHYSICS

MATERIALS

NANOTECHNOLOGY

NUCLEAR ENGINEERING

NUCLEAR PHYSICS

PARTICLE PHYSICS

QUANTUM

104 *Coupled Multiphysics of Advanced Molten-Salt Nuclear Reactors*

106 *Learning Low-Dimensional Feature Dynamics of Turbulent Flows Using Deep Convolutional Recurrent Autoencoders*

108 *Electron Dynamics of Ion-Irradiated Graphene*

110 *Magneto-Optical Kerr Effect of Antiferromagnetic Materials in External Magnetic Fields*

112 *Petascale Simulation of High-Energy-Density (HED) Plasmas*

114 *Resolution Effects and Extreme Events in Petascale Turbulence Simulations*

116 *Advanced Digital Technology for Materials and Manufacturing*

118 *Multiphase Turbulent Flow Modeling of Gas Injection into Molten Metal to Minimize Surface Defects in Continuous-Cast Steel*

120 *High-Resolution Numerical Simulation of Flow and Sediment Transport through Aquatic Vegetation*

122 *High-Fidelity Numerical Simulations of the Reacting Flow Field of an Annular Combustor Model of an Aircraft Engine*

124 *Fully Three-Dimensional Kinetic Simulations of Unsteady Shock-Boundary Layer Interactions*

126 *Mapping Proton Quark Structure—Looking Inside the Proton: How Do Quarks Spin?*

128 *Enabling Discovery at the Large Hadron Collider through Data-Intensive Computation and Machine Learning*

130 *High-Energy Physics on Blue Waters*

132 *Epigenetic Identification and Mapping Using Solid-State 2D Nanopores*

134 *Computational Modeling of Rotating Detonation Engines*

136 *Collaborative Research: Advancing First-Principle Symmetry-Guided Nuclear Modeling for Studies of Nucleosynthesis and Fundamental Symmetries in Nature*

138 *Novel Metal–Carbon Materials*

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

142 *Properties of Dense Hydrogen*

144 *Leading Future Electronics into the Nano Regime Using Quantum Atomistic Simulations in NEMO5*

146 *Modeling Millions of Atoms and Their Electrons: Tight-Binding Using a Spectral Projection Purification Method*

148 *Simulating the Emergent Phenomena Arising from Strongly Correlated Systems*

150 *Pressure Fluctuations Induced by Supersonic Turbulent Boundary Layers*

152 *Topologies of Electronic States of Nanocrystalline Materials*

154 *Understanding Hydrogen Storage in Metal–Organic Frameworks Using Massively Parallel Electronic Structure Calculations*

156 *Numerical Simulations of the Inertial Collapse of Individual Gas Bubbles near a Rigid Surface*

158 *High-Resolution Modeling of Turbulence in Wave Boundary Layers*

160 *Coarse-Grained Force Field for Ionic Liquids*

162 *High-Fidelity Numerical Simulations of Turbine Vane Heat Transfer under Freestream Turbulence*

164 *Mechanistic Insights into Hydration of Solid Oxides*

166 *Machine Learning for High-Energy Physics: Particle Identification and Regression Using Deep Neural Networks*

168 *Petascale Simulations of Graphene Nanoribbon Device Structures*

170 *Interfacial Load Transfer Mechanisms in Carbon Nanotube–Polymer Nanocomposites*

172 *Improving Virtually Guided Product Certification with Implicit Finite Element Analysis at Scale*

174 *Running PlasCom2 on Blue Waters with OpenMP Offloading*

176 *Quantum Magnets and Models*

COUPLED MULTIPHYSICS OF ADVANCED MOLTEN-SALT NUCLEAR REACTORS

Allocation: Blue Waters Professor/60 Knh
PI: Kathryn Huff¹
Collaborator: Alexander Lindsay²

¹University of Illinois at Urbana-Champaign
²Idaho National Laboratory

EXECUTIVE SUMMARY

The Advanced Reactors and Fuel Cycles Group (ARFC) models and simulates nuclear reactors and fuel cycles with the aim of improving the safety and sustainability of nuclear power. In the context of high-performance computing, this work couples multiple physics at multiple scales to improve the design, safety, and performance of advanced nuclear reactors. In particular, thermal-hydraulic phenomena, neutron transport, and fuel reprocessing couple tightly in nuclear reactors. Detailed spatially and temporally resolved neutron flux, temperature distributions, and isotopic compositions can improve designs, 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, thermal hydraulics, and fuel reprocessing in multiple advanced, molten-salt-fueled nuclear reactor designs.

RESEARCH CHALLENGE

Nuclear power provides 19% of the total electricity generated in the United States and is our largest source of clean energy. The current state of the art in advanced nuclear reactor simulation focuses primarily on traditional light-water reactor designs. 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 already unparalleled safety and sustainability of nuclear power. High-fidelity simulation of




Figure 1: The coupled fast neutron flux and temperature in a 3D cuboidal model of a molten-salt reactor. Moltres coupled results have been validated against experimental results from the Molten-Salt Reactor Experiment.

dynamic advanced reactor performance requires development of models and tools for representing unique materials, geometries, and physical phenomena.

METHODS & CODES

We developed two new simulation tools to enable advanced reactor and fuel-cycle simulations. First, the finite-element-based physics application Moltres [1,2] couples the thermal-hydraulics and neutronics of molten-salt flow in high-temperature liquid-fueled reactor designs. Second, we have developed a Python package for modeling fuel-salt reprocessing, SaltProc [3].

SaltProc includes fission product removal, fissile material separations, and refueling for time-dependent analysis of fuel-salt evolution. It relies on full-core high-fidelity Monte Carlo simulations to perform depletion computations that require significant computational time and memory.

Moltres is a collection of physics kernel extensions and material definitions for the highly scalable, fully implicit, Multiphysics Object-Oriented Simulation Environment (MOOSE) framework from Idaho National Laboratory [4]. These physics kernels enable Moltres to solve arbitrary-group neutron diffusion, temperature, and precursor governing equations in up to three dimensions and on an arbitrary number of processing units.

Moltres is built on top of the MOOSE framework, which itself leans on the LibMesh [5] finite element library and the PETSc [6] toolkit for solving systems of discretized partial differential equations (PDEs). MOOSE and LibMesh handle translation of Moltres-defined weak PDE forms into residual and Jacobian functions that PETSc solves simultaneously via Newton-Raphson routines. All codes use MPI for parallel communication and are easily deployed on Blue Waters. MOOSE applications like Moltres use monolithic and implicit methods that are ideal for closely coupled and multiscale physics.

Finally, these simulations are memory-intensive, so MOOSE employs a hybrid shared and distributed memory parallel model. The exceptional memory capability of the Blue Waters resource has been essential to performant simulation times.

RESULTS & IMPACT

Current interest in advanced nuclear energy systems and molten-salt reactor (MSR) concepts has illuminated a need for tools that model these systems. By developing such applications

in the open, ARFC enables both transparency and distributed collaboration on promising nuclear reactor concepts. Detailed spatially and temporally resolved neutron fluxes, temperature distributions, and changing isotopic compositions can improve designs, help characterize performance, inform reactor safety margins, and enable validation of numerical modeling techniques for unique physics.

ARFC has conducted simulations of the Molten-Salt Reactor Experiment (MSRE) and the conceptual Molten-Salt Breeder Reactor (MSBR) with deterministic multiphysics and Monte Carlo methods, respectively. Steady-state, transient, and fuel-cycle analysis simulations have been run in 2D as well as 3D by leveraging the Moltres and SaltProc tools developed in our group. Future work may include similarly challenging materials and geometries such as those in sodium-cooled, gas-cooled, and very-high-temperature reactor designs, which also boast improved safety and sustainability.

The model problem presented in Fig. 1 has a 3D-cuboidal geometry with heterogeneous group constants for fuel and moderator regions. Fuel-volume fraction and fuel-salt composition are based on the MSRE design. Neutron fluxes show expected cosinusoidal shapes in radial and axial directions with visible striations between fuel and moderator regions.

Meanwhile, Fig. 2 demonstrates fuel-cycle dynamics obtained from depletion and SaltProc reprocessing simulations for a 10-year time frame. An MSBR full-core safety analysis was performed at the initial and equilibrium fuel-salt compositions for various reactor safety parameters such as effective multiplication factor (shown), neutron flux distributions, temperature coefficients, rod worths, power, and breeding distributions [7,8].

WHY BLUE WATERS

Simulations that faithfully capture this coupling at realistic spatial and temporal resolution are only possible with the aid of high-performance computing resources. To assess nuclear

reactor performance under a variety of conditions and dynamic transients, the ARFC group must conduct myriad 2D and 3D 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 & DATA SETS

Lindsay, A., G. Ridley, A. Rykhlevskii, and K. Huff, Introduction to Moltres: An Application for Simulation of Molten Salt Reactors. *Annals of Nuclear Energy*, 114 (2018), pp. 530–540.

Lindsay, A., and K. Huff, Moltres: Finite Element Based Simulation of Molten Salt Reactors. *The Journal of Open Source Software*, 3:21 (2018), DOI:10.21105/joss.00298.

Ridley, G., A. Lindsay, M. Turk, and K. Huff, *Multiphysics Analysis of Molten Salt Reactor Transients*, (2017), pp. 0–12.

Ridley, G., A. Lindsay, and K. Huff, An Introduction to Moltres, an MSR Multiphysics Code. *Transactions of the American Nuclear Society*, (2017):

Lindsay, A., K. Huff, and A. Rykhlevskii, arfc/moltres: Initial Moltres release. *Zenodo*, (2017), DOI:10.5281/zenodo.801823.

Rykhlevskii, A., A. Lindsay, and K.D. Huff, Full-Core Analysis of Thorium-Fueled Molten Salt Breeder Reactor Using the SERPENT 2 Monte Carlo Code. *Transactions of the American Nuclear Society*, (2017).

Rykhlevskii, A., A. Lindsay, and K.D. Huff, Online Reprocessing Simulation for Thorium-Fueled Molten Salt Breeder Reactor. *Transactions of the American Nuclear Society*, (2017).




Figure 2: Simulated dynamics of reactor criticality (multiplication factor) in the Molten-Salt Breeder Reactor with SaltProc simulation of Rb, Sr, Cs, and Ba online reprocessing over 3,435 days.

LEARNING LOW-DIMENSIONAL FEATURE DYNAMICS OF TURBULENT FLOWS USING DEEP CONVOLUTIONAL RECURRENT AUTOENCODERS

Allocation: Exploratory/50 Knh
PI: Maciej Balajewicz¹
Collaborator: Francisco Gonzalez¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Dynamical systems are used to describe the rich and complex evolution of many real-world processes. Modeling the dynamics of physical, engineering, and biological systems is thus of great importance in their analysis, design, and control. For systems with existing models based on first principles, high-fidelity solutions are possible through direct numerical simulations. However, these generally yield sets of equations with approximately 10^{6-9} degrees of freedom. Even with recent advances in computational power, solving these high-fidelity models is computationally intractable for multiquery and time-critical applications such as design optimization, uncertainty quantification, and model predictive control. Moreover, some systems may have an abundance of data but could lack the governing laws necessary for accurate modeling. Motivated by this problem, we seek to develop deep learning-based model-reduction approaches, wherein both the identification and evolution of low-dimensional features are learned from numerical and experimental data sets.

RESEARCH CHALLENGE

In recent years, the rise of machine learning and big data have driven a shift in the way complex spatiotemporal systems are modeled. The abundance of data has facilitated the construction of so-called data-driven models of systems lacking high-fidelity governing laws. In areas where high-fidelity models do exist, data-driven methods have become an increasingly popular approach to tackle previously challenging problems wherein solutions are learned from physical or numerical data [1].

In model reduction, machine learning strategies have recently been applied to many remaining challenges, including learning stabilizing closure terms in unstable proper orthogonal decomposition (POD)–Galerkin models and data-driven model identification for truncated generalized POD coordinates [2,3]. A more recent approach considers learning a set of observable functions spanning a Koopman invariant subspace from which low-order linear dynamics of nonlinear systems are modeled [4]. While many of these approaches show great promise, a number of significant issues remain. Notably, the issue of scalability arises when considering training deep neural networks on large-scale simulation data.

To avoid this curse of dimensionality, we instead propose a deep learning method that combines important innovations in dimensionality reduction and arbitrary dynamics modeling to perform robust deep learning-based model reduction. First, a deep neural network architecture called a convolutional autoencoder (Fig. 1a) is used to learn low-dimensional, abstract features of the high-dimensional input data. A modified version of a long short-term memory (LSTM) recurrent neural network (Fig. 1b) is then used to learn the *a priori* unknown dynamics of these features. Both networks are trained jointly in an end-to-end fashion, resulting in a completely data-driven model that offers significant advantages over both linear model reduction strategies and vanilla implementations of neural-network-based model-reduction approaches.

METHODS & CODES

The model proposed in this work consists of a four-layer strided convolutional encoder network followed by a two-layer dense encoder network, both of which learn at each layer a lower-dimensional abstract representation of the input data. This results in a low-dimensional feature vector, $h \in \mathbb{R}^k$, which can be thought of as a nonlinear, location-invariant generalization of the generalized POD coordinates. To efficiently evolve these features, the LSTM network is designed to scale with the reduced dimension, relying only on the current state to make a future state prediction. Both the LSTM and encoder networks use the decoder network, which consists of a two-layer dense network followed by a four-layer strided convolutional transposed network, to reconstruct the full state from the low-dimensional feature vector.

The code used in this project is written in Python using TensorFlow, Google’s open source library for building, training, and serving deep neural network models, which utilizes the CUDA CuDNN deep learning library for acceleration with NVIDIA GPUs [5]. The model is trained using ADAM, a variant of stochastic gradient descent.

RESULTS & IMPACT

We trained our deep convolutional recurrent autoencoder model on a number of illustrative examples. Here, we restrict our attention to the problem of a statistically stationary lid-driven cavity flow at a high Reynolds number. In particular, the Navier–

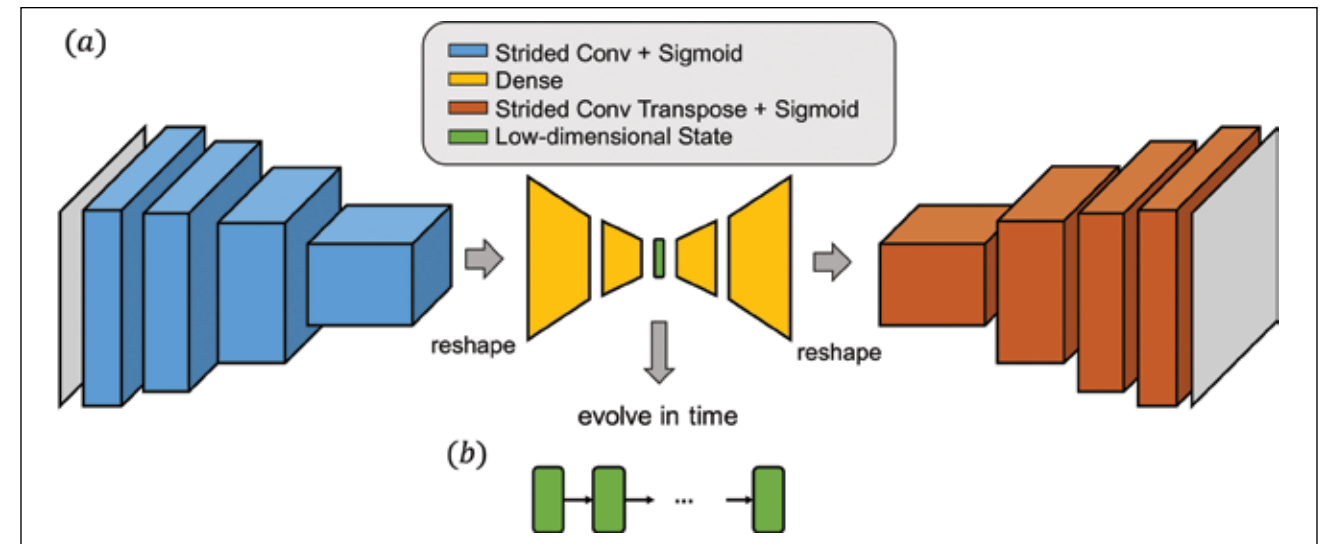


Figure 1: (a) The deep convolutional autoencoder model proposed in this work. In this model, a 2D input snapshot is fed on the left and is processed down to a low-dimensional state (or feature vector). (b) The long short-term memory (LSTM) then evolves this feature vector, where the full state is reconstructed with the decoder portion of the convolutional autoencoder.

Stokes equations in streamfunction–vorticity formulation are discretized in space using Chebyshev polynomials and integrated in time using a semi-implicit second order scheme at a Reynolds number of $Re=2.5 \times 10^4$. At such high Reynolds numbers, this flow exhibits complex spatiotemporal behavior and serves as a well-known benchmark for validation of reduced-order models.

Next, we assembled a data set of finite-time solution snapshot sequences where each snapshot consists of feature-scaled streamfunction fluctuations around the temporal mean. At every step in the training procedure, every individual snapshot is reduced to a low-dimensional feature vector and reconstructed,

and the evolution of the feature vector is compared against the current optimal compression. Fig. 2 compares final evolved output feature-scaled streamfunction fluctuations and the corresponding u-velocity and vorticity fields at three different stages during training using a feature vector $h \in \mathbb{R}^k$, with $k=8$. With no training, the model simply outputs noise from the random initialization of the model parameters. At 7,000 training steps, the model begins to learn the evolution of the system, and by 600,000 training steps the model nearly captures the exact solution.

In this work, we have successfully demonstrated the feasibility of using deep neural network architectures for learning and evolving low-dimensional features of high-dimensional systems through the example of a high-Reynolds-number lid-driven cavity flow. The incorporation of machine and deep learning strategies for constructing smarter and more efficient reduced-order models is still a nascent field, but it is one that could have a significant impact on areas ranging from design optimization, uncertainty quantification, and model predictive control. To this end, we are pursuing a number of different directions including incorporating physics-based constraints to learning low-order dynamics, and learning low-order feature dynamics from heterogeneous data sets.

WHY BLUE WATERS

Training deep neural network models is an inherently data-intensive process. With larger simulations and more sophisticated sensors, there is no shortage of data from which deep learning models of physical systems can be trained. The petascale resources available via Blue Waters, and in particular its large number of GPU-equipped nodes and fast shared parallel storage, have made developing and training deep neural network-based reduced-order models possible.




Figure 2: Evolution of the model output feature-scaled streamfunction fluctuations (top row), its corresponding u-velocity field (middle row), and vorticity field (bottom row) at three different stages during training: (a) no training, (b) after 7,000 training steps, and (c) after 600,000 training steps. The exact solutions are shown in (e).

DI

ELECTRON DYNAMICS OF ION-IRRADIATED GRAPHENE

Allocation: Illinois/550 Knh
PI: André Schleife¹
Co-PI: Alina Kononov¹

¹University of Illinois at Urbana-Champaign

TN

EXECUTIVE SUMMARY

Two-dimensional materials like graphene possess a variety of unique properties desirable for diverse applications. However, developing novel technologies and devices often depends on high-resolution techniques for imaging and manipulating these promising materials. Such techniques include focused ion beams, making it critically important to understand the response of 2D materials to ion irradiation. Although the response of bulk materials to ion irradiation has been well-studied, a comprehensive description of the response of thin materials and material surfaces is lacking. We investigated the dependence of charge transfer, energy transfer, and secondary-electron yield on graphene thickness and ion species, charge, velocity, and trajectory when an ion traverses graphene. Our results provide unprecedented insight into these dynamical processes, and inform the improvement of ion-beam techniques for thin materials.

FS

BI

RESEARCH CHALLENGE

The high conductivity, transparency, mechanical strength, flexibility, and chemical stability of graphene make it a promising candidate for a wide variety of applications, including solar cells, supercapacitors, transparent electrodes, flexible electronics, water desalination, and more [1,2]. However, its properties are extremely sensitive to defects, nanopores, functionalization, and other types of nanostructure, whether intentional or not [3,4]. Thus, developing graphene-based devices requires efficient, reliable, and scalable imaging and patterning techniques for high-resolution characterization and manipulation of its nanostructure.

A variety of materials science imaging and processing techniques involve focused beams of charged particles interacting with materials. Typical microscopy techniques detect secondary electrons emitted by the sample after excitation by energetic particles. On the other hand, typical patterning techniques use ion beams to eject atoms from the sample and/or implant defects, modifying the nanostructure. Since the response of the material depends strongly on the charge, mass, and velocity of the ions used, these parameters must be optimized for the desired application.

Optimizing ion beam parameters for imaging and processing graphene and other two-dimensional materials requires a detailed understanding of the ultrafast electron-ion dynamics occurring during ion irradiation of these materials. As a first-principles study of the electronic response of graphene to irradiation by a variety of ion species across a range of velocities, this project provides unprecedented insight into the ultrafast electron-ion dynamics occurring during the interaction of graphene with various types of ion beams. Our work not only informs ion beam techniques for graphene but also lays the foundation for systematic prediction of the most effective imaging and processing methods for other two-dimensional materials and their heterostructures. This ultimately will accelerate the development of innovative technologies based on these materials, which could revolutionize personal electronics and the energy industry.

METHODS & CODES

We use Qbox/Qb@ll [5], our highly parallel implementation of Ehrenfest molecular dynamics and real-time, time-dependent density functional theory [6], to perform accurate first-principles simulations of excited-electron dynamics. This approach treats nuclei as classical point charges interacting electrostatically with electrons. Electrons are treated quantum-mechanically; their quantum orbitals, represented in a plane-wave basis, are

governed by the Kohn-Sham equations. We use the common adiabatic local-density approximation for the exchange-correlation potential that describes the quantum correction to the electron-electron interaction.

The simulations generate a time-dependent electron density, which is further analyzed to extract the secondary electron yield and the charge transferred to the projectile, among other quantities. Fig. 1 shows a visualization of the change in electron density, or the local charge, in the graphene about 0.5 femtoseconds after impact by a 25 keV proton.

Starting with the lowest energy configuration of the material as the initial condition, the electronic orbitals are propagated through time by numerically integrating the time-dependent Kohn-Sham equations. We tested several different integrators, concluding that enforced time-reversal symmetry (ETRS) is the most efficient option, proving to be more stable and much more accurate than the popular fourth-order Runge-Kutta method. We also laid the groundwork for the future exploration of better integrators for our code. We acknowledge support from the National Science Foundation under Grant No. OAC 17-40219.

RESULTS & IMPACT

We examined the dependence of the secondary electron yield, charge transfer, and energy transfer on projectile type, velocity, and trajectory, allowing us to evaluate the suitability of certain parameter combinations for imaging two-dimensional materials. For example, as illustrated in Fig. 2, we found that secondary electron yield remains nearly constant for 25–80 keV protons impacting graphene. Since energy transfer increases with proton energy in this energy range, these findings suggest that proton

energies of 25 keV and below are optimal for imaging applications, where high secondary-electron emission coupled with low damage to the material are desirable.

The data we obtained informs the controlled application of ion-beam techniques, which promise higher resolution and versatility than traditional tools for imaging and nanostructuring of graphene and other two-dimensional materials. Our approach enables predicting parameter combinations that are suitable for nondestructive imaging of graphene or for intentionally introducing defects in thin materials. By simultaneously studying charge transfer, secondary electron emission, and their ultrafast dynamics, our work also advances the understanding of the fundamental physics occurring during the interaction of a material with ion irradiation.

WHY BLUE WATERS

Blue Waters enabled us to conduct the long simulations of large systems involved in this project. To accurately model the secondary electron emission from a 2D material under irradiation, we must have the capability to evolve thousands of electrons over thousands of timesteps in an elongated simulation cell containing a large vacuum outside the material. These aspects make the simulations computationally expensive; they are only possible with a massively parallel implementation of the first-principles approach and a high-performance supercomputer. Blue Waters allowed us to compute detailed information about more than 20 setups with different graphene thicknesses and projectile species, charges, velocities, and trajectories.




Figure 1: Visualization of local charge in graphene about 0.5 femtoseconds after a proton impacts the material. Blue indicates negative charge while red indicates positive charge.




Figure 2: Visualization of local charge in graphene about 0.5 femtoseconds after a proton impacts the material. Blue indicates negative charge while red indicates positive charge.

MAGNETO-OPTICAL KERR EFFECT OF ANTIFERROMAGNETIC MATERIALS IN EXTERNAL MAGNETIC FIELDS

Allocation: Blue Waters Professor/245 Knh

PI: André Schleife¹

Co-PI: Kisung Kang¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

This work investigates the relationship between light and magnetically ordered materials, and particularly metallic antiferromagnets. Their weak response to applied external magnetic fields and terahertz frequency spin dynamics are interesting characteristics that make antiferromagnetic materials potential candidates for memory devices. In this context, the Magneto-Optical Kerr Effect (MOKE) is a useful phenomenon that provides information about magnetization through optical response; however, antiferromagnetic materials with a compensated spin configuration have no net magnetization. Based on the magnetic susceptibility, the spin response under applied fields can be predicted and related to tilted spin configurations that arise in external magnetic fields. MOKE signals are, hence, present for tilted spin configurations. In this project, we carried out first-principles calculations based on density functional theory to study MOKE and the magnetic susceptibility of antiferromagnetic materials. This will be useful guidance for experimentalists; e.g., to identify the wave length of maximum MOKE response in new materials.

RESEARCH CHALLENGE

Antiferromagnetic materials have interesting properties that can potentially lead to fast switching in future memory devices. Unlike ferromagnetic materials, they show a weak response to external magnetic fields, and do not show a magnetic field outside of the material due to spin compensation. However, this can become a double-edged sword. Due to the weak response under applied external fields, it is difficult to control the internal spin configuration. Spin compensation makes detecting the spin information of antiferromagnetic materials using optical measurements challenging. Measurements of optical response under external magnetic fields have been developed. However, it is still necessary to understand the fundamental behavior of antiferromagnetic materials better. In order to understand basic principles an experiment can benefit from theoretical guidance, e.g., to maximize the detected signal, in order to more easily measure the response.

METHODS & CODES

We used first-principles calculations based on density functional theory to understand the magnetic properties of antiferromagnetic materials using the Vienna Ab-Initio Simulation Package (VASP)

[1]. Based on magnetic structures from neutron scattering, we performed structural relaxations in the presence of magnetism. We explicitly studied spin-tilted cases with noncollinear magnetism, based on a fully relativistic band structure and dielectric tensor [2]. MOKE signals were computed using the Kerr equation and the off-diagonal elements of the dielectric tensor. At the same time, we computed magnetic susceptibility from the relationship between total energies and net magnetic moments of spin-tilted magnetic unit cells. We further studied Néel temperatures based on magnon dispersion calculations using the spin spiral approximation and the frozen magnon model.

RESULTS & IMPACT

In this work, we used first-principles calculations to study magneto-optical properties of various antiferromagnetic metals. MOKE signals measured in experiment are usually small on the nanoradian scale, which is easily overwhelmed by noise. Since it

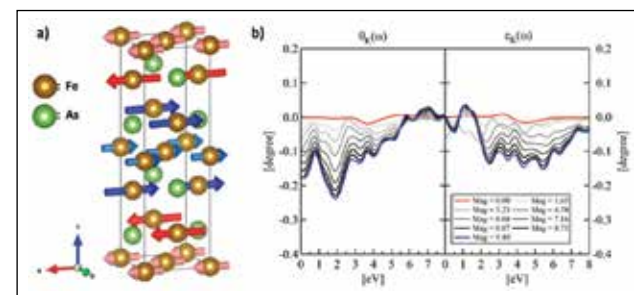


Figure 1: (a) Magnetic unit cell of Fe_2As , and (b) Wave-length dependence of MOKE signals under external magnetic fields. Kerr rotation (θ_K) and Kerr ellipticity (ϵ_K) are maximized at 2 eV (620 nm) and 2.4 eV (520 nm), respectively.




Figure 2: (a) Magnetic unit cell of MnPt; (b) magnetization and energy change of MnPt spin-tilted calculation for corresponding external magnetic fields. Red marker line in (b) presents ground states under corresponding fields and (c) magnon dispersion of MnPt.

is difficult to tune the laser system every time to find the optimal wavelength for excellent signal-to-noise ratios, first-principles calculations based on density functional theory can provide this information. In this project, wavelength-dependent MOKE spectra are computed for MnPt and Fe_2As . We find that MnPt shows a maximum Kerr rotation signal at 730 nm and Kerr ellipticity at 540 nm. For Fe_2As , Kerr rotation and ellipticity are maximized at 620 nm and 520 nm, respectively. These spectra give useful guidance to experimentalists, who seek it to find the wavelength that maximizes MOKE signals.

At the same time, we also computed magnetic susceptibility, which can provide the spin response under applied external magnetic fields. For MnPt along the c -axis, which is perpendicular to the spin direction, computed magnetic susceptibility is $\chi = 4.4 \times 10^{-4}$ which is close to $\chi = 4.742 \times 10^{-4}$, as determined experimentally [3]. Based on this information, magnetic moments of spin-tilted magnetic-structure calculations can be converted to a response to external magnetic fields. Thus, MOKE signals under the applied field can be predicted.

In order to confirm the spin configuration stability at room temperature, we performed Néel temperature calculations. Through the mean-field approximation, the computed Néel temperature of MnPt is 928.6 K, while experimental measurement of MnPt finds 953 K [3]. The results of magnetic susceptibility and Néel temperature calculations illustrate that this first-principles computational approach can provide reasonable predictions for magnetic properties of antiferromagnetic materials. Thus, our calculations are helpful guidance to understand the behavior of antiferromagnetic materials and to develop advanced magnetic devices. For this work, we acknowledge funding through the Illinois Materials Research Science and Engineering Center, which is supported by the National Science Foundation MRSEC program under NSF Award Number DMR-1720633.

WHY BLUE WATERS

Even though first-principles calculations are a great opportunity to thoroughly investigate antiferromagnetic materials, optical and magnetic calculations still require significant computational resources. Unlike nonrelativistic cases, a fully relativistic band structure and its dielectric function calculation with magnetism are demanding processes. Spin-tilted calculations are challenging because noncollinear calculations need more computational resources than collinear calculations. In addition, these require more computational resources because magnetic unit cells of antiferromagnets are usually larger than chemical unit cells. Magnon dispersion calculations for Néel temperatures also require dense Brillouin zone sampling to achieve convergence to an appropriate accuracy. In order to address those challenges, Blue Waters is well suited due to its fast communication and large amount of memory per node. Thus, Blue Waters provides a unique chance to unveil the unknown properties of antiferromagnetic materials.

PETASCALE SIMULATION OF HIGH-ENERGY-DENSITY (HED) PLASMAS

Allocation: NSF PRAC/4,050 Knh

PI: Warren Mori¹

Co-PI: Frank S. Tsung¹

Collaborators: L. O. Silva², R. A. Fonseca²

¹University of California, Los Angeles

²Instituto Superior Técnico

EXECUTIVE SUMMARY

The UCLA Plasma Simulation Group has been using Blue Waters since its inception to apply our own suite of kinetic simulation tools, including the particle-in-cell (PIC) code OSIRIS and the Vlasov–Fokker–Planck (VFP) code OSHUN to probe fundamental physics relevant to current HED (high-energy-density) plasma experiments, including those at the Facility for Advanced Accelerator Experimental Tests at the U.S. Department of Energy’s (DOE) SLAC National Accelerator Laboratory and the National Ignition Facility at Lawrence Livermore National Laboratory. In the past 12 months, the results of these petascale simulations have appeared in high-impact journals such as *Physics of Plasmas* and *Physical Review Accelerators and Beams*. Blue Waters has allowed the UCLA Simulation Group to perform petascale simulations in a timely manner and to train students in HED physics and high-performance computing.

RESEARCH CHALLENGE

The goals of the UCLA Plasma Simulation Group and its collaborators continue to be focused on three key areas funded by NSF and DOE. The key research questions are:

- 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 interactions be controlled or even harnessed in inertial fusion energy-relevant 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

In the above problems, the systems are highly nonlinear and cannot be easily described by fluid models, and kinetic effects can play important roles. Therefore, particle-in-cell (PIC) models are ideally suited for their study, where Maxwell’s equations are solved on a grid using current and charge densities, and the particles’ orbits are calculated using Newton’s laws. The UCLA Plasma Simulation Group and its collaborators at IST in Portugal maintain a large number of PIC codes, including OSIRIS, QuickPIC, and

UPIC. These codes are all developed locally, share many of the same algorithms and data structures, and have been optimized for heterogeneous leadership-class supercomputers such as Blue Waters. All of these codes are open access, and both QuickPIC and UPIC are open source (on GitHub).

RESULTS & IMPACT

Using the tools described here, we continue to perform large-scale simulations to study the various issues facing current experiments in plasma-based accelerators and inertial confinement fusion. In the past 12 months, we have performed very large 3D and quasi-3D simulations of laser wakefield accelerators (LWFAs) in plasma down-ramps. Our simulations showed that the plasma down-ramp can decelerate the energetic electrons in the transverse direction, leading to a beam with very low emittance. Furthermore, the density down-ramp has a chirp in the accelerating gradient that compensates the energy chirp in the accelerated electrons, leading to an overall low energy spread. Our 3D simulations show that in a density down-ramp, electrons with ultra-high brightness ($> 10^{20} \text{A}^{-2} \text{rad}^{-2}$) and low overall energy spread ($< 1 \text{MeV}$) can be generated using either an electron beam or a laser pulse driver, which makes this a promising candidate to drive X-ray free-electron lasers (FELs) with nanometer wavelengths [1].

Furthermore, 2D OSIRIS simulations with realistic laser beam optics [including beam smoothing techniques such as SSD (smoothing by spectral dispersion) or ISI (induced spatial incoherence)] have shown that, for sufficiently large temporal bandwidth (where the frequency bandwidth of the laser beam is comparable to the growth rate of the laser–plasma instability), laser–plasma instabilities (LPI) can be reduced under conditions relevant to current inertial fusion energy experiments [2]. These simulations are very relevant to ongoing experiments, yet at the same time, they explore fundamental kinetic effects in plasma physics, making them ideal for training students and postdocs.

WHY BLUE WATERS

The UCLA Plasma Simulation Group has been a user on Blue Waters since the very beginning. Blue Waters has continued to provide a very stable high-performance platform for the study of kinetic effects in high-energy-density plasmas. This stability has allowed us to perform a large number of petascale simulations

that will help experimentalists produce brighter X-ray sources using X-FEL (using LWFAs) and produce higher-yield targets in inertial fusion experiments.

PUBLICATIONS & DATA SETS

Xu, X.L., et al., High quality electron bunch generation using a longitudinal density-tailored plasma-based accelerator in the three-dimensional blowout regime. *Phys. Rev. Accel. Beams*, 20 (2017), p. 111303.

Wen, H., F.S. Tsung, B.J. Winjum, A.S. Joglekar, and W.B. Mori, Kinetic Simulations of Reducing Stimulated Raman Scattering with Temporal Bandwidth in Inertial Confinement Fusion. In preparation (2018).

Wen, H., Petascale kinetic simulations of laser plasma interactions relevant to inertial fusion—controlling laser plasma interactions with laser bandwidth. *The 3rd International Conference on Matter and Radiation at Extremes*, Qingdao, China, May 2018.




Figure 1: Spatial distribution of beam intensity from 3D OSIRIS simulation of laser–plasma interactions relevant to inertial confinement fusion. The pattern mimics the effects of optical elements that are present in current fusion experiments and is important for achieving quantitative agreement among simulations and experiments.

RESOLUTION EFFECTS AND EXTREME EVENTS IN PETASCALE TURBULENCE SIMULATIONS

Allocation: NSF PRAC/3,400 Knh

PI: Pui-Kuen Yeung¹

Co-PI: K.R. Sreenivasan²

Collaborator: S.B. Pope³

¹Georgia Institute of Technology

²New York University

³Cornell University

EXECUTIVE SUMMARY

Substantial advances have occurred in both domain science and computing for fluid turbulence since access to the Blue Waters resource began. Recent analyses have focused on a critical examination of the effects of insufficient or limited resolution in both space and time for the study of intermittency, which often is expressed via intense but localized fluctuations. Two important indicators of the fine-scale turbulence structure are the dissipation rate and enstrophy (the integral of the vorticity), both of which tend to exhibit extreme events at thousands of the mean value or more. Recently, we have obtained a new understanding of how the probability distributions of these two variables compare, both of which are of stretched-exponential form but with different parameters. Further, we have obtained some data for short periods of time at resolutions beyond 1 trillion grid points.

RESEARCH CHALLENGE

Intermittency, or the occurrence of intense fluctuations localized in time and space, is a fundamental but still not well-understood property of turbulence at high Reynolds number that is used to help predict flow patterns in different fluid flow situations and arises in numerous fields of science and engineering [1]. For an intermittent flow variable, such as energy dissipation rate, a basic question is how large the fluctuations can be, with what likelihood, and whether the dependence on Reynolds number can be quantified reliably. In principle, reliable answers can be obtained from direct numerical simulations of the instantaneous turbulent flow in simplified geometries, such as periodic domains in three-dimensional space. While attaining a high Reynolds number is always desirable, theoretical models have also suggested a need to resolve the small scales better than is often practiced [2].

Although direct numerical simulations based on exact equations for the basic physical laws of conservation of mass and momentum are, in principle, closest to the truth, the physical fidelity of the results is sensitive to the effects of errors of a numerical or statistical nature. In addition, it is unavoidable that, in each given simulation, some statistics are highly accurate while others are less reliable. For example, high accuracy for extreme events in the energy dissipation rate are very difficult to achieve if the Reynolds number is to be at least moderately high. A critical examination of

requirements for accuracy is also especially timely in the transition toward the next generation of multi-petaflop architectures.

METHODS & CODES

In our simulations we integrated the Navier–Stokes equations over a large number of timesteps using Fourier pseudospectral methods in space and finite differencing in time. The probability distribution of fluctuations of energy dissipation rate (quadratic measure of strain rates) and enstrophy (likewise, of the rotation rate), can be computed by postprocessing of large data sets archived on mass-storage systems, or calculated on the fly at many time-instants during the simulation. We have found the latter approach to be beneficial for the purpose of studying high-intensity, small-scale fluctuations that evolve on very short time scales. Besides computing the probability distribution, which can be averaged in time, we also extracted peak values (of intermittent variables), which serve as a simple diagnostic of how strong any extreme events can be.

If a particular quantity in a given simulation is significantly contaminated by numerical errors, it is expected to change substantially when a new and larger simulation at improved resolution is performed. However, before conducting a more expensive simulation with finer grid spacing or shorter timesteps, it is useful to first identify which quantities are more problematic. We have developed an approach to obtain this information by progressively filtering out spectral content at high wavenumbers where aliasing errors associated with nonlinear terms in the equations of motion may not be fully removed.

RESULTS & IMPACT

We have recently conducted a detailed study [3] of the effects of errors due to insufficient spatial and temporal resolutions using a combination of filtering and short simulations where statistics of dissipation and enstrophy are sampled very frequently. Although dissipation is less intermittent than the enstrophy, results show conclusively that, because of the role of incompressibility, dissipation is more sensitive to errors caused by insufficient or limited resolution. The expectation that higher resolution allows larger gradients (hence more intense dissipation and enstrophy) to be captured is clearly demonstrated only if the resolution in time

is sufficiently good to suppress the influence of aliasing errors. Because of the advective nature of momentum transport, the proper measure of temporal resolution is not in terms of the time scales of the small-scale motion, but instead of the Courant number, which is usually interpreted as a criterion for numerical stability of the time-integration procedure.

The observation that dissipation is affected by resolution effects more than enstrophy suggests past statements [4,5] concerning comparative behaviors of extreme events in these quantities may require revision. In [3] we came to the conclusion that the probability density functions (PDFs) of these two variables do not coincide in the range where extreme events occur. However, these two PDFs do possess a remarkable degree of similarity, with both being well described by stretched-exponential fitting functions differing in only one parameter. This result is seen in Fig. 1, where the PDF of normalized enstrophy agrees very closely with that of twice the dissipation for reasons that can be investigated using the theory of multifractals [6]. Furthermore, a parallel investigation [7] suggests that the time scale of the extreme events in dissipation rate behaves similarly as a power law in the Reynolds number.




Figure 1: Probability density functions of normalized dissipation (red) and enstrophy (blue) in isotropic turbulence, at higher spatial and temporal resolution than in previous work. Dashed lines show very close fits with stretched exponentials. Line in green shows results for twice the dissipation.

WHY BLUE WATERS

The highest grid resolutions reported in [4,5,7] were all at the level of $8,192^3$. Although not included in our original project plans, we have also obtained some results at $12,288^3$ (which exceeds 1 trillion grid points) with good resolution in time. Further analyses are expected to lead to more publications.

PUBLICATIONS & DATA SETS

Yeung, P.K., X.M. Zhai, and K.R. Sreenivasan, Extreme events in computational turbulence. *Proceedings of the National Academy of Sciences*, 112 (2015), pp. 12633–12638.

Iyer, K.P., K.R. Sreenivasan, and P.K. Yeung, Reynolds number scaling of velocity increments in isotropic turbulence. *Physical Review E*, 95 (2017), p. 021101R.

Yeung, P.K., K.R. Sreenivasan, and S.B. Pope, Effects of finite spatial and temporal resolution on extreme events in direct numerical simulations of incompressible isotropic turbulence. *Physical Review Fluids*, 3 (2018), p. 064603.

Buaria, D., A. Pumir, E. Bodenschatz, and P.K. Yeung, Extreme velocity gradients in turbulent flows. Submitted to *Physical Review X* (2018).

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 proposal provided the computing resources to four graduate students, funded by the NCSA Materials & Manufacturing (M&M) group, to explore how their research can be furthered through the use of advanced digital technology to address large-scale problem solving. Two of the four research activities conducted via the allocation are reported below.

The first project, titled “Simulation of Reference Point Indentation on Cortical Bone,” was conducted by Ashraf Idkaidek. He used two different instruments that utilize the Reference Point Indentation technique—BioDent and Osteoprobe. The second project—“Mechanics of Materials with a Focus on Accelerated Design and Structure-Processing-Property Relations of Materials via High-Scale Computations”—was conducted by Fereshteh A. Sabet. In this work, he investigated and compared the performance of implicit and explicit solvers for trabecular (inner layer) bone using Abaqus.

Project 1—Simulation of Reference Point Indentation on Cortical Bone

RESEARCH CHALLENGE

Osteoporosis is a bone disease responsible for two million broken bones and \$19 billion per year in related costs in the United States, alone. By 2025, osteoporosis is expected to be responsible for three million fractures and \$25.3 billion in related costs per year, according to the U.S. National Osteoporosis Foundation. Assessing the relationship between bone fracture resistance and bone material properties is important for the diagnosis and treatment of bone diseases. Using traditional material testing approaches such as compression, tension, or three- and four-point bending to measure bone mechanical properties is *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 whole-bone fracture resistance. The Reference Point Indentation (RPI) technique was invented to allow *in vivo* evaluation of bone properties. There are two instruments that

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 (using the commercial Abaqus software) is highly nonlinear where geometric nonlinearity, material nonlinearity, and contacts are needed to be accounted for to preserve the accuracy of the simulation results.

RESULTS & IMPACT

We have related each of the 10 outputs of the BioDent RPI instrument to bone material properties by using the finite element method [1]. We have also evaluated the simulation of bone fracture using the extended finite element method on a single-osteon cortical bone sample [2,3]. Further, we are currently developing a study to relate Osteoprobe RPI output to bone material properties and fracture resistance.

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 relate Osteoprobe device output to different bone mechanical properties. Each of the Osteoprobe RPI simulation iterations demands high computational power and time. Therefore, completing such a study using multicore Blue Waters clusters is essential.

PUBLICATIONS & DATA SETS

Idkaidek, A., and I. Jasiuk, Toward high-speed 3D nonlinear soft tissue deformation simulations using Abaqus software. *J Robot Surg*, 9:4 (2015), pp. 299–310.

Idkaidek, A., and I. Jasiuk, Cortical bone fracture analysis using XFEM—case study. *Int J Numer Method Biomed Eng*, 33:4 (2017), DOI:10.1002/cnm.2809.

Idkaidek, A., S. Koric, and I. Jasiuk, Fracture analysis of multi-osteon cortical bone using XFEM. *Computational Mechanics*, 62:2 (2018), pp. 171–184.

Project 2—Mechanics of Materials with a Focus on Accelerated Design and Structure-Processing-Property Relations of Materials via High-Scale Computations

RESEARCH CHALLENGE

Bone has a hierarchical architecture spanning from atomistic to macroscopic scales. At the scale of one to a few millimeters, the bone tissue is composed of cortical (outer) and trabecular (inner) bone. Osteoporosis is a bone disease characterized by low bone density, which often leads to an increased risk of fractures that mainly occur 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 the design of improved implant systems [1,2].

METHODS & CODES

Modeling of trabecular bone entails a highly nonlinear mechanical behavior along with contacts. This leads to increased ill-conditioning of global stiffness matrices and difficulties converging, especially in the post-yield regime [3]. As a result, it is of considerable interest to assess the effectiveness and efficiency of an explicit solution method. In this project, we used the implicit and explicit solvers of Abaqus to analyze nonlinear 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, by using a similar setup for the model (e.g., element type, loading type, etc.) when using implicit and explicit solvers, there is a perfect match between micro-CT FE model results using implicit and explicit solvers. Fig. 1 shows an example of stress–strain curves obtained from nonlinear micro-CT FE modeling of trabecular bone using implicit and explicit solvers along with experimental results. As can be observed, with the explicit solver we are able to reach a higher applied strain without convergence difficulties. There is also a good match between local stresses obtained using implicit and explicit solvers, as shown in Fig. 2. In addition, we observed that implicit and explicit solvers scale similarly, while 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 reduced computational time. Each of our models has many millions degrees of freedom and nonlinearities, making them impossible to solve without the use of the Blue Waters supercomputer.

PUBLICATIONS & DATA SETS

Sabet, F.A., O. Jin, S. Koric, and I. Jasiuk, Nonlinear micro-CT based FE modeling of trabecular bone—Sensitivity of apparent

response to tissue constitutive law and bone volume fraction. *International journal for numerical methods in biomedical engineering*, 34:4 (2017), p. e2941.

Sabet, F. A., S. Koric, and I. Jasiuk, A comparison between implicit and explicit finite element simulations of trabecular bone. In preparation (2018).




Figure 1: Comparison of apparent response resulting from implicit and explicit solvers (Project 2).




Figure 2: Comparison of contact pressure obtained using implicit and explicit solvers (Project 2).

MULTIPHASE TURBULENT FLOW MODELING OF GAS INJECTION INTO MOLTEN METAL TO MINIMIZE SURFACE DEFECTS IN CONTINUOUS-CAST STEEL

Allocation: Illinois/75 Knh

PI: Brian G. Thomas¹

Co-PIs: Seong-Mook Cho², Surya Pratap Vanka¹, Seid Koric^{1,3}, Ahmed Taha³, Hyunjin Yang¹, Matthew Zappulla²

¹University of Illinois at Urbana-Champaign

²Colorado School of Mines

³National Center for Supercomputing Applications

EXECUTIVE SUMMARY

This project aims to develop comprehensive, sophisticated, and computationally intensive models to simulate transient multiphysics phenomena. In addition, we aim to understand defect formation mechanisms in continuous steel casting and to find practical ways to improve the process, which could greatly impact the steel industry. We have developed a new hybrid multiphase flow model to calculate argon gas behavior and bubble size distribution during the continuous casting process. In addition, we have simulated argon bubble transport and capture into the steel shell using Large-Eddy Simulations (LES) coupled with a Discrete Phase Model (DPM) for particle transport and capture. Furthermore, we have investigated the effect of moving magnetic fields on transient steel–argon flow patterns, surface velocity and level, and argon bubble distribution in the mold using LES simulations coupled with DPM and the magnetic induction magnetohydrodynamics (MHD) model. These simulations have been validated with lab-scale models and plant measurements and applied to reveal deeper insights into defect formation, enabling the investigation of optimal process conditions.

RESEARCH CHALLENGE

Continuous casting is used to manufacture more than 96% of steel worldwide [1]. Many defects in final steel products are related to complex multiphysics phenomena during continuous casting, including turbulent multiphase flow, particle transport and capture, magnetohydrodynamics, heat transfer, solidification, and thermal–mechanical behavior, as shown in Fig. 1a and 1b. This process is extremely difficult to study with lab-scale models and plant experiments owing to the hostile environment of the molten steel and the many process variables, including the thermal properties of steel and slag, process geometries, and process conditions. Thus, this work applies transient multiphase flow simulations on Blue Waters to quantify argon gas behavior, bubble size distribution, and bubble transport and capture, which influence bubble defect formation. We then studied the effect of moving magnetic fields on the flow pattern in order to investigate ways to reduce defects related to entrapped bubbles and inclusions.

METHODS & CODES

We developed a new hybrid multiphase flow model [2,3] that couples the Eulerian–Eulerian two-fluid model approach together with DPM and applied it to simulate bubble transport and behaviors, including gas pocket formation, breakup, coalescence, and size distribution evolution in the nozzle and mold. In addition, we used LES coupled with DPM and MHD models to simulate transient molten steel–argon bubble flow with and without an Electromagnetic Level Stabilizer (EMLS) moving field. We conducted these simulations using the commercial CFD program, ANSYS Fluent High-Performance Computing (HPC) on Blue Waters XE nodes. To calculate bubble transport and capture into solidifying steel shells in the mold, we implemented the LES–DPM–MHD model coupled with the advanced force balance (on each bubble at the solidification front) capture criterion model [4,5] into Blue Waters XK nodes with the multi-GPU based in-house code CUFLOW [6].

RESULTS & IMPACT

We simulated the complex gas behavior including gas pocket formation, shearing off of the gas pocket into bubbles, volumetric expansion, breakup, coalescence and evolution of the bubble size distribution in the steel continuous caster using the hybrid multiphase flow model, which has been validated by comparison

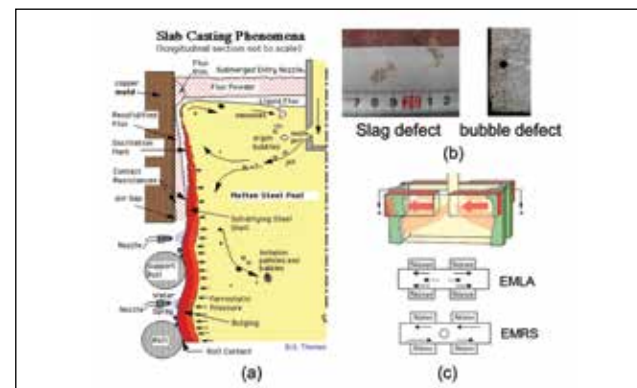


Figure 1: (a) Schematic of multiphase flow phenomena in continuous casting (CC) mold, (b) defects in CC, and (c) electromagnetic level accelerating (EMLA) and rotating stirring (EMRS) systems.

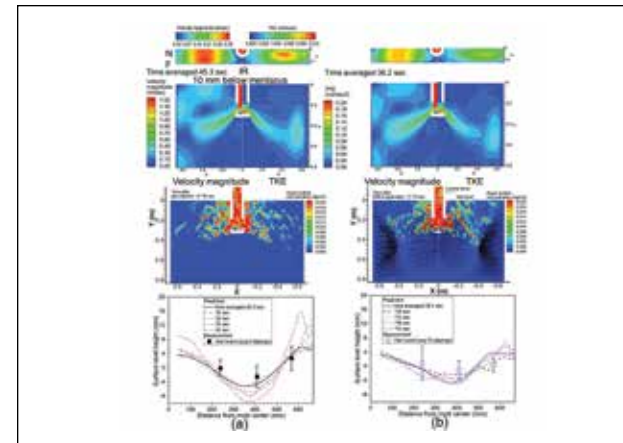


Figure 2: Multiphase flow patterns, argon bubble distribution, and surface level variations in continuous casting (CC) mold (a) without and (b) with EMLS (Electromagnetic Level Stabilizer).

with several lab-scale experiments [7]. The model was able to simulate realistic argon bubble distributions in the molten steel pool in the continuous caster, which enabled more accurate prediction of bubble defects, including their size and location in the steel product.

Argon bubble transport and capture into the steel shell in the mold during continuous slab casting with and without double-ruler Electromagnetic Braking (EMBr) were simulated using LES–DPM–MHD models with particle capture. The models showed good agreement with measurements of size distribution and location of captured bubbles in the steel slabs. The models also revealed the effect of bubble size on residence time of the bubbles in the mold, with and without double-ruler EMBr, which is an important factor influencing bubble defect formation that is being quantified in this work [4,5].

Furthermore, we quantified transient flow patterns, surface velocity and level, and argon bubble distribution in the mold, with and without an EMLS moving magnetic field, using the LES–DPM–MHD model as shown in Fig. 2. The LES models showed good agreement with measurements of surface velocity, surface level profiles, and their fluctuations from nail dipping tests conducted in the commercial caster. The EMLS reduced up-and-down wobbling of the steel jet and made the jet deflect slightly downward due to the strong electromagnetic forces induced near the narrow face. Thus, the jet flow had a longer path toward the mold top surface, resulting in lower and more stable surface velocity and level, which can reduce slag-entrapment defects (Fig. 1b). In addition, the strong Lorentz forces near the narrow faces braked the jet and transported the argon bubbles farther away from the steel shell, which decreased the chance of bubble-capture defects.

The validated transient multiphase flow models on Blue Waters can be applied in parametric studies to find optimum process conditions to reduce defects. This could enable significant savings in the cost of steel production and improve product quality.

WHY BLUE WATERS

Development of accurate multiphysics models of steel continuous casting for a more detailed understanding of defect formation mechanisms and improvement of the process requires great computational resources to be accurate. Specifically, to capture the complex and interrelated phenomena on a micrometer scale, the large domain size (over 1 m³) and long timescale of some of the macroscale flow phenomena (up to a minute) require many computational cells (~100 million) and many timesteps, owing to the very small timestep size (smaller than ~10⁻⁴ seconds). In the current work, we achieved an over 3,000 times faster calculation with ANSYS FLUENT HPC on Blue Waters XE nodes, compared to an ordinary workstation PC. Furthermore, the multi-GPU in-house CUFLOW codes accomplished a good parallel scalability on Blue Waters XK nodes, showing only 48 hours of wall-clock time for 30 seconds of fully developed LES–MHD flow in a 14-million-cell domain. Thus, the parallel supercomputing environment on Blue Waters is greatly contributing to accurately quantifying the complicated multiphysics phenomena with high resolution in order to improve understanding of defect formation in this complex commercial process.

PUBLICATIONS & DATA SETS

Jin, K., S.P. Vanka, and B.G. Thomas, Large Eddy Simulations of Electromagnetic Braking Effects on Argon Bubble Transport and Capture in a Steel Continuous Casting Mold. *Metallurgical and Materials Transaction B*, 49B (2018), pp. 1360–1377.

Yang, H., S.P. Vanka, and B.G. Thomas, A Hybrid Eulerian–Eulerian Discrete-Phase Model of Turbulent Bubbly Flow. *Journal of Fluids Engineering*, 140:10 (2018), pp. 101202–101212.

Yang, H., S.P. Vanka, and B.G. Thomas, Hybrid Eulerian Eulerian Discrete Phase Model of Turbulent Bubbly Flow, *Proceedings of ASME IMECE 2017*.

Yang, H., S.P. Vanka, and B.G. Thomas, Modeling of Argon Gas Behavior in Continuous Casting of Steel. *Proceedings of TMS 2018* (Phoenix, Ariz., March 11–15, 2018).

Zappulla, M.L.S., and B.G. Thomas, Thermal–Mechanical Model of Depression Formation in Steel Continuous Casting. *146th Annual Meeting & Exhibition Supplemental Proceedings, The Minerals, Metals & Materials Series* (Springer 2017), pp. 501–510.

Cho., S.-M., et al., Effect of EMLS Moving Magnetic Field on Transient Slab-Mold Flow. *CCC Annual Report* (2017).

Kim, H.-S., S.-M. Cho, S.-H. Kim, and B.G. Thomas, Mold Flow and Shell Growth in Large-Section Bloom Casting using CON1D. *CCC Annual Report* (2017).

Yang, H., S.P. Vanka, and B.G. Thomas, Hybrid Model of Multiphase Flow Including Gas Pockets, Gas Bubbles, Breakup and Coalescence. *CCC Annual Report* (2017).

Zappulla, M.L.S., and B.G. Thomas, Surface Defect Formation in Steel Continuous Casting. *Proceedings of THERMEC* (2018), accepted.

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

Allocation: Illinois/200 Knh
PI: Rafael Tinoco Lopez¹
Co-PI: Paul Fischer¹
Collaborators: Som Dutta¹, Pallav Ranjan¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Aquatic vegetation exerts significant influence in the hydrodynamics of both fluvial and coastal systems. It extracts energy from the flow, which has been found to impede near-bed flow velocities, to modify the structure of the velocity field, and consequently influence sediment erosion and resuspension. Aquatic vegetation influences different processes to such an extent that vegetation has been earmarked as a key component influencing aquatic ecosystems. Our study is geared toward increasing our understanding of the interactions among vegetation, flow, and sediment. We have conducted Direct Numerical Simulations (DNS) and Large Eddy Simulations (LES) through different arrays of idealized vegetation, represented as cylinders, using the higher-order spectral element-based computational fluid dynamics (CFD) solver Nek5000. We simulated different arrangements of cylinders, and the high-fidelity turbulent simulations have shed light on details of the flow features, increasing our understanding of different hydro- and morphodynamic processes.

RESEARCH CHALLENGE

Aquatic vegetation is known to stabilize its environment, thus engineering different aspects of its own ecosystem [1]. Aquatic vegetation provides a wide range of ecosystem services [2], ranging from an increase in water quality through nutrient uptake and oxygen production to creation of habitats through spatial heterogeneity of the flow velocity to dampening erosion.

Past computational studies have mainly focused on using CFD models, based on Reynolds Averaged Navier–Stokes equations (RANS), which only provide an averaged approximation of the flow-field. The few studies that have used high-fidelity LES have been limited by the number of vegetation elements. Most of the existing data in this area come from laboratory experiments [3], recreating conditions closer to nature but often lacking the spatial and temporal resolution required to capture some fundamental processes in detail.

The current study is geared to complement experiments being conducted at the Ven Te Chow Hydrosystems Laboratory at the University of Illinois at Urbana-Champaign. The study conducts numerical simulations at an unprecedented scale, resolving details that, in conjunction with the experiments, is providing unforeseen insights into the fundamental dynamics of flow and transport in the presence of aquatic vegetation [4]. These large-scale computations will not only help improve lower-order models of the processes but will also help inform better experimental design and measurement practices.

The scale of the experimental setup to be modeled is a paramount 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. This results in a reduction in the number of cases that can be run, thus constraining the insights a broader range of parameters could yield.




Figure 1: Instantaneous velocity magnitude for flow through random arrays of idealized aquatic vegetation at Reynolds number $\sim 10,000$. The arrays have the same volumetric frontal area but different cylinder diameters: (a) 1/8 inch, (b) 1/4 inch, (c) 1/2 inch, and (d) 1 inch. Despite having the same frontal area, the number of high-velocity regions is substantially higher for larger diameters, which also show a more heterogeneous velocity distribution.




Figure 2: 3D simulation of turbulent flow through a staggered array of emergent cylinders with about 150 million computational points. (a) Plots the instantaneous velocity magnitude at three planes. (b) Visualizes the continuous shedding of vortices using the isosurface of the negative eigenvalue of the vorticity field.

To increase variable space, a wide range of conditions are first simulated in two dimensions for the whole domain (~ 6 million computational points). This gives an overview of the effect of different parameters such as Reynolds number, vegetation density, and spatial heterogeneity. For a limited set of conditions, we are conducting three-dimensional simulations (~ 200 million computational points) for only a part of the domain. The conditions to simulate in 3D are informed by the 2D simulations, where the 3D domains are big enough to accurately capture the general dynamics within manageable computational costs.

METHODS & CODES

High-resolution LES and DNS of the flow at different configurations of the idealized vegetation are conducted using the open-source, spectral element-based, high-order incompressible Navier–Stokes solver Nek5000 [5,6]. In the simulations with sediment, sediment transport is modeled under the Eulerian framework using the advection-diffusion equation [7], making our study one of the first to look at the complex interaction between sediment-flow and vegetation using high-fidelity CFD simulations.

RESULTS & IMPACT

We conducted 2D simulations for the full experimental domain for different Reynolds numbers and varying random configurations of vegetation elements for random arrays having the same volumetric frontal area but different cylinder diameters (see Fig. 1). It is evident that, even though the volumetric frontal area is the same, different diameters result in different porosities, thus increasing the heterogeneity in the flow with an increase in diameter of the vegetation elements. We monitored drag on the vegetation arrays and observed that, for the same Reynolds number, the drag in the transverse direction increases with an increase in diameter of the vegetation elements.

We also conducted 3D simulations for different Reynolds numbers, ranging between 12,000 and 20,000, for a staggered array of vegetation elements (see Fig. 2). The simulated setup is similar to experiments currently being conducted at the laboratory. We conducted simulations using up to 200 million computational points, running on 16,384 processors at a time. The simulation

results show clear evidence of the dynamics of the flow being driven primarily by the vegetation elements, with clear preferential flow-paths through the array of elements. We visualized the vortices shedding by the cylinders using the isosurface of the negative eigenvalue of the gradient velocity tensor that is decomposed into its symmetric and antisymmetric parts [8]. As expected, the vortex tubes being generated scale to the diameter of the vegetation elements. As part of the analysis, we are investigating the turbulent kinetic energy and turbulence intensity in order to link them to variations in bed shear stresses and intermittent turbulent events altering the sediment transport capacity of the flow.

WHY BLUE WATERS

The study pushes the limits of the scale at which high-resolution simulations are used to study complex multiphase flow in environmental fluid mechanics, requiring computational resources with sustained computing power at an unprecedented scale, such as Blue Waters. We have conducted simulations for up to 200 million computational points, with the code scaling strongly up to 16,384 MPI ranks. Without access to petascale high-performance computing, it would be impossible to complete the simulations within a realistic timeframe. In addition, since visualization of a phenomenon is an effective way to understand and explain its mechanics, we will work with the Blue Waters project staff to create animations of the phenomenon using data from the simulations.

PUBLICATIONS & DATA SETS

Dutta, S., et al., 2D High-Resolution numerical investigation of flow through surrogate emergent vegetation canopies: Part 1. Unidirectional Flow. In preparation for *Environmental Fluid Mechanics* (2018)

Dutta, S., et al., 2D High-Resolution numerical investigation of flow through surrogate emergent vegetation canopies: Part 2. Oscillatory Flow. In preparation for *Environmental Fluid Mechanics* (2018).

Ranjan, P., et al., Flow and sediment transport through aquatic vegetation. In preparation for *Geophysical Research Letters* (2018).

HIGH-FIDELITY NUMERICAL SIMULATIONS OF THE REACTING FLOW FIELD OF AN ANNULAR COMBUSTOR MODEL OF AN AIRCRAFT ENGINE

Allocation: Director Discretionary/100 Knh

PI: Ahmed Taha¹

Co-PIs: Seid Koric¹, Daniel Mira²

Collaborators: Oriol Lehmkuhl², Ricard Borrell², Albert Coca², Samuel Gómez², Mariano Vázquez², Guillaume Houzeaux², Madhu Vellakal¹

¹National Center for Supercomputing Applications

²Barcelona Supercomputing Center

EXECUTIVE SUMMARY

Our research is dedicated to performing high-fidelity numerical simulations of conditions representative of aeronautical combustors from turbofan engines using wall-resolved Large-Eddy Simulation (LES). The challenges associated with the modeling of these systems correspond to the full characterization of the reacting layers, the near-wall treatment, and the ability to predict pollutant emissions. This project is defined to show the aerospace and aeronautical community that high-performance computing (HPC) codes can be used to perform real-engine simulations, and can go one step beyond the state-of-the-art in modeling large-scale combustion systems for propulsion and power applications.

RESEARCH CHALLENGE

The challenges associated with the modeling of these systems correspond to the full characterization of the reacting layers, the near-wall treatment, and the ability to predict pollutant emissions.

METHODS & CODES

The code used for the project is called Alya [1], which is a computational fluid dynamics code of the PRACE Benchmark Suite for HPC applications. It has been highly optimized and independently tested on most European supercomputing platforms. The Alya code was developed at the Barcelona Supercomputing Centre and is a parallel multiphysics software package using the finite element method to run applications on HPC facilities. For combustion problems, Alya showed excellent scalability on the Blue Waters supercomputer (Fig. 1—bottom) to 100,000 cores on meshes of up to 4.22 billion elements. The algorithms and models employed to solve the individual physical problems have been assessed and validated in previous research [1–4].

The governing equations describing the reacting flow field correspond to the low Mach number approximation of the Navier–Stokes equations with the energy equation represented by the total enthalpy. The chemical state is described by two controlling variables: the mixture fraction and the reaction progress [3]. The combustion process is assumed to take place in the flamelet regime for all computing cases, and we employ a turbulent combustion model based on presumed probability density function [3,4]. The modeling approach that we follow in this project is valid for partially premixed flames, as this approach can recover the structure of premixed flames in regions dominated by flame propagation, as well as the structure of diffusion flames in the postflame regions [3,4].

RESULTS & IMPACT

The main results of this project are manifold as shown below and in Fig. 2:

- Development of an efficient computational framework to perform large-scale simulations of reacting flows under aero-engine-like conditions.

- Numerical characterization of a technically premixed swirling flame.
- Numerical characterization of both a piloted and nonpiloted turbulent nonpremixed jet flame.

The primary impact of the project is the development of a computational framework that can be used to study complex combustion systems based on high-fidelity simulations of turbulent reactive flows.




Figure 2: Snapshots of turbulent flames—(a) nonpremixed jet flame, (b) technically premixed jet flame, and (c) nonpremixed piloted jet flame.

WHY BLUE WATERS

The full characterization of the reacting layers needs a very fine mesh to capture all of the flow details. In addition, achieving convergence at reasonable timescales requires running the code at scale. Using Blue Waters makes conducting the parametric study with such large meshes and complex physical modeling possible in a timely manner.

PUBLICATIONS & DATA SETS

Borrell, R., et al., Parallel mesh partitioning based on space filling curves. *Computers & Fluids*, in press (2018), DOI:10.1016/j.compfluid.2018.01.040.

Sandia D flame (LES data)

DLR Jet flame A (LES data)

Swirling combustor with central axial jet injection (LES data)




Figure 1: Scaling tests on MareNostrum (top) and Blue Waters (bottom).

FULLY THREE-DIMENSIONAL KINETIC SIMULATIONS OF UNSTEADY SHOCK-BOUNDARY LAYER INTERACTIONS

Allocation: Illinois/284 Knh
 PI: Deborah Levin¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Hypersonic flow over double-wedge configurations displays what is known as complex Edney type IV interactions. Whether they are modeled using continuum or particle methods, these flows are unsteady in nature. However, modeling of such flows provides an opportunity to study flow stability mechanisms and identify near-transition behavior. Given the importance of continuum breakdown in strong shocks, a kinetic particle treatment of such flows is crucial. However, the continuum-like free stream conditions create 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 geometry of the double wedge also presents challenges with respect to load imbalance among the processors, which can cause significant increases in communication time.

RESEARCH CHALLENGE

Hypersonic compressible flows are characterized by significant gradients in gas density, high temperatures, and a large degree of nonequilibrium. Simulation of such flows requires due consideration of shocks, the chemistry of diatomic molecules, and the interactions of gases with embedded surfaces. 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. These have a significant role in determining 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 and provides the highest fidelity in strong shock regions. 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. Since the DSMC approach is amenable to parallelization, however, these limitations can be overcome with an efficient use of state-of-the-art petascale resources such as Blue Waters and computational tools that have been designed for these resources.

METHODS & CODES

We developed a three-dimensional MPI-parallelized DSMC code known as Scalable Unstructured Gasdynamic Adaptive mesh Refinement (SUGAR). SUGAR 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

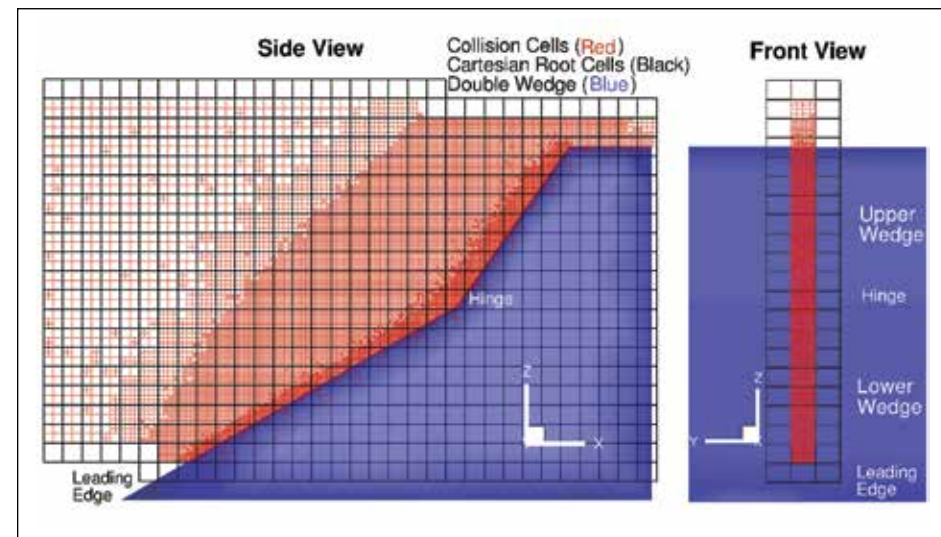


Figure 1: Decomposition of the computational domain held by a single of the 32 processors. Collision and root grids are shown in red and black, respectively. Root grid elements that are blank represent roots held by neighboring processors. The processor needs the location code array data held in both filled and blank roots.




Figure 2: SWBLI features shown based on gradient density magnitude in the symmetry plane based on flow field parameters sampled from ~0.1 to 0.108 ms.

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. We have performed a great deal of work to improve the scalability of the code [1] and have achieved the highest performance in the world for such problems: 87% weak scaling for 8,192 processors for flow over a hemisphere with 24 billion particles.

RESULTS & IMPACT

Satisfying the DSMC requirement for a SWBLI case at a Knudsen number of the order of 10^3 requires billions of computational particles and cells, which necessitates very sophisticated memory optimization tactics. The results for the actual case presented here were simulated using 20,000 processors, a number that would only be possible on Blue Waters. Furthermore, the gradients in the flow are so severe that there are many AMR/octree roots and each can have seven levels of refinement. This requires 8 MB of memory per processor to store the location code array that correlates the position of DSMC particles with leaf nodes. The memory requirement is made more tractable by storing this array for only those roots that lie in the portion of the domain assigned to the processor and the neighboring roots, as shown in Fig. 1.

The unsteady simulated Schlieren of the complex shock structure formed in the symmetry plane on the double wedge at ~0.1 ms is shown in Fig. 2. We can see that an oblique and bow shock are formed at the leading edge of the lower wedge

and in front of the upper wedge, respectively. As the flow moves up the lower wedge, an adverse pressure gradient is formed and the boundary layer separates, thus generating a separation shock. These three shocks meet at the triple point and with their interaction form a reflected shock, also known as the trailing shock. As the reflected shock reaches closer to the boundary layer, the compression is nearly isentropic with a continuous transition to subsonic velocities. At the reattachment location of the boundary layer, a succession of compression waves are formed that coalesce into a secondary shock. As the flow passes through the oblique, separation, and reflected shocks, the increase in density and temperature results in a reduction in the thickness of the boundary layer downstream of the reattachment location. We observed strong recirculation in the separation zone embedded between the hinge, separation point, and the reattachment point. 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 although the velocities and temperatures differ.

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, 24 billion computational particles, and 2.2 billion collision cells to reach one millisecond. Blue Waters is one of the few computer architectures that can host and execute these simulations.

MAPPING PROTON QUARK STRUCTURE—LOOKING INSIDE THE PROTON: HOW DO QUARKS SPIN?

Allocation: NSF PRAC/9,440 Knh

PI: Caroline Riedl¹

Co-PIs: Matthias Perdekamp¹, Naomi Makins¹

Collaborators: Vincent Andrieux¹, Robert Heitz¹, Riccardo Longo¹, Angelo Maggiora², Artem Petrosyan³

¹University of Illinois at Urbana-Champaign

²National Institute for Nuclear Physics (Torino, Italy)

³PNPI (Moscow, Russia)

EXECUTIVE SUMMARY

COMPASS (COMmon Muon Proton Apparatus for Structure and Spectroscopy) is a high-energy physics experiment that probes proton substructure by scattering high-energy pion and muon beams off of nuclear targets at CERN in Geneva, Switzerland. The experiment explores 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 and 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 the detailed simulation of COMPASS detector properties and environmental effects.

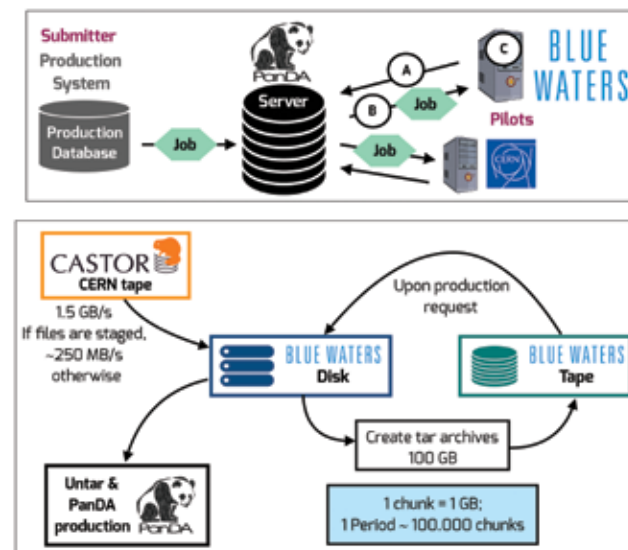


Figure 1: Top—work flow of COMPASS data production using PanDA. Bottom—management of raw experimental data on Blue Waters. Transfers between CERN and BW are handled by FTS3; transfers within BW by Globus Online.

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 is one of the few Nuclear Science Advisory Committee [1] milestones for U.S. Department of Energy- and National Science Foundation-funded research in nuclear physics. (Sivers functions arise from correlations between proton spin and quark transverse momentum and thus appear connected to quark orbital motion inside the proton.) 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 [2]: the negatively charged pion beam from the Super Proton Synchrotron was impinged on a target of transversely polarized protons.

With the 2016 and 2017 Generalized Parton Distributions (GPD) runs, COMPASS has added valuable observables in Deeply Virtual Compton Scattering (DVCS) to constrain spin-independent GPDs in the so-far unexplored kinematic domain between HERMES and the Jefferson Lab experiments on the one hand, and the HERA collider experiments on the other. The global community of model-fitters awaits the future COMPASS results in DVCS.

METHODS & CODES

For experimental data production, about 1 petabyte of raw COMPASS data collected at CERN has been transferred to Blue Waters. The average throughput speed is up to 1.5 GBs using the File Transfer System FTS3 [3], a bulk data mover created to distribute the Large Hadron Collider (LHC) data globally. More data—about 4 petabytes—will be transferred in 2018. The data are tared and stored on tape; upon production request, they are retrieved from tape (see Fig. 1, bottom).

For each triggered event in COMPASS, the information of the detectors is recorded by the Data Acquisition system. The COMPASS Reconstruction Analysis Library (CORAL) software performs the transition from raw data information to physical quantities. CORAL's function is to reconstruct particle trajectories and momenta, as well as 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 CORAL and PHAST reconstruction codes as experimental data.

RESULTS & IMPACT

COMPASS accumulates a raw experimental data set of about 1 petabyte 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. Approximately 6.5-million CPU hours and 50 days at the CERN computing cluster are needed for one annual production pass. Our team has adapted the COMPASS version of PanDA (Production ANd Distributed Analysis) [5], a data production and monitoring system developed for ATLAS–LHC, to Blue Waters (see Fig. 1, top). We expect to be able to process an annual COMPASS data set on Blue Waters within five days. Long-term, we are planning on four annual data sets, each with two passes.

Apart from the processed experimental data, simulated Monte Carlo (MC) data are an essential ingredient of the data analysis. Simulations of the detectors play a central role in understanding subtle detector effects and in removing background events from the data sample. With the available resources at CERN and collaborating institutions, the CPU-intensive part of the

MC calculation—the simulation of the detector properties with GEANT—often cannot be afforded for extensive studies; for example, event pile-ups or time-dependent detector efficiencies. Two examples of MC activities are shown in Fig. 2: an environmental simulation that led to a modification in the experiment's shielding setup, and the simulation of the physics signal and its backgrounds for the Drell–Yan analysis.

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 and is a step toward building a community capable of using petascale computing.

WHY BLUE WATERS

With the petascale resources of Blue Waters, COMPASS experimental and Monte Carlo data can be processed significantly faster (10 to 25 times faster compared to other computing resources available to COMPASS), which will allow the completion of publications and PhD theses in a timely manner. In the case of simulations, the data can also be generated in greater detail, delivering the high precision that keeps systematic uncertainties at the smallest possible levels. An example is the realistic pile-up of events in the particle collision. Blue Waters enables novel explorations; for example, detector resolutions in kinematic binnings and two-dimensional detector efficiency maps. In addition, Blue Waters staff provide essential guidance in terms of job flow, load distribution, data transfer, and BW-specific technical features.

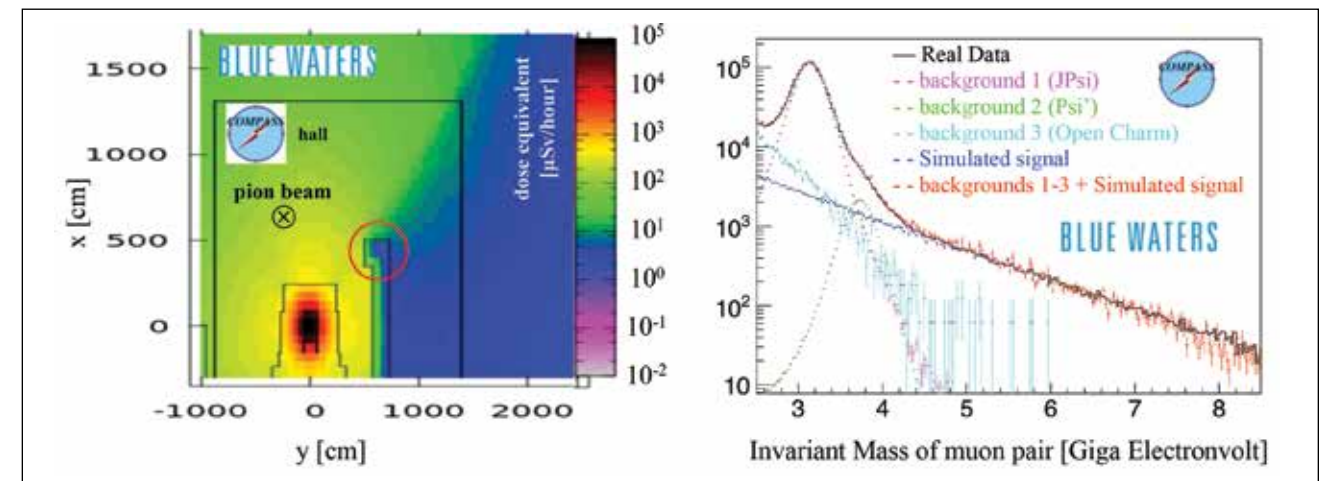


Figure 2: Left—BW FLUKA [6] simulation of hourly radiation dose in the COMPASS experimental hall. Additional concrete shielding blocks (red circle) were installed to satisfy CERN's radiation-protection requirements. Right—BW PYTHIA [7] simulation of muon pairs reconstructed with the COMPASS spectrometer in comparison with the experimental data.

ENABLING DISCOVERY AT THE LARGE HADRON COLLIDER THROUGH DATA-INTENSIVE COMPUTATION AND MACHINE LEARNING

Allocation: Illinois/140 Knh
PI: Mark Neubauer¹
Collaborators: Robert Gardner², David Lesny¹, Dewen Zhong¹

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

EXECUTIVE SUMMARY

The Large Hadron Collider (LHC) is the world’s most powerful particle accelerator, designed to study the fundamental nature of matter and the forces that govern its interactions by colliding beams of protons at the highest-available energies. We are using Blue Waters to process, simulate, and analyze high-energy proton–proton collision data produced by the ATLAS experiment at the LHC and to improve our sensitivity to new phenomena by developing novel approaches to identifying Higgs bosons produced with high momentum at the LHC by using machine learning techniques.

RESEARCH CHALLENGE

The goal of particle physics is to understand the universe at its most fundamental level, including the constituents of matter, their interactions, and the nature of space and time itself. This quest is one of the most ambitious and enduring of human endeavors.

The Standard Model (SM) of particle physics describes all known fundamental particles and their interactions, including the Higgs boson, which was discovered at the LHC [1,2] in 2012,

with significant contributions by the Illinois (Neubauer) Group. The discovery led to François Englert and Peter W. Higgs receiving the 2013 Nobel Prize in Physics. The SM has withstood the last 40 years of experimental scrutiny, with important exceptions being neutrino mass, dark matter, and dark energy. Recent developments in particle physics and cosmology raise the exciting prospect that we are on the threshold of a major step forward in our understanding.

It is an enormous challenge to process, analyze, and share the 15 petabytes of data generated by the LHC experiments each year with thousands of physicists around the world. To translate the observed data into insights about fundamental physics, the important quantum mechanical processes, and the detector’s responses to them, need to be simulated to a high level of detail and with a high degree of accuracy.

A key thrust of this project is to use the recently discovered Higgs boson to search for new physics in novel ways enabled by the Blue Waters supercomputer. The enormous energy available in proton–proton collisions at the LHC leads to the production of particles with very high velocity relative to the ATLAS detector

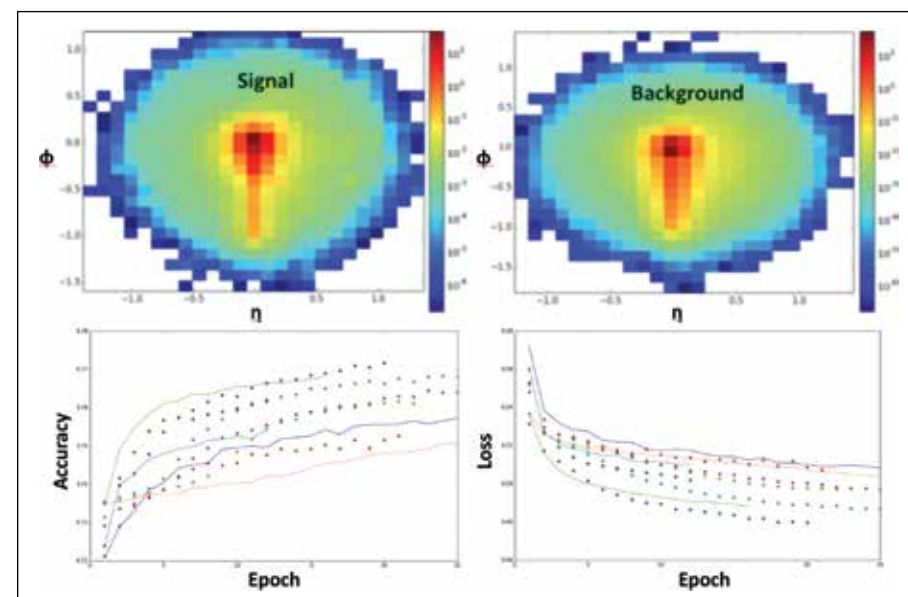
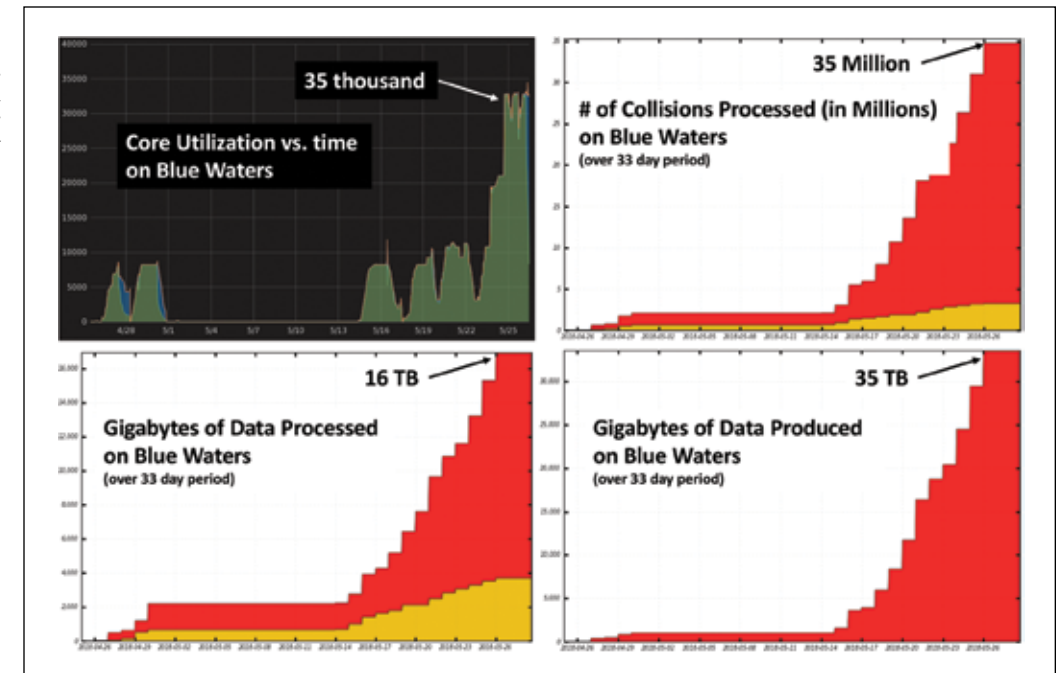


Figure 1: (Top) “Images” created by jets of charged particles in Higgs decay signal events and background events that are detected by the ATLAS tracking system. (Bottom) Higgs boson identification accuracy and signal loss as a function of the number of training epochs for a variety of CNN configurations and hyperparameter settings.

Figure 2: CPU core utilization, number of collision events processed, data consumed and simulated, and refined detector data generated during a 33-day period during the project.



(lab frame). Even massive particles like the Higgs boson can have a large momentum and, therefore, large Lorentz factor (γ) in the lab frame. When these “boosted” particles decay, their decay products are highly collimated and not easily distinguished in the detector instrumentation (e.g., by calorimeters). This limits the sensitivity of searches for new physics such as $X \rightarrow hh$, where X is a new massive ($\sim \text{TeV}/c^2$) particle.

METHODS & CODES

We have integrated Blue Waters into our production processing environment to simulate and analyze massive amounts of LHC data. Blue Waters resources are made available to the ATLAS computing fabric using a system called ATLAS Connect [3], which is a set of computing services designed to augment existing tools and resources used by the U.S. ATLAS physics community. Docker Images are delivered via Shifter to create an environment on Blue Waters’ nodes that is compatible with the ATLAS job payload.

The approach we are currently taking to identify boosted Higgs bosons is to use a convolutional neural network (CNN) trained using “images” created by jets of charged particles in Higgs decay events, as shown in Fig. 1 (top). This work was done in collaboration with Indiana University and the University of Göttingen.

The images shown are the particles’ angles in the detector as centered on the Higgs boson, with the color represented by the particles’ momentum transverse to the proton beamline. We have successfully performed CNN training using GPUs on Blue Waters with these images via a sequential Keras model with a TensorFlow backend.

We have also studied deep neural networks, again using Keras with TensorFlow. We are developing an alternative approach

that uses the particles’ four-momentum and jet-clustering history rather than images fed into a recursive neural network, which draws inspiration from natural language processing.

RESULTS & IMPACT

Fig. 2 shows the CPU core utilization, data consumed and generated, and the number of collision events processed during a 33-day period of the project. The job output was made available to the rest of the ATLAS collaboration for use in analysis of the LHC data to improve measurements of the SM and to search for new physics. Fig. 1 (bottom) shows the Higgs boson identification accuracy and signal loss as a function of the number of training epochs for a variety of CNN configurations and hyperparameter settings. We are also using Hyperas, a convenience wrapper using Hyperopt with Keras models, on Blue Waters to automate the scanning of hyperparameters in a variety of machine and deep learning approaches to improve the Higgs boson identification over backgrounds. The techniques show promise in addressing the challenges of boosted Higgs boson identification and improving the sensitivity of new physics searches at the LHC.

WHY BLUE WATERS

Blue Waters, as a large CPU and GPU resource with high data-throughput capabilities, greatly facilitated our research. The strong support for containers allowed us to deploy our science application on Blue Waters nodes. Also, Blue Waters provided a means for a highly parallelized and automated scanning of free parameters in our machine learning configurations and, therefore, rapid optimization of our boosted Higgs boson identifier.

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, the MIMD Lattice Computation Collaboration, and RBC Collaborations, which among them contain almost all 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 leading roles 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 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), the theory of the strong interactions. The standard model has been enormously successful in explaining a wealth of data over the past 40 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 to obtain a deeper understanding of the standard model and to search for physical phenomena that go beyond it.

METHODS & CODES

QCD is formulated in the four-dimensional spacetime 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 so.

Numerical studies of QCD current use a number of different formulations of quarks on the lattice, 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. Domain wall fermions 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.

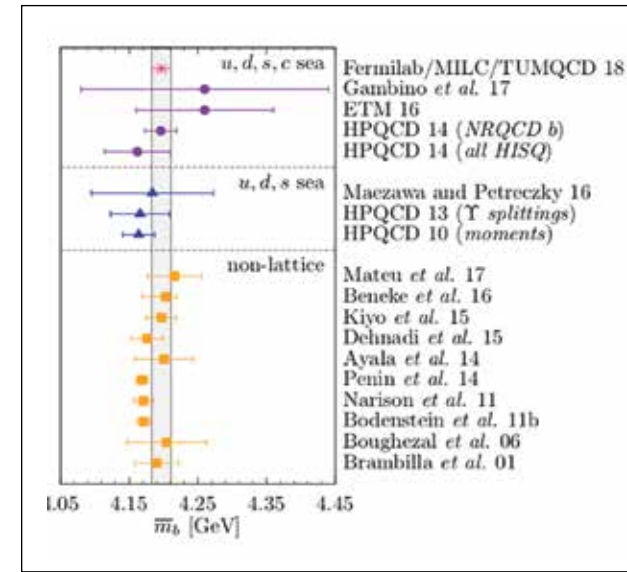


Figure 1: Comparison of our result (magenta burst) for the mass of the bottom quark in the \overline{MS} -bar renormalization scheme to results from other calculations and from nonlattice methods.

RESULTS & IMPACT

We have used our allocation under this grant to generate DWF and HISQ gauge configurations that are among the most challenging produced to date. The initial applications of the DWF configurations are in the study of three processes that are highly suppressed in the standard model and, therefore, favored as places where physics beyond the standard model may emerge: 1) the calculation of the long-distance contributions to the indirect

violation of Conjugation Parity symmetry in the decay of neutral kaons; 2) the calculation of the long-distance contributions to the rare kaon decay of the charged kaon into a pion and a neutrino–antineutrino pair; and 3) the determination of the mass difference between the two neutral kaon-decay eigenstates, which is the smallest particle mass difference ever measured. The first calculation of this final quantity with physical quark masses is now underway using the ensemble to which our Blue Waters calculation is contributing [1].

The HISQ gauge configurations generated under this allocation have been used in the determinations of quark masses [2] and leptonic decay constants [3] of unprecedented precision. In particular, our results for the charm quark mass matches the precision of the most accurate previous calculations, and those for the other quark masses and their ratios are the most precise to date. Similarly, our results for the decay constants of D and B mesons are the most precise to date. Some comparisons with previous calculations are shown in Figs. 1 and 2.

WHY BLUE WATERS

Lattice QCD calculations have made major progress in the last few years with a limited number of calculations reaching precision of a fraction of a percent, and techniques in place to determine many more quantities to this level of accuracy. Such precision is needed to test the standard model and 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.




Figure 2: Comparison of our results (magenta bursts) for the leptonic decay constants of the B^+ and B_s mesons with those from earlier calculations. These quantities play an important role in testing the standard model, and when coupled with experimental results yield some of the basic parameters of the standard model.

EPIGENETIC IDENTIFICATION AND MAPPING USING SOLID-STATE 2D NANOPORES

Allocation: Illinois/418 Knh

PI: Jean-Pierre Leburton¹

Co-PI: Lav Varshney¹

Collaborators: Nagendra Athreya¹, Aditya Sarathy¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Epigenetic modifications, such as heritable alterations of the human genome, are believed to play a critical role in gene regulation, causing diseases such as cancer and various autoimmune and neurological disorders. In the present work, we develop computational techniques to understand and assess the efficiency of epigenetic detection of methylation sites and their mapping on the DNA strands with the use of 2D atomically thin nanopore membranes in electrolytic cells. We consider various detection scenarios involving ionic current blockade of the pore as well as monitoring the transverse electronic current variations across the membrane. The simulated current signatures are obtained by coupling all-atom molecular dynamics (MD) simulations to a combination of self-consistent Poisson-Boltzmann electrostatics and electronic transport calculations. Additionally, to overcome the inherently low signal-to-noise ratio (SNR) during these detections, we have developed statistical signal processing algorithms recognizing and distinguishing DNA nucleotides and various methylated sites on the DNA strands.

RESEARCH CHALLENGE

Aside from sequencing DNA molecules, the identification of traits of the human genome, such as methylation, is crucial for diagnosis of epigenetic diseases. Recent experimental evidence

of DNA methylation alterations linked to tumorigenesis suggests that DNA methylation plays a major role in causing cancer by silencing key cancer-related genes [1–3]. Until now, detection and mapping of such DNA methylation patterns using solid-state nanopores have been unsuccessful due to rapid conformational variations generated by thermal fluctuations that result in low SNR. To overcome these drawbacks, a versatile, general sensor technology for detecting methylation patterns is desirable. For this reason, we propose an integrated approach that combines MD with device physics-based electronic modeling and statistical signal processing techniques to assess the resolution limit of solid-state nanopore sensing. In addition, we further develop algorithms for epigenetic marker classification at the fundamental limits of SNR improvement for biodetecting membranes.

METHODS & CODES

Our research consists of a two-step process that first uses MD simulations with the latest NAMD version, and then exploits the MD data to calculate the current variations due to DNA translocation through the nanopore via electronic transport modeling. The system is built, visualized, and analyzed using VMD [4]. The protein and DNA are described by the CHARMM22 force field with CMAP corrections [5] and the CHARMM27 force




Figure 1: Schematic of the experimental setup for epigenetic detection and mapping with a 2D nanopore membrane. Frozen (ideal) and nonfrozen (noisy) signatures of three different biomarker proteins (right) binding to the methylation sites are considered for use in the matched-filter algorithm.




Figure 2: Normalized correlations of noisy electronic sheet current signals obtained from the matched-filter algorithm. Two signals, one corresponding to single-MBD1 (blue) and the other corresponding to multiple-MBD1 (red), are aligned and the difference between them (green) is used to determine and map the presence of a second protein on the DNA.

field [6], respectively. An external electric field is applied to the system to drive the DNA-protein complex through nanopores.

For each frame of the trajectory files obtained from the MD simulations, ionic current blockade is calculated instantaneously [7]. Further, electrostatic potential induced by the biomolecule around the pore is obtained by using the self-consistent Poisson-Boltzmann equation (PBE) formalism. The PBE is solved numerically using the multigrid method until the convergence criterion is met. Once the electrostatic potential co-planar to the membrane is obtained, the transverse conductance across the 2D nanopore, graphene, or MoS₂ membrane is computed by using the nonequilibrium Green's function formalism [8] or a semiclassical thermionic emission technique [9], respectively. The PBE and the electronic transport code are written and maintained by the Leburton Group at the University of Illinois at Urbana-Champaign.

RESULTS & IMPACT

Previously, we showed the ability of 2D material nanopores to detect DNA methylation sites labeled by MBD1 proteins by two techniques: through the ionic current blockade and transverse electronic sheet conductance. Further, our combined MD device-modeling approach showed that multiple methylation sites could be distinguished in a single ionic current measurement, provided that they are separated by at least 15 base-pairs (bps), whereas single transverse sheet current measurements resulted in a better identification resolution of 10 bps. In the latter case, the superior performance of electronic detection is due to the ability of the single transverse sheet current method to capture local protein charge variation within the membrane nanopore [10,11].

In our scenario, we built systems where the methylated cytosines are complexed by attaching either a methyl-CpG binding domain (MBD-1) protein or a methyl CpG binding protein 2 (MeCP2), whereas methylated adenines are attached to an oligosaccharide, γ -cyclodextrin (γ -GCD) [12]. “Noise-free” electronic currents (ideally obtained by frozen biomolecules artificially translocated through the nanopore) were calculated for all the biomarkers. These signatures were compiled into a set of dictionary signals for each of the marker proteins. When a target noisy signal includes

the stochastic conformational fluctuations obtained from MD simulations and electronic transport calculations whose marker protein is unknown is fed into the matched filter, it is correlated with the different dictionary signals to identify the marker-protein type. We anticipate that this algorithm can be extended to the detection of multiple markers attached to the same DNA molecule.

These biomarkers are important for the recognition of different cancer segments. MBD1 and MeCP2 are proteins in humans that are capable of binding to hypermethylated sites along the DNA strand and that also repress transcription from methylated gene promoters [13]. MeCP2 mutations are thought to be responsible for Rett syndrome, and polymorphisms of MBD1 are associated with increased lung cancer risk. Alternatively, hypomethylation, identified using γ -GCD (a synthetic biomarker for the N6-methyladenine), has been linked to cancers of the stomach, prostate, breast, pancreas, and kidney [14]. Therefore, identification and differentiation of these different proteins are critical, as their interactions with DNA play important roles in breast, lung, and other kinds of cancers.

WHY BLUE WATERS

Investigation of the interactions of biomolecules with solid-state materials, characterization of the stochastic structural fluctuations of the epigenetic biomarker complexed with DNA translocating through solid-state nanopores, and acquisition of the electronic response using all-atom MD simulations coupled with electronic transport calculations are only possible with petascale computing resources such as Blue Waters. Our systems are about 500,000 atoms in size, each requiring multiple MD simulation (NAMD) runs. With NAMD code efficiently deployed on XE/XK nodes to run highly parallel simulations of large biomolecular systems, Blue Waters is well suited for our research needs.

PUBLICATIONS & DATA SETS

Sarathy, A., N.B. Athreya, L.R. Varshney, and J.-P. Leburton, Classification of Epigenetic Biomarkers with Atomically Thin Nanopores. *Journal of Physical Chemistry Letters*, 9 (2018), pp. 5718–5725.

COMPUTATIONAL MODELING OF ROTATING DETONATION ENGINES

Allocation: GLCPC/840 Knh
PI: Venkat Raman¹

¹University of Michigan

EXECUTIVE SUMMARY

We have used the Blue Waters supercomputing resource for full-system simulations of rotating detonation engines (RDEs) with an aim of improving the efficiency of power generation and propulsion systems. The RDE system provides pressure-gain combustion, which increases the efficiency of the thermodynamic cycle. A reliable RDE design can increase cycle efficiency by up to 30% compared to a conventional combustor system. Within the RDE, a continuous detonation wave propagates through a fuel–air mixture. Due to the extreme conditions within the device, numerical simulations are particularly useful in understanding the complex flow physics while minimizing the use of expensive experiments. To this end, we developed a massively parallel compressible flow solver, UMDetFOAM, to provide high-fidelity simulation results to better understand RDE dynamics. We used numerical simulations to understand the dynamical features of the flow field, including the nature of detonation processes.

RESEARCH CHALLENGE

The conventional combustor, based on deflagration flame propagation, has been actively studied in the previous century. While this system is mature and reliable, significant improvements in efficiency are unlikely. Thus, pressure-gain combustion—the mode used by RDEs—has recently gained traction as a next-generation combustor system [1]. RDEs use a detonation-based thermodynamic process that allows for increased thermodynamic efficiency over the traditional Brayton cycle combustor, and is seen as a more robust combustor concept than the previously explored pulse detonation engines (PDEs). In RDEs, a continuous detonation wave propagates through a fuel–air mixture to achieve combustion and increase pressure inside the constant volume, mitigating the pressure losses that occur in conventional engines. As opposed to PDEs, the continuous pressure gain associated with RDEs reduces the complexity of coupling the combustor with a downstream turbine system in realistic applications. Advancements in related technologies, such as high-temperature materials science and flow control systems, have led to an increased interest in detonation-based combustion as a viable method for increasing fuel efficiency and reducing the environmental impact of propulsion systems.

Although a series of experiments has been conducted within the community, the detailed flow physics, such as the nonuniformity effect on the detonation structure and the turbulent mixing process, are unknown [1,2]. In order to augment experimental

work, high-fidelity computational tools need to be developed to simulate the complex physics within these engines and accelerate the design process. Insights gained from the combination of high-fidelity simulations and experiments can then be used to optimize RDE design.

The focus of this study is the effect of fuel injection design on the robustness of the detonation process. In this context, the RDE configuration presents unique numerical challenges. First, a detailed chemical kinetics capability is required to simulate the nonpremixed injection of fuel and oxidizer. Second, the solver must utilize unstructured grids to mesh the complex injection geometries. Third, the presence of detonation fronts requires numerical tools that reduce dispersive as well as dissipative errors. These challenges need to be overcome to perform reliable simulations of the RDE configuration. The goal of this research lies in developing a compressible flow solver for simulating RDEs with complex geometries, creating detailed models for combustion processes, and using high-performance computing to enable simulation-aided design.

METHODS & CODES

We developed an in-house compressible flow solver, UMDetFOAM, for this research. Open-source platforms such as OpenFOAM [3], for its finite volume method, and Cantera, for its chemistry mechanism readers, are core foundations of the solver. We realized two critical achievements through these platforms: (1) the ability to define unstructured grids in a parallel environment; and (2) the ability to utilize any detailed chemistry mechanism of interest. We performed the discretization of the convective terms and the time integration of the Navier–Stokes equations using a Monotonic Upwind Scheme for Conservation Laws-based Harten–Lax–van Leer–Contact scheme and a second-order Runge–Kutta method, respectively. The chemical source terms are handled explicitly through a detailed multistep mechanism for hydrogen and air with nine species and 19 steps. Additionally, we implemented detailed chemical mechanisms for the combustion of hydrocarbons.

RESULTS & IMPACT

In this study, we simulated the RDE geometry experimentally studied by the Air Force Research Laboratory. The flow configurations, such as the fuel and oxidizer injection pressure, correspond to their experiment.




Figure 1: Structure of detonating flow inside an RDE. A detonation wave propagates in the azimuthal direction by consuming unreacted fuel–air mixture. (Purple: detonation front; green: oblique shockwave; blue: unreacted fuel gases; red and orange: detonation product species.)

This research confirms that stable detonations present within the RDE configuration are consistent with experimental observations. For example, the wave speed is slower than the premixed theoretical values due to the flow nonuniformity effect. With regard to the injection recovery process, the fuel injector recovers faster than the oxidizer, but both recover to only 75% of the initial profile within a quarter cycle of a detonation wave. Furthermore, the use of detailed chemical kinetics provides a complete view of the highly three-dimensional detonation structure.

Because meshes created by CAD (Computer-Aided Design) data can be directly simulated, our solver can dramatically accelerate the optimization process for engineering design. Further, the above result shows that the detailed RDE dynamics, and the nonuniformity effect on the detonation structure, can be extracted by simulating the complex injection system. This helps the RDE community better understand the physics and optimize this type of combustion system. In conjunction with experimental work, we generated high-fidelity data to assist in the realization of the next-generation combustor.

WHY BLUE WATERS

The Blue Waters high-performance computing resource greatly accelerated our research studies because we could use several thousand computing nodes and exploit the parallelizability of our solver, which is necessary to simulate large systems. Second, the working directory allowed us to manage and store the large amounts of data generated by numerical simulations without




Figure 2: Fluid stream trace near the injection system illustrating the complex flow patterns that induce turbulent mixing. Two streams of the oxidizer gases from the left and the fuel gases from the bottom are mixing in the detonation chamber.

concerns over exceeding a storage quota. Finally, multiple job queues catered to our requirements and deadlines. The computing resources allowed us to implement, test, and deploy computational tools quickly and achieve the results we desired. In addition, the project staff were readily available to assist in the implementation of our numerical tools.

PUBLICATIONS & DATA SETS

Sato, T., S. Voelkel, and V. Raman, Detailed Chemical Kinetics Based Simulation of Detonation-Containing Flows. *The ASME Turbo Expo* (Oslo, Norway, June 11–15, 2018).

COLLABORATIVE RESEARCH: ADVANCING FIRST-PRINCIPLE SYMMETRY-GUIDED NUCLEAR MODELING FOR STUDIES OF NUCLEOSYNTHESIS AND FUNDAMENTAL SYMMETRIES IN NATURE

Allocation: NSF PRAC/3,430 Knh

PIs: J.P. Draayer¹, W.M. Tang²

Co-PIs: K.D. Launey¹, T. Dytrych^{1,3}, B. Wang²

Collaborators: D. Langr⁴, M. Kocourek⁴, T. Oberhuber⁴, R. Baker¹, G. Sargsyan¹, A. Mercenne¹

¹Louisiana State University

²Princeton University

³Czech Academy of Sciences

⁴Czech Technical University

EXECUTIVE SUMMARY

The objective of this project is to provide nuclear wave functions of unprecedented accuracy by combining the cutting-edge computational power of the Blue Waters (BW) system with advanced techniques of nuclear structure calculations. This enables large-scale modeling of light- and intermediate-mass nuclei along with, for the first time, medium-mass nuclei. This is key to addressing the two most challenging questions in physics today; namely, the origin of elements and whether the neutrino is its own antiparticle. The work also supports and informs current and projected experimental efforts at state-of-the-art radioactive beam facilities, including the upcoming Facility for Rare Isotope Beams. Breakthrough theoretical advances [1,2], coupled with BW's cutting-edge computational power, have opened a new region—the intermediate-mass nuclei from fluorine to calcium isotopes—for first investigations with *ab initio* (“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 a dual challenge; namely, the nonperturbative nature of QCD in the low-energy regime and the complexity of many-particle nuclei. Because short-lived nuclei, currently inaccessible to experiment, are often found to be key to understanding processes in extreme environments ranging from stellar explosions to the interior of nuclear reactors, first-principle nuclear models that hold predictive capabilities will have 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” [3–5], 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 BW system, allow us to reach medium-mass nuclei that are inaccessible experimentally and to other *ab initio* methods (see Fig. 1) [2,6].

RESULTS & IMPACT

The nuclei of interest represent a considerable challenge requiring computational power of nearly the entire BW machine and its system memory. Two graduate students have carried forward these studies and had the unique opportunity to work




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.

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 intermediate-mass nuclei, such as ¹⁹Ne, ²⁰Ne, ²⁴Ne and ²⁴Si [2], with *ab initio* wave functions to be provided as input to our *ab initio* reaction code under development. Such nuclei, along with (α,p) and proton capture reactions, are key to further understanding the production of heavy elements, and especially the X-ray burst nucleosynthesis. Of special interest are short-lived nuclei (such as ²⁴Si) that are difficult, and often impossible, to study experimentally.
- We have studied emergent deformation and clustering in nuclei, from first principles, for Mg isotopes and their mirror nuclei (²¹Mg and ²¹F, ²²Mg and ²²Ne, with work in progress on ²³Mg and ²⁴Mg). While enhanced deformation and cluster substructures are difficult to describe from first principles, the BW system has allowed the first *ab initio* descriptions of deformed nuclei using chiral internucleon interactions [6]. This is important for providing accurate predictions for deformed and, in the future, heavy nuclei of interest to understanding the r process nucleosynthesis, one of the most challenging problems in astrophysics today.
- We have performed first-principle simulations of ⁴⁸Ca and ⁴⁸Ti 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 DUNE (Deep Underground Neutrino Experiment), which represents one of the most fundamental problems in physics today (see Fig. 2).

Large investments have been made in new generations of radioactive beam facilities to enable important discoveries in nuclear science. These calculations, enabled by the BW system, aim to make significant contributions in this area, supporting and informing current and upcoming experiments at these facilities. 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's PAID program, may have wider impact, as multiphysics simulations in the areas of nuclear energy and national security have similar needs.

WHY BLUE WATERS

Currently, only the BW system provides resources required for the *ab initio* studies of medium-mass isotopes with cutting-edge accuracy. To illustrate the level of complexity, applications to medium-mass nuclei require more than hundreds of exabytes of memory to store the Hamiltonian matrix. In order to capitalize on advances feasible with the SA-NCSM and BW capabilities, with the help of the BW staff, we managed to improve scalability and performance of our code. As a result, our largest production

runs utilized efficiently 715,712 concurrent threads running on 22,366 Cray XE6 compute nodes to solve the nuclear eigenvalue problem with a memory footprint of up to 750 TB of data. Clearly, the BW system represents a unique computational platform that already plays a crucial role in advancing *ab initio* nuclear theory toward new domains.

PUBLICATIONS & DATA SETS

Draayer, J.P., et al., Untangling simple patterns in intricate atomic nuclei. *Walter Greiner Memorial Volume*, World Scientific Publishing Co. (2018).

Langr, D., T. Dytrych, K.D. Launey, and J.P. Draayer, Accelerating Many-Nucleon Basis Generation for High Performance Computing Enabled *Ab Initio* Nuclear Structure Studies. *Intl. J. High Performance Computing Applications*, submitted (2018).

Pan, F., X. Ding, K.D. Launey, and J.P. Draayer, A simple procedure for construction of the orthonormal basis vectors of irreducible representations of O(5) in the OT(3) ⊗ ON(2) basis. *Nucl. Phys. A*, 974 (2018), p. 85.

Pan, F., et al., Exact solution of mean-field plus an extended T=1 nuclear pairing Hamiltonian in the seniority-zero symmetric subspace. *Physics Letters B*, 780 (2018), p. 1.

Burrows, M., et al., *Ab initio* translationally invariant nonlocal one-body densities from no-core shell-model theory. *Phys. Rev. C*, 97 (2018), p. 024325.

Launey, K.D., et al., *Ab initio* Picture of Nuclei: Shapes, Rotations, and Vibrations from Chiral Potentials. *Bulg. J. Phys.*, 44 (2017), p. 345.

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 (2017), p. 044312.

Draayer, J.P., T. Dytrych, and K.D. Launey, Symmetry-adapted no-core shell model—the *ab initio* picture of nuclear collectivity. *Emergent Phenomena in Atomic Nuclei from Large-scale Modeling: A Symmetry-guided Perspective*, ed. K.D. Launey, World Scientific Publishing Co. (2017); ISBN: 978-981-3146-04-4.

Draayer, J., et al., No-core Symplectic Model: Exploiting Hidden Symmetry in Atomic Nuclei. *Journal of Physics: Conference Series*, 1st ed., 863 (2017), p. 012008.

Dreyfuss, A.C., et al., Simple Patterns in Nuclei. *Proceedings of the 6th International Conference on Fission and properties of neutron-rich nuclei* (Sanibel Island, Fla., 2017).




Figure 2: First *ab initio* simulations of the medium-mass ⁴⁸Ti nucleus of importance to studies of the neutrinoless double-beta decay, with a view toward reducing large uncertainties in the nuclear structure matrix elements. This will, in turn, allow us to determine if the neutrino is its own antiparticle.

NOVEL METAL–CARBON MATERIALS

Allocation: Exploratory/50 Knh

PI: Iwona Jasiuk¹

Co-PI: Erman Guleryuz²

¹University of Illinois at Urbana-Champaign

²National Center for Supercomputing Applications

EXECUTIVE SUMMARY

We used the Blue Waters supercomputer to model novel metal–carbon materials called covectics, which are metals that are combined with carbon in a new way. Experimental results have shown that covectics have superior mechanical, thermal, and electrical properties compared to base metals. However, the role of the added carbon on materials performance is still not well understood. A goal of this project is to improve our understanding of covectics.

This project focused on copper covectics. The objectives were: (1) to build atomistic models using the Large-Scale Atomic Molecular Massively Parallel Simulator [2] (LAMMPS) molecular dynamics code to obtain a fundamental understanding of the material response of covectics and to provide inputs for continuum models; and (2) to develop experimentally-based continuum-level finite element and micromechanics models of covectics (using the Abaqus [4] software package) to create predictive tools for their bulk properties. Blue Waters is necessary because such multiscale modeling requires very high computational resources due to the hierarchical structure of these materials.

RESEARCH CHALLENGE

Experimental results have shown that covectics exhibit superior multifunctional properties compared to corresponding base metals, including improved mechanical properties, higher thermal and electrical conductivities, and higher corrosion resistance, making them excellent candidates for various technological applications. However, they need to be more fully studied and understood before they can be widely accepted and used commercially.

METHODS & CODES

Covectics are metals infused with carbon (up to 10 wt%) in amounts far exceeding the solubility limits predicted by phase diagrams. The metal is combined with carbon in a new way, which leads to a strong carbon–metal bond. To understand the nature of this bond, we performed powerful atomistic simulations. Atomistic simulations are computational techniques that can provide unique insights on materials characteristics and mechanisms that shape the material response. The most common methods in use are Monte Carlo algorithms, Molecular Dynamics (MD), and Density Functional Theory (DFT).

We conducted DFT calculations on covectics by assuming various copper–carbon configurations and studying lowest energy states.

Our simulations were guided by high-resolution transmission electron microscopy observations that an epitaxial rearrangement of a lattice structure forms in the covectics and that the carbon forms alternating layers of graphene between atomic planes of metal. We extended these studies to other geometries of graphene ribbons intercalated between metal layers. We are studying the most promising lowest-energy candidates determined in these simulations on larger systems.

We also employed MD modeling, which has not been used previously on covectics. MD allows us to study larger material volumes due to relatively reduced computational cost compared to DFT and is critical in studying experimentally observed structures. We used LAMMPS classical molecular dynamics code, which has a modular design; the community code model allows users to extend the code at the source level. LAMMPS uses message-passing and spatial domain decomposition techniques and is highly efficient on large-scale computational platforms. LAMMPS' extreme level of parallel scalability (Fig. 1) was previously demonstrated on Blue Waters [3].

Stresses and their corresponding strains within the elastic limit are related by elastic constants. In this work, we calculated the elastic constants after relaxing the geometry for systems with and without carbon. In addition, we are studying structural




Figure 1: The Lennard–Jones benchmark runs on Blue Waters XE6 nodes. The parallel speedup curve is plotted for a fixed-sized test case (strong scaling test) with 200 million atoms.




Figure 2: Orientation maps from white beam Laue diffraction for 0 wt% C (left) and 3 wt% C (right) copper–carbon covectics.

transformations occurring in the system (such as dislocations and void formation) for the most stable systems.

Our approach to creating continuum models involves computational modeling of a bulk mechanical response of copper covectics. To model the macroscopic mechanical response, we incorporated details from the micro scale (grain level) obtained from scanning electron microscopy, electron backscatter diffraction (SEM–EBSD), and X-ray microdiffraction imaging into the finite element program (Fig. 2). Data on texture, crystal orientations, and grain shapes obtained by SEM–EBSD and local mechanical properties measured by nanoindentation serve as inputs for such computational models.

We also modeled elastic and plastic materials responses by carrying out Finite Element Method (FEM) simulations with Abaqus, a commercial software package. Combined nonlinearities arising from inelastic material properties and geometric nonlinearities from large deformations leads to increased ill-conditioning of the global stiffness matrices. This makes the FEM analysis of 3D models difficult even using the latest high-performance computing platforms. To solve such a problem within each quasi-static timestep, a system of nonlinear equations must be linearized and solved in Abaqus with a Newton–Raphson iteration scheme, which requires several linear solver solutions or global equilibrium iterations. For this work, we used the direct multifrontal solver in Abaqus/Standard with hybrid Message-Passing Interface/Threaded parallelization.

RESULTS & IMPACT

This research is still in progress. However, our initial atomistic and continuum models led to the first predictive multiscale models for covectics [1].

WHY BLUE WATERS

The Blue Waters supercomputer and NCSA researchers were essential for this research project. Blue Waters is necessary because such multiscale modeling requires very high computational resources due to the hierarchical structure of covectics.

HIGH-ORDER METHODS FOR TURBULENT TRANSPORT IN ENGINEERING AND GEOSCIENCES

Allocation: Blue Waters Professor/300 Knh
PI: Paul Fischer¹
Collaborators: Som Dutta¹, Ketan Mittal¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Turbulent flows are prevalent in both natural and engineered environments; understanding them provides insight into the transport and mixing processes in a variety of systems. Here, we describe two ongoing turbulence simulation projects based on the scalable open-source spectral element code Nek5000. The first study explores the effect of biofilms on turbulent transport of fine-scale particles. Using measured biofilm morphology and new particle-tracking capabilities, these simulations have indicated mechanisms through which the structure of the biofilms influence the transport of fine particles and nutrients. The second project develops a scalable overlapping mesh method for high-fidelity simulations of turbulence in complex industrial and natural systems. In the development of this new capability, we have simulated turbulence in several configurations involving complex static and moving meshes.

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, the combustion in automotive and aerospace applications, sediment and pollutant transport in rivers and oceans, and atherogenesis in arterial blood flow. Our objective is to develop robust algorithms that can address research questions through direct numerical and large-eddy simulation of turbulent flows 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., at high speeds).

As mentioned, the first project studies the 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 those studies, however, had the spatial and temporal resolution to unravel the interaction between the biofilm structure and the flow, or the 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, we developed new methods for generating the computational mesh.

The second project is geared toward developing a scalable overset mesh method for large-scale high-fidelity turbulent simulations in complex geometries. We have developed the methodology to address problems that either have complicated meshes, involve disparate spatial scales, or have moving boundaries that would require the remeshing of standard conforming meshes.

METHODS & CODES

The turbulence simulations were based on the open-source spectral element code Nek5000 [2]. 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 [3]. The local tensor-product structure allows low-cost

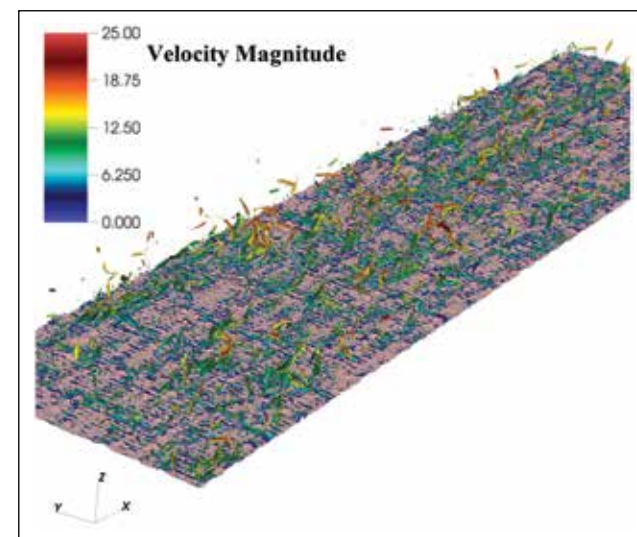


Figure 1: The image shows the coherent turbulent structures in the flow caused by the biofilm. Apart from the hairpin vortices, one may observe the vortical tubes formed behind the roughness elements, which are absent in the case of the smooth bed. The vortical structures have been colored using the magnitude of velocity.

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. 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. We conducted for the first time spectrally accurate DNS simulations for a channel with complex natural roughness. This was possible because of the development of sophisticated mesh-smoothing algorithms for Nek5000 [4]. For the second project, we used a recently developed method known as NekNek that has proven the validity of the overset grid method in context of SEM [5]. This new algorithm has been further developed and tested at scale on Blue Waters for large complex industrial problems [6].

RESULTS & IMPACT

We conducted simulations of the flow over biofilm with a bulk Reynolds number of 8,000 with mesh resolutions ranging from 20 million to 200 million grid points. The number of points depended on the details with which the biofilm morphology was captured. The results show the coherent turbulent structures in the flow caused by the biofilm (Fig. 1). Apart from the hairpin vortices, one may observe the vortical tubes formed behind the roughness elements, which are absent in the case of a smooth bed. This vortex formation results in a relative increase in the turbulence kinetic energy near the bed. Applicability of the current study goes beyond biofilms, as the methodology developed to simulate the flow over the biofilms could be utilized for any kind of natural surface roughness encountered in industrial and aerospace engineering applications.

For the second study, we simulated several turbulent flows to test the applicability of NekNek to complex industrial problems. The cases reported here include a jet flowing into a tank (Fig. 2), which is an idealized version of a setup encountered across different industries. We can observe that the domain has been decomposed into two subdomains, orange and black, with orange accounting for the incoming jet and the black for the rest of the domain. The results showed appreciable agreement among simulations with Nek5000 and NekNek, with around 20% reduction in computation cost. Scaling studies with the new NekNek code indicate that it retains the good strong-scale capabilities of the native Nek5000 [6].

WHY BLUE WATERS

The two studies we have reported here not only have large-scale computation requirements, but they also require fast turnaround times for the simulations, especially for the second study, in which we are testing a range of setups. This requires access to a dedicated petascale high-performance computing system, and Blue Waters fits that bill perfectly. We have conducted simulations using up to 296 million computational points, with the code demonstrating linear speed-up for this problem out to 32,768 MPI ranks. Since visualization of fluid flow is important to understanding its mechanics, we also work with Blue Waters project staff to create animations of the phenomena using data from the simulations.

PUBLICATIONS & DATA SETS

Mittal, K., and P.F. Fischer, Mesh smoothing for spectral element method. *J. Scientific Computing*, accepted for publication (2018).

Mittal, K., S. Dutta, and P.F. Fischer, Nonconforming Schwarz–Spectral element methods for incompressible flow. *Parallel CFD2018* (Indianapolis, Ind., May 14–17, 2018).




Figure 2: The figure shows velocity magnitude from a NekNek simulation of a turbulent jet entering a container. It also shows the two subdomains (orange and black) decomposed from the global domain. Adopting NekNek allows a reduction of the computational points required to simulate the problem.

PROPERTIES OF DENSE HYDROGEN

Allocation: Blue Waters Professor/240 Knh

PI: David Ceperley¹

Collaborators: Carlo Pierleoni², Marcus Holzmann³, Yubo Yang¹

¹University of Illinois Urbana-Champaign

²University of L'Aquila

³French National Center for Scientific Research

EXECUTIVE SUMMARY

The phase diagram of high-pressure hydrogen is of great interest for fundamental research, planetary physics, energy applications, and as a benchmark for computational quantum simulation methods. Using highly accurate coupled electron–ion Monte Carlo simulation methods, we predicted transitions between a molecular insulating fluid and a monoatomic metallic fluid. The existence and precise location of the transition line is relevant for planetary models. Experiments have reported contrasting results about the location of the transition; however, this year the most precise experiments have confirmed our predictions, validating the simulation methods for materials that cannot be measured in the laboratory. We have also explored the properties and structures of molecular and atomic solid hydrogen at lower temperatures.

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 for understanding the giant planets, Jupiter and Saturn, but experiments under the relevant conditions are challenging. Even though hydrogen is the first element in the periodic table, calculating its properties is not simple since both the electronic and protonic motions are quantum and correlated. Nonetheless, hydrogen is a valuable benchmark for computational quantum simulation methods.

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. Previously, we predicted a first-order phase transition in the fluid phase between a molecular insulating fluid and a monoatomic metallic fluid [2]. Recent experiments [3,4] reported contrasting results regarding the location of the transition that differed 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. Our recent experiments have also 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 among 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, as well as the difficulty in treating the hydrogen bond-breaking and the large van der Waals interactions. We model hydrogen with as many as 100 electrons and protons in a periodic cell and use special methods to extrapolate accurately to the thermodynamic limit. With our method, we simulated hydrogen for temperatures in the range of 200K up to 5000K, and at pressures of 100 GPa to 700 GPa.

RESULTS & IMPACT

For temperatures below 2000K we observed [5] a transition between an insulating molecular liquid and a more dense metallic atomic liquid. Our predicted transition pressures are intermediate between the two experimental observations [3,4]. Our work stimulated extensive experiments and further calculations on dense hydrogen. The most recent experiments confirm our predictions.

In work published in 2017 [7], we performed a detailed study of solid molecular hydrogen at temperatures between 200K and 414K and pressures between 180 GPa and 520 GPa to understand the detailed crystal structures, the melting temperature of the solid, and the changes in optical properties during the transitions. A controversial report, also published in 2017, concerns the experimental observation of this transition [6].

Our calculations are needed both to validate our computational method and to resolve the different experimental measurements. For progress in the high-pressure community to occur, it is essential to resolve this difference between experiments and computation. After validation, the method can be used with more confidence in modeling the wide variety of astrophysical objects being observed, which are 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 simulation of these systems.

WHY BLUE WATERS

Without access to Blue Waters we could not have performed this study because of the computational demands from the simultaneous treatment of quantum electrons and protons.

PUBLICATIONS & DATA SETS

Holzmann, M., et al., Theory of Finite Size Effects for Electronic Quantum Monte Carlo Calculations of Liquids and Solids. *Phys. Rev. B*, 94 (2016), DOI:10.1103/PhysRevB.94.035126.

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

Rillo, G., M.A. Morales, D.M. Ceperley, and C. Pierleoni, Coupled Electron–ion Monte Carlo simulations of hydrogen molecular crystals. *J. Chem. Phys.*, 148 (2018), p.102314.

Pierleoni, C., G. Rillo, D.M. Ceperley, and M. Holzmann, Electron localization properties in high pressure hydrogen at the liquid–liquid phase transition by Coupled Electron–Ion Monte Carlo. *Proceedings of CCP2017*, arXiv:1712.00392 (2017).

Pierleoni, C., M. Holzmann, and D.M. Ceperley, Local structure in dense hydrogen at the liquid–liquid phase transition by coupled electron–ion Monte Carlo. *Contributions to Plasma Physics*, 58 (2018), pp. 99–106.

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

Allocation: NSF PRAC/1,660 Knh

PI: Gerhard Klimeck¹

Co-PI: Tillmann Kubis¹

Collaborators: James Charles¹, Fan Chen¹, Jim Fonseca¹, Junzhe Geng¹, Xinchun Guo¹, Hesameddin Ilatikhameneh¹, Daniel Mejia¹, Bozidar Novakovic¹, Michael Povolotskyi¹, Rajib Rahman¹, Prasad Sarangapani¹, Archana Tankasala¹, Yu Wang¹

¹Purdue University

EXECUTIVE SUMMARY

With a revenue of \$338.9 billion in 2016, and as an enabler for larger economy chains such as 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 its dependent product chains and economies. Needed improvements in the performance of transistors 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, multiphysics phenomena for nanoscale technology through efficient computational approaches and enables quantitative study of new generations of nanoelectronic devices even beyond transistors [2,3].

RESEARCH CHALLENGE

The United States has always been a world leader in the semiconductor industry, with 40% of the worldwide semiconductor device-related patents originating here [4]. The industry is one of the nation's largest and most strategic, and the U.S. holds one-third of the global semiconductor device market, which is 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 devices reduce power consumption?




Figure 1: Cylindrical gallium arsenide nanowire transistor represented in a mixed atomistic and solid mechanics basis. The number of atoms in such a device is countable, and the local atom arrangements are critical. The source and drain of the nanowire transistor are conceptualized as thin. The wrap-around gate is indicated by the green solid.

NEMO5 was developed and used by the Institute for NanoElectronic Modeling (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.




Figure 2: Atomistic representation of a silicon-germanium (SiGe) ultra-thin-body transistor. SiGe is a random alloy with atomistic local disorder. The source and drain are indicated by the rounded shapes. The top gate regulates the electron flow from source to drain.

METHODS & CODES

iNEMO's research on Blue Waters encompasses multiphysics atomistic quantum simulations implemented in NEMO5. The needed physics vary from one device to another, but all can be simulated inside NEMO5. For example, for transistors, quantum transport simulations are performed using the self-consistent Poisson with a nonequilibrium Green's function (NEGF) approach employing semiempirical tight-binding methods. The current mechanism of tunneling field effect transistors (TFETs) is through interband tunneling rather than thermionic emission in typical MOSFETs (metal-oxide semiconductor field-effect transistors). For realistic performance predictions, incoherent scattering effects are also included.

Several classes of devices demand a multiscale decomposition into regions of high- and low-electron densities. These different charge densities imply very different scattering mechanisms and thermalization processes. For example, for nitride-based light emitting diodes and alloy-engineered nitride TFETs, we use a multiscale quantum transport model, which treats high-density regions as local charge reservoirs where each reservoir serves as a carrier injector/receptor to the next/previous reservoir.

Atomistic basis sets are mapped from fundamental *ab initio* methods such as Density Functional Theory into reduced-order tight-binding or localized basis sets. With such reduced-order models, we can represent realistically extended devices.

RESULTS & IMPACT

Interconnect modeling: Orientation effects on the specific resistance of copper grain boundaries are studied systematically with two different atomistic tight-binding methods. A methodology is developed to model the specific resistance of grain boundaries at the ballistic limit using the embedded atom model, tight-binding methods, and nonequilibrium Green's functions. The methodology is validated against first-principles calculations for thin films with a single coincident grain boundary resulting in a 6.4% deviation in the specific resistance compared to DFT. A statistical ensemble of 600 large, random structures with grains is studied. For structures with three grains, we found that the distribution of specific resistances is close to normal. Finally, we constructed a compact model for grain-boundary-specific resistance based on a neural network.

Quantum device modeling: The discovery of integer (1980) and fractional (1982) quantum Hall effects revolutionized the field of condensed-matter physics by providing a platform for testing the effects of strong interactions among electrons, which are difficult to observe elsewhere. Strong interelectron interactions create a liquid-like state in which the current is carried by fractionally charged quasiparticles known as composite-Fermions (CFs). Further interactions among CFs could lead to the formation of quasiparticles that follow non-Abelian statistics, meaning they would hold a memory of the way in which their world lines are braided. These yet-unobserved non-Abelian quasiparticles could provide a platform for fault-tolerant topological quantum computing applications. In this work, we developed a computational model for devices that could be used to observe these quasiparticles and to propose better designs.

These devices consist of gallium arsenide-based nanostructures that are used for confining electrons in two dimensions. We applied a perpendicular magnetic field to limit electron/CF flow to the edges of the sample in one-dimensional channels. Such channels can be brought together from across the device and made to interfere with each other, like photons using gate-defined one-dimensional constrictions called quantum-point contacts. Observation of these interference oscillations can shed light on the properties of the non-Abelian quasiparticles. These quasiparticles must be traveling fast enough to maintain their phase while interfering and show a strong interference pattern.

We propose new designs for quantum-point contacts with high electron velocity along the edge channels using our modeling tool. Blue Waters has provided us with a unique and reliable computing environment to support these kinds of simulations.

WHY BLUE WATERS

Quantum transport simulations are very computationally expensive and demanding of memory 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

Continued on page 147

MODELING MILLIONS OF ATOMS AND THEIR ELECTRONS: TIGHT-BINDING USING A SPECTRAL PROJECTION PURIFICATION METHOD

Allocation: Exploratory/50 Knh
PI: Harley Johnson¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

In this project, we demonstrated a new method for accurately computing properties of material systems containing up to 10 million atoms and their electrons. The method is based on a linear scaling version of tight binding, which is a semiempirical approach for computing, from a model, the energy of a collection of atoms for the interactions among the electrons and the nuclei. We computed the density matrix—a fundamental property of such a system—using a spectral projection purification method. The method scaled well on Blue Waters to millions of atoms over thousands of processors. As an example, we computed the electronic properties of an SiO₂ crystal, which we reported in the form of the density of states. The results show the feasibility of the approach as a general-purpose total-energy electronic structure computational tool for materials science.

RESEARCH CHALLENGE

Our interest is in testing more efficient methods for computing the electronic structure of systems with millions and tens of millions of atoms. Properties of material domains of this size are important in many different technologies, from microelectronics to medical science to energy technologies. Our interest is in the properties of dielectric materials of plasma generation, which are important in combustion, materials processing, and catalysis. Traditional electronic structure calculations for these kinds of materials have cubic scaling in N —the number of atoms in the system—making the calculations prohibitively computationally expensive for systems larger than 10⁴ atoms.

METHODS & CODES

We used the code we developed to compute the density matrix, a fundamental electronic structure quantity for a system of atoms and their electrons. Our computational method [1] combines the advantages of two existing linear scaling methods: the Kernel Polynomial Expansion (KPE) [2] and second-order spectral projection purification (SP2) [3]. The KPE method is computationally efficient and can be easily expressed in terms of sparse matrix–vector multiplications (SpMVs), but cannot satisfy one of the required constraints on its own. On the other hand, the SP2 method is highly accurate but can be prohibitively costly in terms of memory and communication when expressed in terms of sparse matrix–matrix multiplications (SpGEMM). When

expressed in terms of SpMVs, the SP2 method scales exponentially with the number of iterations required to converge. An advantage of SP2, however, is that near the correct solution P , the method converges quadratically [3]. Thus, we have constructed a hybrid method that takes the inexpensive KPE solution and purifies P_{kpe} with a few SP2 iterations that are more expensive.

RESULTS & IMPACT

Our recent work has shown the feasibility of this method for computing properties of million-atom systems [1]. We will apply the method and have impact in the area of surface and dielectric properties for combustion, catalysis, and materials processing. The method can also have impact in other areas of physics and materials science. With the scaling observed during the outcome of the Blue Waters work, we have demonstrated that this method is fast, accurate, and efficient in 10-million-atom systems and beyond, which will increase its impact significantly.

WHY BLUE WATERS

Access to the Blue Waters system made this work possible by permitting our studies on a single platform that offered both large parallelism and resources for memory-intensive computations. Depending upon the tuning of the method, it is possible to carry out the entire computation on a single node in a very memory-intensive approach or, by distributing the computation over many processors, it is possible to perform the algorithm in a massively parallel way. The Blue Waters system allowed us to study the approach and plan for future studies with more optimized tuning.

Continued from page 145

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 a 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 (20 × 20) nm², which includes the core semiconductor and the surrounding gate material. Such devices of larger sizes are especially an issue due to the $O(n^2)$ 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 assist with code performance and scaling.

PUBLICATIONS & DATA SETS

Sahasrabudhe, H., et al., Optimization of edge state velocity in the integer quantum Hall regime. *Physical Review B*, 97 (2018), DOI:10.1103/PhysRevB.97.085302.

Valencia, D., et al., Grain-Boundary Resistance in Copper Interconnects: From an Atomistic Model to a Neural Network. *Physical Review Applied*, 9:4 (2018), p. 044005.

Chen, C.-Y., et al., Channel thickness optimization for ultra thin and 2D chemically doped TFETs. Submitted (2018), arXiv:1804.11034.

Ilatikhameneh, H., et al., Sensitivity Challenge of Steep Transistors. *IEEE Transactions on Electron Devices*, 65:4 (2018); DOI:10.1109/TED.2018.2808040.

Tankasala, A., et al., Two-electron states of a group V donor in silicon from atomistic full configuration interaction. *Phys. Rev. B*, 97 (2018), DOI:10.1103/PhysRevB.97.195301.

Long, P., et al., Atomistic Modeling trap-assisted tunneling in hole tunnel FETs. *Journal of Applied Physics*, 123 (2018), DOI:10.1063/1.5018737.

Gen, J., et al., Quantitative Multi-Scale, Multi-Physics Quantum Transport Modeling of GaN-Based Light Emitting Diodes. *Phys. Status Solidi A*, 215:9 (2018), DOI:10.1002/pssa.201700662.

Ameen, T., et al., Combination of Equilibrium and Nonequilibrium Carrier Statistics Into an Atomistic Quantum Transport Model for Tunneling Heterojunctions. *IEEE Transactions on Electron Devices*, 64:6 (2017), pp. 2512–2518.

SIMULATING THE EMERGENT PHENOMENA ARISING FROM STRONGLY CORRELATED SYSTEMS

Allocation: Blue Waters Professor/250 Knh

PI: Bryan Clark¹

Collaborators: Xiongjie Yu¹, Di Luo¹, Eli Chertkov¹, Benjamin Villalonga¹, Dmitrii Kochkov¹, Jahan Claese¹, Krishna Kumar¹, Eduardo Fradkin¹, Hitesh Changlani², Vadim Oganesyan³, Gil Rafeal⁴, David Luitz⁵

¹University of Illinois at Urbana-Champaign

²Johns Hopkins University

³City University of New York

⁴California Institute of Technology

⁵Max Planck Institute for the Physics of Complex Systems

EXECUTIVE SUMMARY

This work focuses on using simulations and developing new algorithms to understand exotic quantum phenomena. For example, a key goal in condensed matter physics is to find Hamiltonian models with interesting physics; we developed a novel inverse algorithm to accomplish this [1]. Instead of the typical forward method of guessing at Hamiltonian models that might produce interesting physics, we start with the interesting physics encoded in a wave-function and automatically find Hamiltonians that support them. Additionally, we have discovered a new special Hamiltonian that sources almost all the phases on frustrated magnetism materials, explaining the menagerie of observed phases there [2] and connecting to the enigmatic spin-liquid [3] on the kagome.

RESEARCH CHALLENGE

The rules of quantum mechanics are simple. The phenomena that arise from these rules is both difficult to simulate (with an exponential cost in system size) and results in complicated phenomena. Both these properties are the result of quantum entanglement.

We address three key research problems in our work:

(1) A key goal in condensed matter is to find models (i.e., Hamiltonians) with interesting physics. A typical approach toward accomplishing this is to try many different Hamiltonians until a good one is found; this is inefficient, especially since each Hamiltonian is exponentially costly to simulate. We developed a novel approach to tackle this key problem.

(2) One universal property of complicated quantum systems is the panoply of competing phases that comes from similar Hamiltonians. Is this a cosmic accident or is there a unifying explanation for this menagerie? We focus on resolving this question via a class of materials called frustrated magnets [4,5]; these are insulating materials whose spin degrees of freedom reside on lattices, such as the triangular or kagome lattice.

(3) Recently, it has been realized that a class of physical systems, coined the many-body localized phase [6,7], have exotic entanglement at high temperature causing statistical mechanics

to break down. We numerically probe the transition between this many-body localized phase and the ergodic phase where statistical mechanics still operate.

METHODS & CODES

(1) Our approach to this problem was to develop an inverse technique: start with some targeted interesting physics encoded in a wave-function and automatically determine a physically reasonable Hamiltonian from which those physics might arise. Our new algorithm, the Eigenstate-to-Hamiltonian Construction (EHC), uses a quantum covariance matrix that is a quantum generalization of the typical covariance matrix from statistics.

(2) We considered the XXZ Hamiltonian on the frustrated kagome lattice and computed its phase diagram. This was done using a highly parallel exact diagonalization code that we developed.

(3) We used a novel algorithm our group recently developed, SIMPS [8], to compute interior eigenstates of a $2^{32} \times 2^{32}$ Hamiltonian matrix in the many-body localized phase and measured their properties at a temperature slightly beneath the ergodic phase.

RESULTS & IMPACT

(1) Our algorithm, EHC, opens up an entirely novel approach to solving condensed matter problems and changes the search for interesting physics from a search in the dark to a targeted one. While the typical forward approach is exponentially slow, EHC is a quadratic algorithm. Although 50 sites is intractable [9] in the forward method—even on Blue Waters—thousands of sites can be simulated using the inverse approach.

(2) While most Hamiltonians have unique ground states, we have discovered a new Hamiltonian, the XXZ0 point [10], which has an exponential number of ground states that can be represented as all possible ways of coloring the kagome lattice with three colors. This Hamiltonian then becomes the source for the menagerie of competing phases because they all have identical energy in the ground state. We showed in a concrete example five explicit phases surrounding the XXZ0 point, including the enigmatic spin-liquid.




Figure 1: Our algorithm, Eigenstate-to-Hamiltonian Construction (EHC), generates Hamiltonians as a linear combination of null vectors of the quantum covariance matrix. Shown here is a color plot of the quantum covariance matrix that EHC computed for a wave-function of singlet dimers.

(3) We discovered that many-body localized eigenstates at low temperature can tell that there is an ergodic phase above them at higher temperatures. This allows us to learn about the transition using the many-body localized eigenstates that are easier to probe.

WHY BLUE WATERS

Without Blue Waters, we would not have been able to perform these calculations. It was an essential component to benchmarking EHC on a series of examples and will become even more important




Figure 2: The XXZ0 Hamiltonian has an exponentially degenerate ground-state space. Each of the exponential states in this ground state can be represented as a product state over three spins (here represented as red, blue, and green). Shown is one prototypical three-coloring ground state on the kagome lattice.

as we begin to enumerate a long list of interesting wave-functions. Our discovery of the macroscopically degenerate point relied on having enough computational power to explore a region of phase space that no one had ever considered before. Then, once it was discovered, it required thousands of phase points to validate that it was connected to the spin-liquid phase. Without computation at the scale of Blue Waters, this would not have been possible. Finally, the examination of the many-body localized phase involves studying samples with disorder. To extract any physics requires averaging many thousands of disordered samples that we run in parallel on Blue Waters.

PUBLICATIONS & DATA SETS

Claes, J., and B. Clark, Finite-temperature properties of strongly correlated systems via variational Monte Carlo. *Physical Review B*, 95:20 (2017), DOI:10.1103/PhysRevB.95.205109.

Pekker, D., et al., Fixed points of Wegner-Wilson flows and many-body localization. *Physical Review Letters*, 119:7 (2017), DOI:10.1103/PhysRevLett.119.075701.

Changlani, H.J., et al., Macroscopically Degenerate Exactly Solvable Point in the Spin-1/2 Kagome Quantum Antiferromagnet. *Physical Review Letters*, 120:11 (2018), DOI:10.1103/PhysRevLett.120.117202.

Villalonga, B., et al., Exploring one-particle orbitals in large many-body localized systems. *Physical Review B*, 97:10 (2018), DOI:10.1103/PhysRevB.97.104406.

Chertkov, E., and B.K. Clark, A computational inverse method for constructing spaces of quantum models from wave functions. *arXiv:1802.01590* (2018).

Yu, X., D. Luo, and B.K. Clark, Beyond many-body localized states in a spin-disordered Hubbard model with pseudo-spin symmetry. *arXiv:1803.02838* (2018).

Sahoo, J., et al., Classical phase diagram of the stuffed honeycomb lattice. *arXiv:1805.09831* (2018).

PRESSURE FLUCTUATIONS INDUCED BY SUPERSONIC TURBULENT BOUNDARY LAYERS

Allocation: NSF PRAC/6,900 Knh

PI: Lian Duan¹

Collaborator: Meelan Choudhari²

¹Missouri University of Science and Technology

²NASA Langley Research Center

EXECUTIVE SUMMARY

The fluctuating pressure on aerodynamic surfaces of flight vehicles plays an important role in vibrational loading during atmospheric reentry; the freestream pressure fluctuations radiated from a supersonic boundary layer is related to the genesis of freestream acoustic noise by walls of a supersonic wind tunnel. This research exploited the cutting-edge computational power of Blue Waters to advance fundamental understanding of the generic statistical and spectral features of boundary-layer-induced pressure fluctuations at the wall and in the freestream. Such an understanding paves the way for developing physics-based models to predict vibratory loading of reentry vehicles, as well as to help define the disturbances caused by the tunnel walls in supersonic/hypersonic wind tunnels and allow more accurate extrapolations of wind-tunnel measurements to free flight.

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 is 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 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.

METHODS & CODES

We conducted Direct Numerical Simulations (DNS) 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. The very-low viscosity fluxes of the governing equations are computed using a seventh-order weighted essentially nonoscillatory (WENO) scheme. Compared to the original finite-difference WENO introduced by Jiang and Shu [1], the present scheme is optimized by means of limiters [2] to reduce the numerical dissipation; WENO adaptation is limited to the boundary-layer region for maintaining numerical stability while the optimal stencil of WENO is used outside the boundary layer for optimal resolution of the radiated acoustic field. A fourth-order central difference scheme is used for the viscous flux terms, and a third-order low-storage Runge–Kutta scheme [3] is employed for time integration, which significantly lowers the memory requirement and is well-suited for time-accurate simulations like DNS. The turbulent inflow can be generated using either a recycling/rescaling method [4] or a digital filtering method [5]. 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 (wing tip to wing tip) direction.

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 transient as a result of inflow boundary conditions. A large number of timesteps are also required for the study of low-frequency behavior of the pressure spectrum. As such, the proposed computational efforts cannot be completed without the world-class computing capabilities of Blue Waters.

PUBLICATIONS & DATA SETS

Duan, L., et al., Characterization of Freestream Disturbances in Conventional Hypersonic Wind Tunnels. *2018 AIAA Aerospace Sciences Meeting, AIAA SciTech Forum*, AIAA Paper 2018–0347, DOI:0.2514/6.2018-0347.

Huang, J., et al., High-Mach-Number Turbulence Modeling using Machine Learning and Direct Numerical Simulation Database, AIAA Paper 2017–0315, *55th AIAA Aerospace Sciences Meeting*, DOI:10.2514/6.2017-0315.

RESULTS & IMPACT

The current work advanced the state-of-the-art knowledge of the global pressure field induced by supersonic turbulent boundary layers across a wide range of Mach numbers. The trailblazing study represents the first-ever attempt to exploit the advances in high-performance computing to overcome the difficulties in experimental measurements, and to provide access to both flow and acoustic quantities that are difficult to otherwise obtain. In particular, the study led to the first successful comparison between numerical predictions and wind tunnel measurements of surface pressure fluctuations underneath hypersonic turbulent boundary layers at Mach 6, 8, and 14 (Fig. 1). The simulations also captured all major features of the freestream disturbance spectra that had been measured during a one-of-a-kind, historical wind tunnel campaign [6] (Fig. 2), helping to clarify the physics of the noise generation process in supersonic/hypersonic wind tunnels. The characterization of wind-tunnel freestream disturbances paved the way for extrapolation to flight from the boundary-layer transition data obtained in noisy wind tunnels.

WHY BLUE WATERS

Direct numerical simulations are used to capture both the broadband turbulence field within the boundary layer and the near-field acoustic disturbances radiated by the boundary layer. In such simulations, extremely fine meshes are required to fully resolve all the turbulence scales in order to obtain the pressure spectra in




Figure 1: Comparison between numerical predictions and wind tunnel measurements of surface pressure fluctuations underneath a hypersonic turbulent boundary layer, with the experimental data measured on the nozzle walls of the Purdue Boeing/AFOSR Mach-6 Quiet Tunnel, the Sandia Hypersonic Wind Tunnel at Mach 8, and AEDC Tunnel 9.




Figure 2: Comparison between numerical predictions and wind tunnel measurements of bulk propagation speed of the freestream acoustic pressure fluctuations radiated from a supersonic turbulent boundary layer, with the experimental data measured during a one-of-a-kind, historical wind tunnel campaign by Laufer [6].

TOPOLOGIES OF ELECTRONIC STATES OF NANOCRYSTALLINE MATERIALS

Allocation: Illinois/50 Knh

PI: Prashant K. Jain¹

Co-PIs: Sudhakar Pamidighantam², Daniel Dumett Torres¹

¹University of Illinois at Urbana-Champaign

²Indiana University Bloomington

EXECUTIVE SUMMARY

The surface electronic states of crystalline solids determine a number of their physical properties. An exemplary class of materials for which surface states are intrinsically tied to a defining characteristic is the topological insulator (TI) family. The resistanceless surface conduction of TIs, which has potential applications in spintronic devices and dissipation-less transistors, originates from spin-protected surface electronic states. Topologically protected electron conduction in bulk TI materials has been the focus of many recent experimental and theoretical studies. However, the potentially diverse characteristics and topologies that these states may exhibit on the nanoscale have not been fully explored. Critical questions remain such as whether topologically protected surface states can exist on the surfaces of small nanocrystals. Blue Waters permits us to answer such questions through computations of realistic models of nanostructures consisting of thousands of atoms. Our computations will drive discoveries of novel semiconductor phases and electronic topologies in nanostructures.

RESEARCH CHALLENGE

Engineered nanocrystals have promise in new functional electronic and optical materials. Our laboratory makes use of unconventional techniques of nanomaterials synthesis that often produce novel compositions and crystal phases that are not found in the bulk phase diagram. Computational electronic structure investigations are allowing us to explore the properties of these unconventional nanomaterials and to discover novel physicochemical phenomena and principles.

As one key example, it is thought that novel electronic topologies may arise on the surfaces of nanocrystalline structures. The surface electronic states can determine key physical properties of a crystalline solid. A class of materials for which surface states are intrinsically tied to a defining characteristic is the topological insulator (TI) family. The resistanceless surface conduction of TIs results from the presence of spin-protected helical electronic states at the surface. However, questions remain as to whether topologically protected surface states can exist on the surfaces of small nanocrystals. If not, are there novel topologies associated with surface electronic states of nanocrystals? Blue Waters is a crucial resource that is permitting us to answer these questions. Our investigations may identify nanostructures with unconventional

electronic transport properties of utility in energy storage and information technology.

METHODS & CODES

We use the open source Quantum Espresso software suite [1] to run electronic structure calculations using the plane-wave density functional theory (DFT) method. This software was deployed on Blue Waters as part of the work conducted with our prior exploratory allocation. We make use of Quantum Espresso's numerous parallelization levels that distribute large matrices and linear algebra operations among groups of nodes with good load-balancing. These parallelization schemes involve sizable and frequent communication among nodes, which relies on the speed of the Blue Waters communications hardware.

As our model system, we chose the $\text{Hg}_x\text{Cd}_{1-x}\text{Se}$ alloy, where our calculations have predicted the presence of an inverted band structure and a nontrivial band gap (Fig. 1), two key attributes of a TI. In this model system, we are exploring whether there are critical differences in the nature of surface states of $\text{Hg}_x\text{Cd}_{1-x}\text{Se}$ bulk slabs versus those of nanocrystals of $\text{Hg}_x\text{Cd}_{1-x}\text{Se}$. What are the resulting implications for the electronic properties of the material? How do the surface states of $\text{Hg}_x\text{Cd}_{1-x}\text{Se}$ nanocrystals change as the nanocrystal size is varied from 2 nm to 10 nm? We are interested in regimes where the nanocrystal electronic structure deviates from that of the bulk due to small-size effects.

Using DFT, we calculate electronic surface states of $\text{Hg}_x\text{Cd}_{1-x}\text{Se}$ bulk slabs. Band-structure calculations coupled with projection of the electron spin permit the protected states of slabs to be identified. Study of size effects requires a different approach. Since nanocrystals below a certain size do not have bands, band-structure calculations cannot be used. We instead use the projected density of states and charge density maps obtained from calculations of whole nanocrystals. The projected density of states will reveal whether size regimes exist wherein quantum confinement enhances the band gap while retaining the protected surface states.

RESULTS & IMPACT

From our band-structure calculations, we are interested in uncovering how nanocrystal electronic structure and topology deviates from that of the bulk due to small-size effects. These deviations may amount to a complete disruption of the TI surface states. Alternatively, it might be found that nanocrystals of the




Figure 1: Our calculated band structures [3] show that in the novel wurtzite form of HgSe, symmetry breaking through elongation of Hg–Se bonds results in the opening of a band gap (right), whereas the conventional form of HgSe has no band gap (left). Hg atoms are shown in purple and the electron density of the highest energy valence band is shown in red (center).

appropriate size can retain the TI properties. From the electronic structure of nanocrystals that do or do not retain TI states, we will elucidate the underlying physics. This will allow us and others to tune electronic topologies by nanostructure design.

The proposed studies will have marked impact on both the scientific community and society at large. Computational results demonstrating TI behavior in nanocrystals will put forth a new class of materials for study of fundamental physics and device applications [2]. Moreover, an understanding of the surface state characteristics will provide guiding principles to expedite new discoveries about unique electronic behavior on the nanoscale.

TIs are proposed to lead to energy-efficient logic devices because they provide a means to conduct electrons with almost zero resistance. This property could be used to reduce the amount of wasteful heat generated by logic devices, lowering the energetic cost of cooling and therefore enabling further scaling up of high-performance computing resources beyond current limitations imposed by overheating.

WHY BLUE WATERS

Our project requires use of the Blue Waters resource due to the computational scale of the bulk slab and nanocrystal calculations. Tens of nodes are required to fulfill the memory requirements of these calculations, exceeding the ability of small departmental clusters. Moreover, our investigation of the alloying and size effects requires not one but numerous such calculations. The computational expense of our effort would be prohibitive were it not for a Blue Waters allocation. Blue Waters will be able to accommodate any necessary increase in the size of our models. Furthermore, the specialized hardware of Blue Waters allows the Quantum Espresso code to run even more efficiently. Quantum Espresso's parallelization schemes involve sizable and frequent communication among nodes, which rely on the speed of the Blue Waters communication hardware.

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 one such possible material, and their hydrogen storage potential and microscopic properties need further investigation. Toward that end, we performed first-principles quantum-mechanical molecular-dynamics calculations to understand the behavior of hydrogen inside MOFs. We included the effect of quantum nuclear motions, which are critical for the properties of hydrogen, the lightest element. These challenging calculations are only possible on large-scale and massively parallel 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 used 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 source of green energy, but, as a fuel it requires efficient storage materials that retain and release a great deal of hydrogen. We studied the properties of hydrogen inside MOFs to understand their physical properties and potentially how to improve MOFs to deliver increased hydrogen storage. We simulated hydrogen inside of MOFs at the atomistic scale

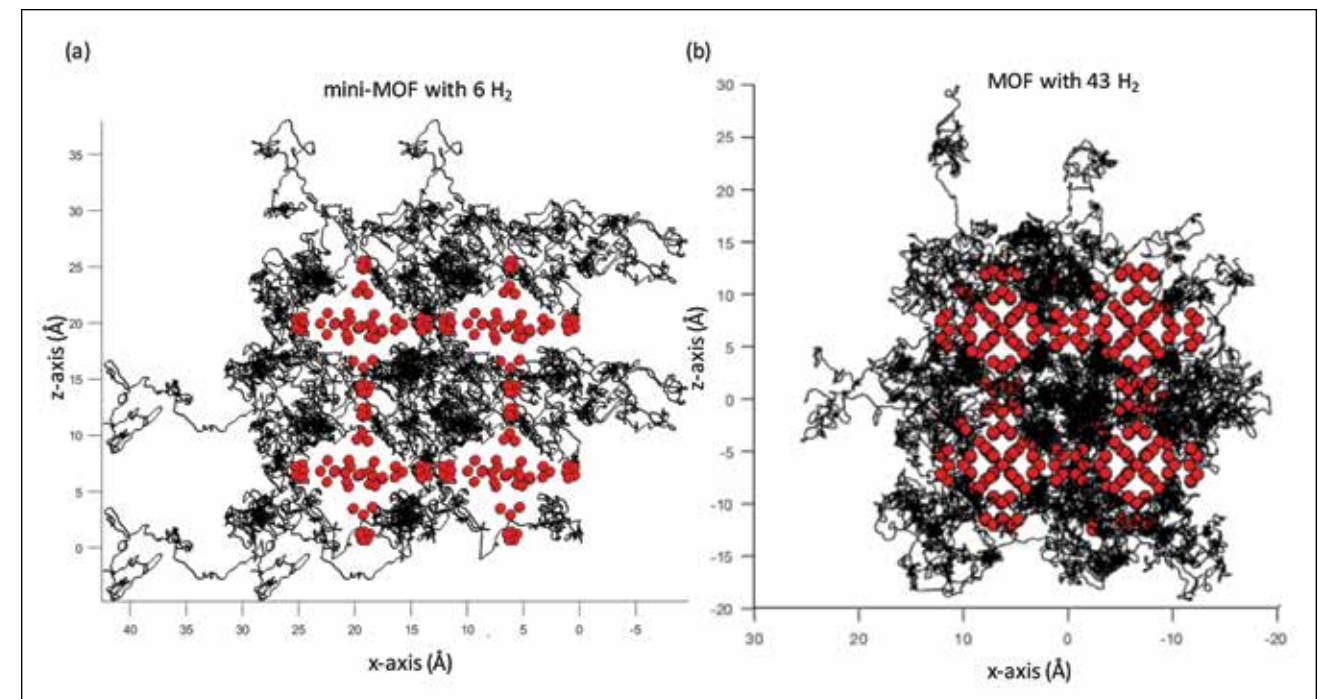


Figure 2: A 2D projection of the MOF backbone atoms (in red) along with the trajectories followed by the center of mass of each H₂ molecule diffusing in (a) the mini-MOF, and (b) the full MOF at 77K. Periodic boundary conditions were imposed to extend the trajectories to neighboring unit cells.

using accurate first-principles quantum-mechanical simulations based on Density Function Theory (DFT). We aim to understand the binding and dynamics of light elements such as hydrogen inside MOFs.

METHODS & CODES

Answering the above questions requires an accurate quantum-mechanical simulation at finite temperature. The method must work beyond the harmonic approximation because 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 employed the Car-Parrinello Molecular Dynamics simulation technique (CPAIDM), which allows the nuclei to move on the Born-Oppenheimer energy surface provided by plane wave-based DFT. By including the electronic degrees of freedom (valence) explicitly, we bypassed force field difficulties.

RESULTS & IMPACT

We spent the majority of the 2017 calendar year performing molecular dynamics simulations of a MOF (the zinc-containing "MOF-5" material) using our OpenAtom first-principles molecular dynamics software. We first obtained simulation on a smaller 1/8-sized subsection of the full MOF (denoted as mini-MOF) and then compared the results with the full MOF. The full MOF simulation has in excess of 510 atoms per simulation cell with

a large amount of "vacuum" (voids) in the structure to provide room for hydrogen diffusion, and a large internal surface area for hydrogen adsorption. The smaller size of the mini-MOF allows us to run very long molecular dynamics simulations and benchmark the results carefully. We studied the dynamics of 43 hydrogen molecules inside this metal-organic molecular system for the full MOF; we studied the diffusion of six H₂ molecules for the mini-MOF. We performed our simulations at 77K and 300K for both systems to understand the effect of temperature on hydrogen diffusion. Fig. 1 shows a comparison at 77K of the mini-MOF and full MOF mean square displacement versus time delay; the slope of a linear fit gives us the diffusion coefficient. Fig. 2 shows a 2D projection of the MOF backbone atoms (in red) along with the trajectories followed by the center of mass for each H₂ molecule diffusing in (a) the mini-MOF, and (b) the full MOF at 77K. Both simulations are long enough to show significant diffusion of the H₂ molecules and thus are statistically meaningful.

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 CPAIDM simulations already require a massively parallel calculation with many hundreds of nodes.




Figure 1: Mean square displacement at 77K of the centers of mass of the H₂ molecules as a function of elapsed simulation time. For both the mini-MOF and full MOF, each colored curve represents a unique H₂ molecule in the simulation cell. The black curve is the average over all the H₂. The slope of the average curve is a direct measure of the diffusion coefficient for H₂. The diffusion coefficient (D) for the mini-MOF is approximately $1 \times 10^{-8} \text{ m}^2/\text{s}$ and for the full MOF it is $8 \times 10^{-9} \text{ m}^2/\text{s}$. Both values are very close to prior simulation results of $7 \times 10^{-9} \text{ m}^2/\text{s}$ where the simulation parameters were fitted to experimental databases [2].

NUMERICAL SIMULATIONS OF THE INERTIAL COLLAPSE OF INDIVIDUAL GAS BUBBLES NEAR A RIGID SURFACE

Allocation: GLCPC/880 Knh

PI: Eric Johnsen¹

Co-PIs: Shahaboddin Alahyari Beig¹, Mauro Rodriguez¹

¹University of Michigan

EXECUTIVE SUMMARY

Cavitation, or the formation of vapor bubbles in a liquid flow, occurs in applications ranging from naval hydrodynamics to medicine to energy sciences. Vapor cavities can collapse violently in an inertial fashion that concentrates energy into a small volume, producing high pressures and temperatures, generating strong shock waves, and even emitting visible light. One of the main outcomes of inertial cavitation is structural damage to neighboring solid objects due to bubble collapse. While cavitation erosion is a negative outcome in naval hydrodynamics, it is exploited in medicine in the context of cavitation-based therapeutic ultrasound; e.g., for pathogenic tissue ablation.

To better understand this phenomenon, we used highly resolved numerical simulations of the inertial collapse of individual bubbles near a rigid surface to predict pressures and temperatures thereby produced. For this purpose, we developed a novel numerical multiphase model combined with high-performance computing techniques to perform large-scale, accurate, and efficient simulations of the three-dimensional compressible Navier–Stokes equations for a binary, gas–liquid system. This knowledge will paint a clearer picture of the detailed physics of such complex flows, it will elucidate the damage mechanisms, and it will inform the development of mitigation strategies for cavitation erosion.

RESEARCH CHALLENGE

Cavitation—the process whereby vapor cavities are produced in a liquid—is a ubiquitous phenomenon in high-Reynolds-number flows of liquids [1]. In contrast with boiling, in which liquid vaporizes as the temperature rises, cavitation happens due to local pressure reductions that lead to the formation of initially small bubbles. These bubbles respond to the surrounding flow field by growing and collapsing, sometimes with extreme violence. During the collapse, the cavitation bubbles undergo a rapid compression such that the bubble volume decreases by several orders of magnitude [1]. This implosion, usually occurring within a few microseconds, concentrates energy into a small volume, creates regions of high pressure and temperature, emits radially propagating shock waves, and is capable of damaging nearby objects [2,3]. The destructive nature of cavitation erosion is a significant challenge in naval hydrodynamics, e.g., eroding turbine blades, propellers, and rudders [2,3]. On the other hand, if controlled, damage can be exploited for therapeutic purposes in biomedical applications. In the context of therapeutic ultrasound,

the pressure pulses from the collapse of cavitation bubbles are employed to fragment kidney stones, a treatment called shock wave lithotripsy [4]. Owing to its wide range of applications, cavitation erosion has been the topic of numerous studies in the past decades.

Cavitation-induced erosion is a multiphysics and multiscale problem at the intersection of fluid and solid mechanics. The interactions of many bubbles with turbulence, the compressibility effects of the multiphase mixture, and the propagation of shock waves produced by bubble collapse and their interactions with material interfaces are challenging nonlinear and multiscale phenomena in fluid dynamics. Diagnosing these flows experimentally is particularly challenging because of the wide range of spatial and temporal scales, difficult optical access, and intrusiveness of measurement devices. Thus, highly resolved numerical simulations have emerged as beneficial complements to experimental studies, providing valuable insight into the detailed dynamics of the inertial collapse of cavitation bubbles.

METHODS & CODES

To perform the desired simulations, we developed a novel numerical algorithm to solve the three-dimensional compressible Navier–Stokes equations for a multiphase system [5]. This numerical framework prevents spurious pressure and temperature oscillations across the material interfaces and is capable of accurately and robustly representing shock waves and high-




Figure 1: Volume rendering of the bubble shape, colored by its temperature. After the collapse, the bubble takes the form of a hot vortex ring moving in the direction of the jet impact toward the rigid boundary.




Figure 2: Scaling of the maximum pressure (left) and maximum temperature (right) along the wall as a function of initial stand-off distance from the wall for different driving pressures.

density-ratio material discontinuities. For discretization, we developed a solution-adaptive central/discontinuity-capturing approach. Our spatial scheme is of high-order accuracy in smooth regions and is nominally nondissipative; high-order discontinuity-capturing is applied only at sharp gradients detected by a discontinuity sensor. This approach was especially designed to simulate nonspherical dynamics of individual bubbles and the resulting shock waves produced during collapse. An in-house computational code, utilizing Message-Passing Interface for parallelization and Hierarchical Data Format for I/O, was developed in C++ to perform the proposed three-dimensional high-resolution simulations. The code was verified and validated using a suite of problems, and its parallel scaling was demonstrated on Blue Waters.

RESULTS & IMPACT

This project focuses on the detailed dynamics of the collapse of individual gas bubbles near a solid object for different geometrical configurations and driving pressures. We have shown that the presence of a rigid boundary breaks the symmetry of the collapse and leads to the formation of a high-speed re-entrant jet directed toward the neighboring wall [6–8]. The jet impact on the distal side of the bubble generates a water-hammer shock impinging on the adjacent surface that can cause structural damage. The bubble then takes the form of a hot vortex ring convecting toward the object (Fig. 1), and if close enough to the object wall, the rise in surface temperature may lead to thermal damage. We further developed scaling for important collapse properties (e.g., wall pressures/temperatures), in terms of the initial stand-off distance and driving pressure. This not only illustrates the universality of nonspherical bubble dynamics but also provides the means to predict these phenomena (Fig. 2).

Since real flows involve many bubbles, we also simulated the inertial collapse of a pair of gas bubbles near a rigid surface to investigate the bubble–bubble interactions and their effects on collapse dynamics. Through this work, we have fully solved the problem of a single bubble collapsing near a rigid boundary. This problem is the central problem to cavitation damage; solving it will enable more accurate prediction of cavitation-induced damage in naval hydrodynamics, therapeutic ultrasound, and other fields. Based on this information, control strategies can be developed.

In naval hydrodynamics, mitigation strategies would improve overall performance and reduce maintenance costs. In medicine, this knowledge would result in the development of safer and more efficient procedures.

WHY BLUE WATERS

In this project, we utilized an in-house production code for the large-scale simulations. The code foundation is based on high-order accurate algorithms, explicit in time and in space, naturally lending itself to massive parallelization. To carry out accurate three-dimensional simulations of the collapse of cavitation bubbles that effectively resolve the small-scale features of the flow, high resolution (of up to 2.5 billion grid points) is paramount. Performing such simulations requires a substantial computational power that is difficult to achieve on any other NSF-funded computing machines. Therefore, a leading-edge petascale high-performance computing system like Blue Waters is essential for the success of the present study. This project will help us to gain valuable insights and understanding of these complex flows that previously were not possible.

PUBLICATIONS & DATA SETS

Beig, S.A., B. Aboulhasanzadeh, and E. Johnsen, Temperatures produced by inertially collapsing bubbles near rigid surfaces. *J. Fluid Mech.*, under review (2018).

Beig, S.A., and E. Johnsen, Inertial collapse of a gas bubble near a rigid boundary. *J. Fluid Mech.*, in preparation (2018).

Beig, S.A., M. Rodriguez, and E. Johnsen, Inertial collapse of bubble pairs near a rigid boundary. *J. Fluid Mech.*, in preparation (2018).

Beig, S.A., and E. Johnsen, Bubble–bubble interactions and wall pressures/temperatures produced by the collapse of a bubble pair near a rigid surface. *10th Int. Cav. Sym.* (Baltimore, Md., 2018).

Beig, S.A., M. Rodriguez, and E. Johnsen, Collapse of cavitation bubbles near solid surfaces. *32nd Sym. Nav. Hydrodyn.*, (Hamburg, Germany, 2018).

Beig, S.A., and E. Johnsen, Inertial collapse of bubble pairs near a solid surface. *Bull. Am. Phys. Soc.*, (Denver, Colo., 2017).

Rodriguez, M., et al., The role of confinement in bubble collapse in a channel. *Bull. Am. Phys. Soc.*, (Denver, Colo., 2017).

HIGH-RESOLUTION MODELING OF TURBULENCE IN WAVE BOUNDARY LAYERS

Allocation: Illinois/300 Knh
PI: Marcelo García¹
Co-PI: Paul Fischer¹
Collaborators: Dimitrios K. Fytanidis¹, Jose M. Mier Lopez¹

¹University of Illinois at Urbana-Champaign

BW

EXECUTIVE SUMMARY

Wave boundary layer flows play an important role in coastal engineering and coastal sediment transport. However, current state-of-the-art models fail to accurately predict the complex mechanics of the turbulence and momentum exchange between seabed and free-stream velocity under oceanic flow conditions. Recent experimental and numerical studies in our laboratory indicate the presence of a phase lag between the time instance when the maximum bed shear stress occurs with respect to the maximum free-stream velocity in transitional oscillatory boundary layer flows. However, the effect of different levels of bed roughness and porosity remains unknown. This work is the first computational effort to simulate the effect of bed roughness height and bed porosity on the maximum bed shear stress phase difference compared to the maximum free-stream velocity value. It also examined the effect of turbulent flow structures on the bed-turbulent flow interaction, the ensuing sediment transport, and the bottom morphodynamics.

RESEARCH CHALLENGE

Recent experimental and numerical studies from our laboratory suggest the presence of phase lag between the time instances when the maximum bed shear stress occurs with respect to the maximum free-stream velocity in transitional oscillatory boundary layer flows. However, the validity of the phase-lag findings with respect to real seabeds, where different levels of roughness and porosity may exist, remains unknown. It is likely that the turbulent spots, which are arrowhead-shaped turbulent flow structures associated with local bed shear stress peaks and strong turbulent bursting, cause a phase lag to exist in the case of rough seabeds as well.

This work is the first computational effort to simulate the effect of bed roughness height and bed porosity on the maximum bed shear stress phase difference compared to the maximum free-stream velocity value. In addition, it is the first numerical study to quantify the effect of the roughness regime on the turbulent characteristics and quadrant analysis under oscillatory flow conditions. It is also the first study of the turbulent flow over various bottoms with varying porosities in the case of oscillatory flow, and among the first studies that have examined the mixing layer and momentum exchange between the free-stream oscillatory flow and the pore-scale flow under unsteady flow conditions.

METHODS & CODES

Extensive experiments using advanced experimental techniques in the Large Oscillatory Water-Sediment Tunnel at the Ven Te Chow Hydrosystems Laboratory at the University of Illinois at Urbana-Champaign 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 flatbed [1]. This observation is extremely important for the field of environmental fluid mechanics and coastal sediment transport, as this study is the first one in the literature that supports 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. While advanced measurement techniques have been developed over the years, there is still no other available measurement technique that can give adequately

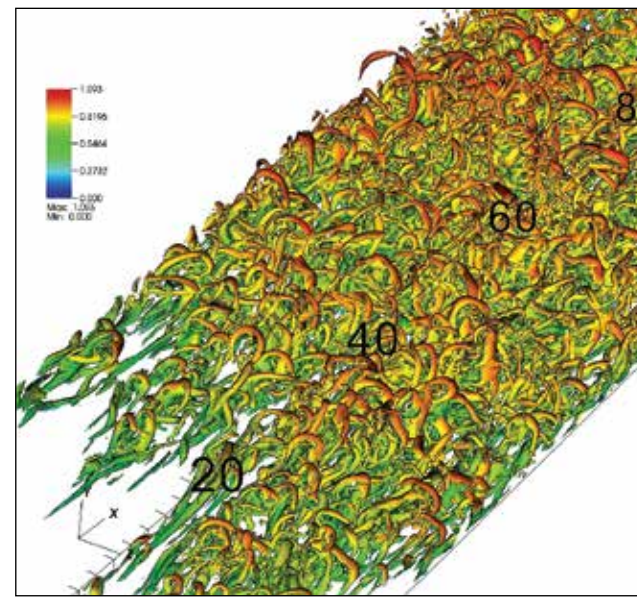


Figure 1: Turbulent flow structures (hairpin vortices) in oscillatory boundary layer flow (colored with the normalized velocity magnitude).




Figure 2: Flow-bed interaction in the case of oscillatory flow above hemispherical roughness elements.

detailed information, both spatially and temporally, to estimate all the needed parameters for the three-dimensional turbulence structures under an oscillatory flow. Thus, turbulence in complex three-dimensional flows is characterized mainly numerically, using a combination of high-order numerical methods and highly scalable numerical codes with backing from accurate experimental observations from the lab.

We developed a Direct Numerical Simulation model capable of simulating the complex oscillatory boundary layer (OBL) flow and sediment transport using the Spectral Element Method (SEM) 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 a Eulerian approach and proper boundary conditions for the sediment mass exchange between the coastal bed and the free-stream flow [3,4].

RESULTS & IMPACT

Our work is the first in the literature that explores the effect of turbulence characteristics of the flow and, particularly, the turbulent flow structures (Fig. 1) on the phase difference between maximum bed shear stress and free-stream velocity. It also explains the complex interaction between flow and bed for various bed

characteristics (roughness and porosity: Fig. 2). It is expected that these previously unexplored physics related to the oceanic boundary layer and sediment transport will have a great impact on improving the accuracy of modern engineering tools and models used in everyday practice.

Our research will lead to a deeper understanding of the interaction of oscillatory turbulent flow, bed shear stress, and sediment mass transport, and eventually, will lead to the development of new, simplified but accurate models for the analysis and design of engineering systems in coastal and oceanic environments.

WHY BLUE WATERS

Such an analysis pushes the limits of turbulence-resolving numerical studies in terms of the computational resources and the high-performance computing facilities it requires and thus it can only be materialized on a petascale supercomputer such as Blue Waters. The work combines the expertise of Prof. Marcelo García's group from the Ven Te Chow Hydrosystems Laboratory of the Civil and Environmental Engineering Department at the University of Illinois at Urbana-Champaign and Prof. Paul Fischer's group from the Illinois Computer Science and Mechanical Engineering Departments with the leading-edge petascale computing resources of Blue Waters.

COARSE-GRAINED FORCE FIELD FOR IONIC LIQUIDS

Allocation: Illinois/200 Knh
PI: Narayana Aluru¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Ionic liquids have potential applications in energy, healthcare, and nanotechnology. However, a fundamental understanding of their behavior is limited by the high computational cost of their simulation using all-atom models. With the goal of furthering that fundamental understanding, we developed a coarse-grained model of ionic liquids to capture their heterogeneous structure. In addition, we developed a new method to compute their thermodynamical properties, such as density and pressure, which in turn enables their simulation in different systems. Furthermore, ionic liquids have long-range electrostatic interactions, which give rise to many distinct physical properties compared to conventional solvents. In addition, we extended our model to include the long-range interactions in a systematic procedure between coarse- and fine-grained models, making it applicable to any class of ionic liquids or other molecular systems. We applied our method to imidazolium-based ionic liquids. We demonstrated that the force field is transferable with high accuracy to different thermodynamic states and to various alkyl chain lengths.

RESEARCH CHALLENGE

Ionic liquids are known as “solvents of the future.” However, their very slow dynamics and long-range interactions hinder fundamental understanding of their behavior using all-atom models, which makes their simulation prohibitively

computationally expensive. Therefore, we developed structure-based coarse-grained force fields capturing their structure and thermodynamic properties. Past studies show that coarse-grained force fields can be used to understand underlying physical phenomena in supercapacitors, nanolubricants, and the like [1,2]. So far, however, coarse-grained models have been developed using hand tuning with little information from a realistic model of ionic liquids.

The method we developed in this study is generalizable for any ionic liquid and solvent, and systematically connects all-atom and coarse-grained models. Recent studies also indicate that the behavior of ionic liquids needs simulation of large systems for long time periods to capture physical phenomena occurring at the nanoscale [3,4].

METHODS & CODES

We performed Molecular Dynamics (MD) simulations on Blue Waters using the GROningen MACHine for Chemical Simulations (GROMACS) MD package. GROMACS allows large-scale simulation of molecular systems with different potential forms. Furthermore, the user can determine potential for specific beads (usually coarse-grained beads) by using the Table option. We calculated the long-range Coulombic interactions with the Particle Mesh Ewald (PME) method for higher accuracy and speed. The all-atom models consist of ~100,000 atoms, while the coarse-grained models consist of ~20,000 atoms. We used the Versatile Object-oriented Toolkit for Coarse-graining Applications (VOTCA), developed at the Los Alamos National Laboratory, for coarse-graining procedures with at least 200 iterations for each model to obtain the coarse-grained force fields.

RESULTS & IMPACT

In this study, we extended the relative entropy coarse-graining method with the addition of a constraint to the objective function and consideration of the long-range electrostatic interactions. We modified the VOTCA code using the Lagrange multiplier method to compute thermodynamic properties in addition to the structure. The thermodynamic properties allow the scientific community to study coarse-grained models in different systems, such as an isothermal–isobaric ensemble with a realistic behavior. Furthermore, we also considered the long-range electrostatic interactions in a systematic manner during the coarse-graining procedure. The electrostatic interaction is the driving force of many biological processes such as protein folding and DNA–protein interactions; therefore, the electrostatic interaction should

be optimized systematically and accurately. According to [5], 70% of the interactions in ionic liquids originate from electrostatic interactions, so optimization of the electrostatic interaction using the current method can pave the way for more rigorous and accurate simulation of ionic liquids.

Understanding many physicochemical phenomena in biological and physical systems is also possible using the current force field parameters with higher accuracy. Many previous studies used toy models to understand these phenomena due to a lack of systematic coarse-grained force fields and the prohibitive computational cost of all-atom models. We have shown that the force fields are transferable for different thermodynamic states and various alkyl chain lengths. The current force field and method can pave the way for a fundamental understanding of ionic liquids and other solvents in a computationally feasible and physically consistent method.

WHY BLUE WATERS

We performed large-scale all-atom molecular dynamics simulations for multiple alkyl chain lengths of imidazolium-based ionic liquids, which are highly charged systems. Therefore, their simulation requires both annealing and equilibrium simulations for long time periods to avoid getting trapped in a nonequilibrium state. All-atom models have at least ~100,000 atoms and are performed for different temperatures and multiple alkyl chain lengths for at least ~30 nanoseconds (ns). We performed coarse-grained molecular dynamics simulations and coarse-grained force field optimization in an iterative manner, requiring simulation of the coarse-grained system with at least ~20,000 beads for at least 5 ns and 200 iterations. GROMACS scalability on Blue Waters and the VOTCA package usage of multiple nodes made the simulation and coarse-graining procedures feasible.

PUBLICATIONS & DATA SETS

Moradzadeh, A., et al., Coarse-Grained Force Field for Ionic Liquids. *Journal of Chemical Theory and Computation*, 14:6 (2018), pp. 3252–3261.




Figure 1: Schematic representation of all-atom and coarse-grained model for [C₆MIM]⁺ [PF₆]⁻ with reduced degrees of freedom and similar statistical behavior.

HIGH-FIDELITY NUMERICAL SIMULATIONS OF TURBINE VANE HEAT TRANSFER UNDER FREESTREAM TURBULENCE

Allocation: Director Discretionary/62.5 Knh
PI: Sumanta Acharya¹
Collaborator: Yousef Kanani¹

¹Illinois Institute of Technology

EXECUTIVE SUMMARY

Designing efficient and reliable gas turbines requires detailed knowledge of the aerodynamics and heat transfer over the turbine surface. Since experiments at realistic engine temperatures (~3000°F), high pressures (20–40 bars) and rotor speeds (~10,000 rpm) are not possible, designers have to rely on data that can be collected over lower parameter ranges or from simulations. Accurate predictions are stymied by the complexities in the flowfield that include freestream turbulence and wakes, high strain rate, high flow acceleration, and the like. Time-averaged equations require modeling and are generally inaccurate. Higher-fidelity simulations, as in the present study, require large computing resources. Rapid advances in the capability of supercomputers such as Blue Waters allow researchers to perform high-fidelity simulations in complex geometries under realistic conditions, and to accurately predict and obtain a fundamental understanding of the flow and heat transfer physics that can then be used for rational design of engineering models.

RESEARCH CHALLENGE

The flow upstream of the turbine airfoil is typically turbulent and characterized by large length scales. The boundary layer on the pressure and suction surfaces of the airfoil experiences this high level of freestream turbulence and can transition into a turbulent boundary layer downstream on the surface. Predicting the transition to turbulence typically requires a very fine grid to resolve the smallest scales in turbulence; this, in turn, requires

a very large mesh (on the order of 50–150 million nodes for a realistic Reynolds number) and can make the computations time consuming, even on large supercomputers. Such computations are, therefore, out of reach for most of those in industry, where using Reynolds-averaged turbulence models is commonplace despite their lower accuracy. The available transitional models are not universally applicable in different geometries and typically perform poorly. To capture the transition to turbulence, the finest scales available in the fluid flow should be resolved, which requires significant computational power.

In the current study, the flow and heat transfer over a turbine vane under realistic engine conditions of high turbulence are simulated using Large-Eddy Simulation (LES). We aim to correctly predict the flow and heat transfer characteristics on both the suction and pressure side of the vane for different levels of elevated freestream turbulence. Experimental observations suggest that the suction-side boundary layer transitions to turbulence, with the transition location moving upstream at higher turbulence levels. Heat transfer also is enhanced on the pressure side of the vane prior to the transition to turbulence.

METHODS & CODES

The vane geometry selected for the computational study is identical to the experiments conducted by Varty and Ames [1]. We studied a linear cascade, which is represented in the computations by a single vane with periodic boundary conditions imposed in the pitch direction. The periodic boundaries are separated by

the pitch $P=0.773L$, where L is the chord length and is selected as the length scale. To capture the secondary flow effects, the numerical domain includes the whole span of the vane between two endwalls with the spanwise length of $0.5L$.

We selected the exit chord Reynolds number of $Re=510,000$ and inflow turbulence levels of 0.7%, 7.9%, and 12.44%, with length scales of 0.16, 0.04, and 0.074 of the chord, also as in the experimental study of Varty and Ames [1]. The inlet boundary is located at $x=-0.51L$ where L is the chord length. We used the Synthetic Eddy Method (SEM) [2] to generate isotropic and homogenous velocity perturbations with the specified turbulence characteristics (i.e., mean velocity, turbulence intensity, and length scale).

We treated the outflow boundary with the homogeneous Neumann condition for both velocity and pressure; it is located $0.4L$ away from the trailing edge axially, which allows the wake to exit the domain in the periodic direction. All walls are treated with the no-slip condition for the velocity. We imposed a constant heat flux on the vane for the temperature solution.

We solved the incompressible Navier–Stokes equations and the temperature equations using the Pressure-Implicit with Splitting of Operators algorithm with the OpenFOAM flow finite volume code. Air is assumed as the working fluid. We discretized all spatial derivatives with a second-order central scheme. The central difference scheme is blended with a second-order upwind scheme for the convective terms to ensure the stability of the numerical simulations. We used the Quadratic Upstream Interpolation for Convective Kinematics method to approximate the convection term in the passive scalar equation. We used the second-order backward scheme to discretize the temporal derivatives. We conducted the LES using the WALE subgrid-scale stress model.

The heat transfer augmentation prediction over the pressure side of the vane is less computationally demanding than the prediction of the transition to turbulence and requires a mesh with 11 million grid points. Simulations capable of predicting transition to turbulence require large grid sizes in the order of 100 million grid points, which is intensive both computationally and in terms of storage.

RESULTS & IMPACT

Current numerical simulations are capable of accurately predicting the heat transfer augmentation on the leading edge and over the pressure side of the vane (Fig. 1). We have clearly described the mechanism on which this augmentation occurs for the first time. The freestream turbulence generates vortex tubes wrapped around the leading edge (Fig. 1), which perturbs the velocity inside the boundary layer and forms low- and high-speed streaks. These streaks retain their signature downstream up to the trailing edge of the vane. There is a strong correlation between the instantaneous heat transfer and the velocity streaks.

We are using a more refined grid to predict the transition to turbulence on the suction side. Formation of turbulent spots on the suction side of the vane is captured with 88 million grid points




Figure 2: (Top) Formation of turbulent spots visualized by instantaneous velocity fluctuations close to the surface and (bottom) its effect on surface heat transfer on the suction side of the vane.

(Fig. 2). Turbulence spots form and grow in size on the suction side of the vane and eventually trigger the entire boundary layer into a turbulent state. Predictions agree with the measurements and transition locations [1]. We have also correctly predicted the heat transfer enhancement due to the turbulence spots.

WHY BLUE WATERS

Blue Waters provides a powerful tool to conduct high-fidelity simulations. Without accessing the computational resources at the scale that Blue Waters provides, performing these simulations is practically impossible. Further, analyzing the temporal behavior of the flow and heat transfer field requires data storage of terabytes of data, which is made possible by both short- and long-term storage solutions provided by Blue Waters.

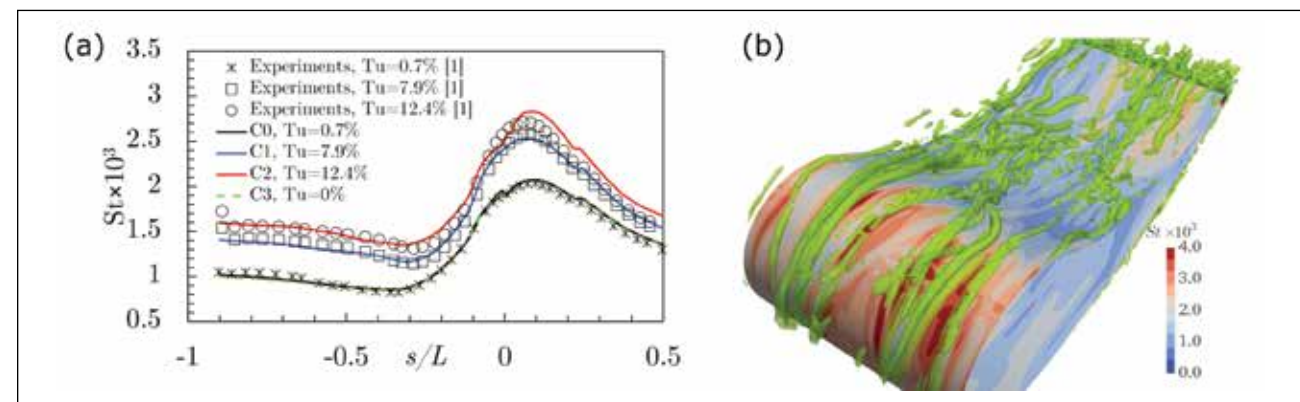


Figure 1: (a) Stanton number distribution near the leading edge and pressure side; and (b) flow structures and instantaneous Stanton distribution over the turbine vane.

MECHANISTIC INSIGHTS INTO HYDRATION OF SOLID OXIDES

Allocation: Strategic/492 Knh
PI: Narayana Aluru¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Some solid oxide materials used in solid oxide fuel and electrolysis cells are known to conduct protons once the materials are hydrated. However, the mechanisms by which solid oxide materials become hydrated are not clear. By performing detailed density functional theory (DFT) calculations, we investigate hydration of two typical solid oxides with a single-crystal structure—a proton-conducting yttrium-doped strontium zirconate (SZY) and an oxide ion-conducting yttria-stabilized zirconia (YSZ). We suggest a four-step process to understand the hydration of solid oxides—water adsorption on the surface, proton migration from the surface to bulk, proton migration in the bulk, and oxide ion vacancy migration in the bulk. Our analysis provides mechanistic insights into the hydration of single-crystal SZY and nonhydration of single-crystal YSZ. Our study not only explains the hydration of materials, but also illustrates the importance of structural rearrangement when a proton is incorporated into the bulk of the solid oxide material.

RESEARCH CHALLENGE

Hydrogen production has attracted great interest as it is a clean and sustainable fuel and is a promising choice for the storage of intermittent renewable energies such as solar and wind power [1]. Water electrolysis using renewable energy is considered to be one of the cleanest methods to produce hydrogen [2], and electrolysis cells are widely used for this purpose. Compared to a commercially available polymer electrolyte membrane (PEM)-based low-temperature electrolysis cell, a solid oxide electrolysis cell (SOEC) operating at a high temperature can take advantage of thermal energy to reduce electrical energy demand. This results in a reduced cost and enhanced efficiency for hydrogen production [3]. However, the design and development of solid oxide electrolytes with sufficient stability and enhanced conductivity for SOEC are still a challenge.

Oxide ion conductors are commonly used as electrolytes in SOECs. For example, ZrO₂ doped with Y₂O₃ (YSZ) exhibits sufficient oxide ion conductivity as well as thermal and chemical stability at high temperature [2]. Comparing SOECs with oxide ion-conducting electrolytes, SOECs with proton-conducting electrolytes such as Y-doped SrZrO₃ (SZY) can produce pure and dry hydrogen. In order to design high-conductivity proton conductors, a fundamental understanding of what determines whether an oxide ion conductor can become a proton conductor is needed. The origin of proton conduction, which is attributed to the hydration process, is of significant importance for the development

of proton conductors. The hydration reaction can be written in Kröger-Vink notation [4] as $H_2O(g) + V_o^{\times} + O_o^{\times} \rightarrow 2OH_o^{\times}$, where V_o^{\times} , O_o^{\times} , and OH_o^{\times} denote the oxygen vacancy, lattice oxygen, and protonic defect in the hydrated structure, respectively. Significant effort has gone into correlating hydration thermodynamics (hydration enthalpy and entropy) and material properties such as the type and concentration of acceptor dopant on the B-site of perovskites [6,7], the difference in electronegativity between B- and A-site constituents in perovskites [5], and chemical expansion of perovskites [8]. However, the existing approaches have not been able to describe the detailed dynamic process of hydration. In this work, density functional theory (DFT) calculations are performed to investigate the hydration behavior of solid oxides.

METHODS & CODES

All DFT calculations are performed using the Vienna Ab initio Simulation Package (VASP) [9–11]. The Perdew–Burke–Ernzerhof [12] exchange–correlation functional is employed based on the projector augmented wave method [11]. The cutoff energy for the plane-wave basis set was 500 eV for all calculations. All the calculations are spin polarized. Bader population analysis is performed to calculate the atomic charges [13]. The migration energy barriers are calculated using the climbing image nudged elastic band method [14].

RESULTS & IMPACT

Based on DFT calculations, we investigated the hydration processes in two different solid oxides with a single-crystal structure—a proton-conducting SZY and an oxide ion-conducting YSZ.

The water adsorption calculations show that both the SrO (001) surface of SZY and (111) surface of YSZ are favorable for the dissociative adsorption of water. The structural minimization and charge analyses indicate that the water molecule is adsorbed as two hydroxide ions on the surface. The calculations for proton migration from the surface to bulk show that the energy barrier for YSZ is larger than that of SZY for both hydroxide ion migration and proton hopping. In addition, the energy barriers for proton migration in both bulk SZY and YSZ structures are relatively small and comparable. The energy barrier for the oxide ion vacancy migration in bulk SZY is slightly larger than that of bulk YSZ, but the results confirm that bulk SZY shows oxide ion conductivity. The migration of the oxide ion vacancy from the bulk to the surface of the oxide ensures that water is not only adsorbed on the surface but also hydrates the interior (in the case of SZY). Our results




Figure 1: NEB results for proton migration from the surface to bulk of perovskite structure.

indicate that the primary difference between the hydration of SZY and YSZ comes from the energy barriers for proton migration from the surface to the bulk of the oxide material.

Our results not only provide mechanistic insights into the hydration of single-crystal SZY and nonhydration of the single-crystal YSZ, but also the importance of structural rearrangement when a proton is incorporated into the bulk of the solid oxide material.

WHY BLUE WATERS

In our work, large-scale *ab initio* simulations are needed to obtain the dynamics of hydration in solid oxides. For the proposed DFT calculations, using eight to 10 nodes (256 to 320 cores) for each job can achieve the best performance, which is attributed to the power of Blue Waters and the support of project staff. Thus, we could run *ab initio* simulations on Blue Waters easily and quickly and speed up our research greatly.

PUBLICATIONS & DATA SETS

Jing Y., H. Matsumoto, and N.R. Aluru, Mechanistic Insights into Hydration of Solid Oxides. *Chem. Mater.*, 30 (2018), pp. 138–144.

MACHINE LEARNING FOR HIGH-ENERGY PHYSICS: PARTICLE IDENTIFICATION AND REGRESSION USING DEEP NEURAL NETWORKS

Allocation: Illinois/125 Knh
PI: Benjamin Hooberman¹
Co-PI: Amir Farbin²
Collaborators: Ryan Reece³, Matt Zhang¹

¹University of Illinois at Urbana-Champaign
²University of Texas at Arlington
³University of California at Santa Cruz

EXECUTIVE SUMMARY

The Large Hadron Collider (LHC) at CERN, in Switzerland, is the world’s most powerful particle accelerator. The LHC recreates the conditions of the universe one tenth of a nanosecond after the Big Bang by colliding together protons traveling at 0.99999997 times the speed of light 40 million times every second. Each proton–proton collision creates up to several hundred particles that pass through one of four detectors situated at the LHC interaction points. Reconstructing the collisions requires identifying these particles using their signatures in the detector. Recent advances in machine learning and artificial intelligence, which are also known as deep learning, have made it possible to apply learning networks to many kinds of problems. In particular, identifying particles from their energy deposition in calorimeter cells bears a strong resemblance to problems in machine vision, in which objects are reconstructed from intensity values in pixel arrays. We exploit deep learning techniques to identify and measure particles produced at colliders and find that they provide improvements in performance with respect to conventional methods.

RESEARCH CHALLENGE

As previously mentioned, the Large Hadron Collider (LHC) at CERN in Switzerland recreates the conditions of the universe a tenth of a second after the Big Bang by colliding together high-energy protons. In 2012, the Higgs boson was discovered in LHC data, completing the standard model of particle physics

and leading to the Nobel Prize in Physics in 2013. This discovery transformed our understanding of the building blocks of matter and the fundamental forces by explaining the origin of the masses of subatomic particles. However, the standard model is not capable of resolving key open questions and thus cannot be the final theory of nature. In particular, it cannot explain the origin of dark matter, which comprises about five times as much total mass in the universe as visible matter but whose nature is not understood. Various beyond-the-standard-model scenarios, including supersymmetry and extra dimensions of spacetime, have been posited to resolve these problems. These scenarios generically predict the existence of exotic new particles, which may be produced at particle colliders such as the LHC. Searching for these particles to understand the nature of physics beyond the standard model is now the highest priority of the LHC physics program and the focus of this project.

Analyzing LHC data to search for physics beyond the standard model requires identifying and measuring the particles produced in proton–proton collisions. Particles produced in collisions traverse detectors, depositing their energy in calorimeters consisting of a granular array of detecting elements (“pixels”). The resulting “image” can be analyzed to distinguish among the six species of stable particles (electrons, photons, charged hadrons, neutral hadrons, and muons) and infer their energies. Electrons and photons are expected signatures of a wide variety of interesting new physics scenarios, but may be mimicked by charged and

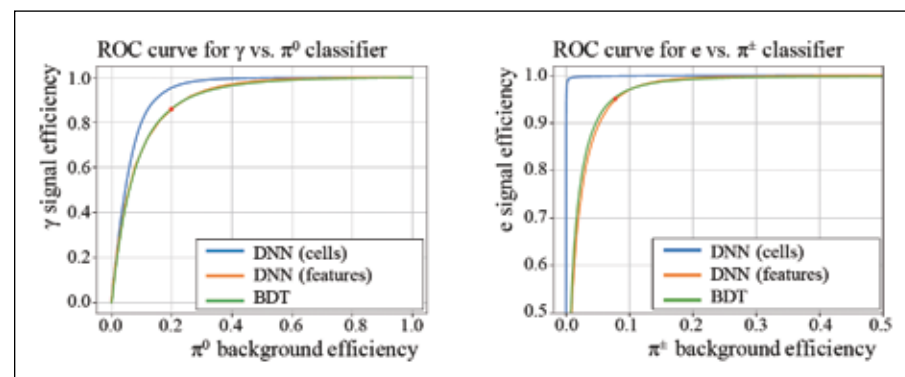


Figure 1: Signal vs. background efficiency Receiver Operating Characteristic (ROC) curves for (left) photon vs. neutral pion and (right) electron vs. charged pion discrimination.

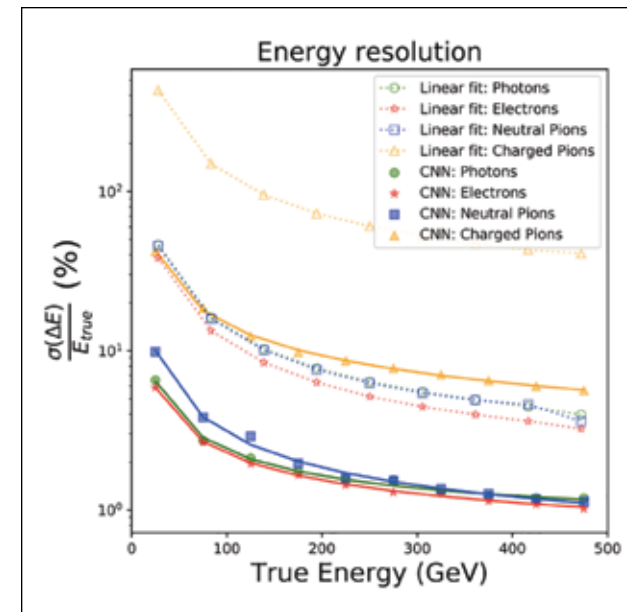


Figure 2: The relative energy resolution of the four particle types versus the true particle energy.

neutral hadrons, which are produced at rates that are higher by several orders of magnitude. Since each collision typically contains thousands of particles, discriminating signals from electrons and photons from hadronic backgrounds is complicated by the presence of additional overlapping particles. Identifying and measuring electrons and photons, especially those with low energy, is thus a major challenge of high-energy physics.

METHODS & CODES

Recent advances in machine learning and artificial intelligence, known as deep learning, have made it possible to apply learning networks to many kinds of problems. These techniques are driven by the emergence of large data sets, powerful Graphical Processing Unit (GPU) processors, and new techniques to train billion-neuron multi-layer artificial neural networks (NN). In computer vision, Deeply-connected Neural Networks (DNN) and Convolutional Neural Networks (CNN) have provided dramatic improvements in performance and speed with respect to conventional algorithms and require minimal engineering.

We employ DNNs and CNNs to distinguish among signals from electrons and photons and hadronic backgrounds and measure particle energies. We simulate samples of individual electron, photon, charged hadron, and neutral hadron images in a simple high-granularity calorimeter detector implemented with the Geant4 simulation toolkit. These images are used to train NNs using PyTorch, which distinguishes among electrons versus charged hadrons and photons versus neutral hadrons and measures the energies of the four particle species. To optimize the network architectures, we vary the NN “hyperparameters,”

including the number of NN layers (“depth”), number of neurons per layer (“width”), and the learning and dropout rates.

RESULTS & IMPACT

We have evaluated the performance of DNNs and CNNs trained on particle images, and compared the results to the current state-of-the-art algorithms widely used in particle physics. These algorithms employ NNs and Boosted Decision Trees (BDTs) to analyze a precomputed set of particle “features” such as the calorimeter shower depth and width. For both classification and energy measurement using regression, we find that the deep NNs provide significant improvements compared to the conventional methods. These results serve as a first step toward implementing deep learning for particle identification and measurement at the LHC.

WHY BLUE WATERS

Optimizing the network performance using hyperparameter scans requires retraining NNs hundreds or thousands of times, which is especially challenging for memory-intensive networks such as GoogLeNet or ResNet. The 4,228 GPU-enabled XK nodes with 25 TB of GPU accelerator memory available on Blue Waters enable training and optimization of neural nets beyond what has previously been achieved, allowing for detailed investigations of their behavior for both particle physics and general applications.

PUBLICATIONS & DATA SETS

Hooberman, B., et al., Calorimetry with Deep Learning: Particle Classification, Energy Regression, and Simulation for High-Energy Physics. *Proceedings of the Deep Learning for Physical Sciences Workshop at the 31st Conference on Neural Information Process (Long Beach, Calif., December 8, 2017).*

PETASCALE SIMULATIONS OF GRAPHENE NANORIBBON DEVICE STRUCTURES

Allocation: NSF PRAC/4,200 Knh

PI: Jerry Bernholc¹

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

Collaborator: Zhongcan Xiao¹

¹North Carolina State University

EXECUTIVE SUMMARY

This project focuses on high-performance calculations for materials and devices of high current interest, as well as on development of petascale methods for such simulations. In FY'18, we concentrated on three applications: (1) metallic contacts to 1D graphene nanoribbons (GNRs), which were established by varying their widths. While 7-atom-wide armchair GNR (7-aGNR) is a semiconductor, 14-aGNR and 21-aGNR are metallic and form a seamless staircase contact to 7-aGNR. (2) We also worked on atomic-level stability of both zigzag and armchair edges of GNRs against oxidation. In a joint experimental–theoretical investigation, we showed that hydroxyl and epoxy species were produced by oxidation, leading to a significantly reduced bandgap of the 7-aGNR. We mapped out the evolution of edge structures and revealed the reaction path and the produced chemical species of both for the zigzag and the armchair edges. (3) We also investigated potential GNR-based devices and proposed a novel structure that can be manufactured with atomic precision and exhibits significant negative differential resistance.

RESEARCH CHALLENGE

Electrical contact to low-dimensional (low-D) materials is a key to their electronic applications. Traditional metal contacts to low-D semiconductors create gap states that can pin the Fermi level (E_F). However, low-D metals possessing limited density of states at E_F can enable gate-tunable work functions and contact barriers. Moreover, a seamless contact with native bonds at the interface that does not introduce in-gap states can act as an optimal contact electrode. Indeed, an ohmic contact has been achieved in 2D transition metal dichalcogenides at the interface of the metallic 1T'-phase and the semiconducting 2H-phase. We demonstrated an all-carbon staircase contact to ultra-narrow graphene nanoribbons (GNRs), a wide bandgap 1D semiconductor. A combined scanning tunneling microscopy and density functional theory study revealed the native covalent-bond nature and the quasi-metallic contact characteristics of the interfaces. The GNR staircase points constitute a promising route to high-performance graphene nanoelectronics.

The stability of graphene nanoribbons (GNRs) against oxidation is critical for their practical applications. We systematically studied the thermal stability and the oxidation process of ambient-exposed armchair GNRs with a width of seven carbon atoms (7-aGNR),

grown on an Au(111) surface. The atomic scale evolution of the armchair edges and the zigzag ends of the aGNRs after annealing at different temperatures were revealed by using scanning tunneling microscopy, Raman spectroscopy, X-ray photoelectron spectroscopy, and first-principles calculations. In the oxidation process, O_2 molecules first dissociate and then react with carbon atoms at the edges. Two different oxygen species were identified at the armchair edges, namely the hydroxyl pair and epoxy with one bridge oxygen bonded to two edge carbons. These oxidation species modify the electronic properties of pristine 7-aGNRs, with a bandgap reduction from 2.6 eV to 2.3 eV and 1.9 eV for the hydroxyl pair- and epoxy-terminated edges, respectively. These findings revealed that both the zigzag and armchair edges of GNRs can be oxidized after thermal annealing and that the oxidation can greatly affect the electronic properties of GNRs, which also opens an opportunity for using GNRs as high-temperature oxygen sensors.

The development of atomically precise synthesis of graphitic nanostructures promises a revolution in device design to deliver novel electronic functions. A prominent example is the negative differential resistance (NDR) device, for which many designs based on graphene nanoribbons (GNRs) were proposed. However, a controllable fabrication method of such devices with atomic precision has yet to emerge. We proposed a practical device structure, based on armchair GNRs, which exhibits strong NDR effect. Our computational evaluation of the traditional resonant tunneling diode uncovers important issues at the atomic scale concerning the need to minimize the direct tunneling current between the leads while achieving high peak current. The proposed device consists of a short semiconducting GNR and its lower-gap intermediate segments, which enables high current by the alignment of levels across the segments while enlarging the tunneling distance between the leads.

METHODS & CODES

The electronic structure calculations used the Real-space-MultiGrid (RMG) and Quantum Espresso codes. The quantum transport nonequilibrium Green's function calculations used the localized orbitals and nonequilibrium Green's function (NEGF) branches of the RMG code.

RESULTS & IMPACT

Our results indicate that GNRs with variable widths and seamless heterostructures can be realized by using only one molecular precursor. This is different from previous reports where type-I and type-II semiconducting heterojunctions were synthesized with several nanometers of GNRs by using two different precursors. The lack of strongly localized interfacial states is highly desirable for making a good electrical contact to the semiconductor material, as verified in 2D heterojunctions between the metallic 1T'-phase and the semiconducting 2H-phase transition metal dichalcogenides (MoS_2 and $MoTe_2$). The conjugated wide GNRs are either quasi-metallic or small-gap semiconductors with limited DOS at the E_F , similarly to the 2D metals. They can be used as electrical contacts to the ultra-narrow semi-conducting aGNRs, which would avoid Fermi-level pinning and a high Schottky barrier. Furthermore, a staircase with gradually diminishing bandgaps can potentially be used to direct charge transport, while the long and wide branches in the GNR heterojunctions (HJs) facilitate their integration with more conventional contact electrodes. The staircase HJs with atomically controlled interfaces thus provide a promising interconnect to the semiconductor gate material.

We found that GNR oxidation leads to the formation of hydroxyl- and epoxy-group-terminated edges, which lead to significantly reduced bandgaps of GNRs. The significant changes of electronic properties suggest that aGNRs can be used as high-temperature oxygen sensors. The finding that both zigzag ends and the armchair edges can be oxidized by thermal annealing may also help in understanding oxidation of other graphitic structures.

Utilizing the type-I band alignment between the 7-aGNR and the polymer-GNR intermediate, we designed an experimentally realizable nanoscale NDR device with a practical peak-to-valley current ratio and peak current. The computationally guided design uncovers new aspects important for atomic scale devices concerning the need to minimize direct tunneling between the leads while maintaining sufficient peak current and peak-to-valley current ratio. Starting from the concept of a resonant tunneling diode and controlling the confinement and interfacial levels energy matching, we discovered a new, broadly applicable design for atomic-scale ballistic NDR devices, based on multiple narrow segments of low and high bandgap atomic-scale constituents.

WHY BLUE WATERS

The applications described above require a very large parallel supercomputer with a high-speed interconnect among the nodes (due to frequent exchange of substantial amounts of data among nodes). Each project required many runs to explore its various scientific issues, with a substantial amount of analysis between the runs. High availability and quick turnaround are thus also very important for timely progress in our research.

PUBLICATIONS & DATA SETS

Ma, C., et al., Seamless Staircase Electrical Contact to Semiconducting Graphene Nanoribbons. *Nano Letters*, 17 (2017), DOI:10.1021/acs.nanolett.7b02938.

Ma, C., et al., Oxidation stability of atomically precise graphene nanoribbons. *Phys. Rev. Mat.*, 2 (2018), DOI:10.1103/PhysRevMaterials.2.014006.

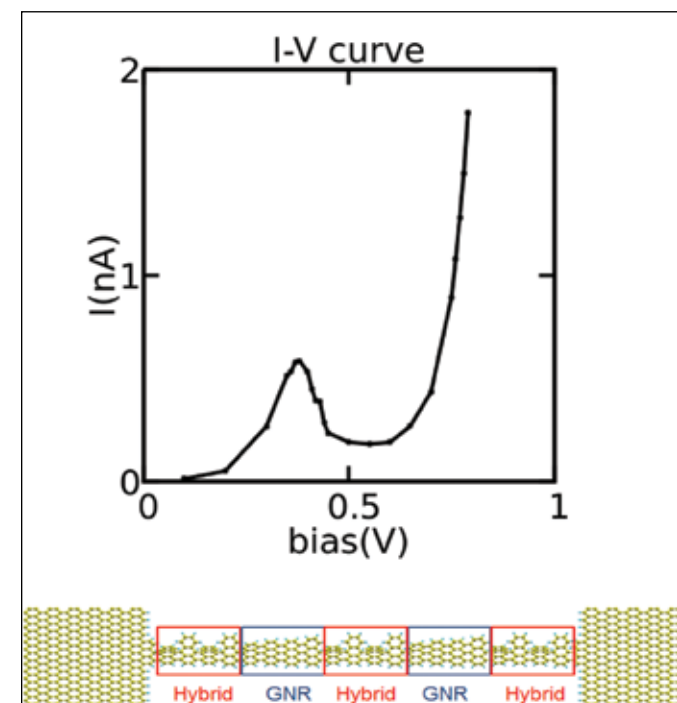


Figure 1: Calculated I-V curve and structure of a negative differential resistance device based on graphene nanoribbon and graphene nanoribbon precursors.

INTERFACIAL LOAD TRANSFER MECHANISMS IN CARBON NANOTUBE-POLYMER NANOCOMPOSITES

Allocation: Illinois/240 knh
 PI: Huck Beng Chew¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Carbon nanotubes (CNTs) are highly promising for strength reinforcement in polymer nanocomposites, but conflicting interfacial properties have been reported by single nanotube pullout experiments. The petascale Blue Waters supercomputing resources have enabled us to investigate the interfacial load transfer mechanisms during pullout of CNTs from poly(methyl methacrylate), or PMMA, matrices, using massively parallel molecular dynamics simulations. By building atomistic models at experimentally relevant length scales, we showed that the pullout forces associated with nonbonded interactions between CNT and PMMA are generally small, and low-density distribution of crosslinks along the interface increases the pullout forces by an order of magnitude. We identified that mechanical unfolding and pullout of single or pair polymer chains attached to the individual crosslink bonds resulted in substantial interfacial strengthening and toughening while contributing to interfacial slip between CNT and PMMA. Our molecular dynamics (MD) simulation results combined with a shear-slip model compared well with experiments and provided mechanistic insights into the design of nanocomposites.

RESEARCH CHALLENGE

Polymer nanocomposites show substantial property enhancements at much lower filler densities compared to conventional polymer composites reinforced with micron-scale fillers, which ultimately results in lower component weight. These property enhancements arise from the orders-of-magnitude higher surface-to-volume ratio of the nanoscale filler particles for interfacial load transfer. Single nanotube pullout tests are

ideal measurements to quantitatively characterize the interfacial load transfer between CNT and the polymer matrix involve the use of atomic force microscopy (AFM) probes *ex situ*, or are conducted *in situ* within a scanning electron microscopy (SEM) chamber. However, conflicting interfacial shear strengths (IFSS) between CNTs and the polymer matrix have been reported by single nanotube pullout experiments conducted over the past decade. Some of these experiments have reported saturation in the critical pullout force with the embedment length of the nanotube, implying a shear lag effect that cannot be predicted by analytical models developed earlier that incorporated pure van der Waals forces. While MD simulations have the potential to provide more realistic insights, most to date have been restricted to relatively short ~10-nm model length of the nanotubes, compared to the actual nanotube embedment lengths spanning ~100 nm to several microns in experiments. The computational barrier to study all-atom models of such CNT-polymer systems at these length sales has hindered the understanding of fundamental mechanisms that could explain the experimental observations.

METHODS & CODES

Our MD simulations were performed using the open source MD code LAMMPS. We adopted the DREIDING force field potential to describe the interactions among all atoms in our single-walled CNT-PMMA system. The atomic configuration of an uncrosslinked, amorphous PMMA matrix was generated using a rapid heating and quenching technique (Fig. 1a, b) by replicating a single PMMA polymer chain generated by Polymer Modeler. The long single-PMMA chain was comprised of 500 monomers and was used to represent PMMA with a molecular weight of

Figure 1: Atomic configurations of a 500-nm-long SWCNT (blue) in a PMMA matrix (grey) at (a) initial stage and (b) heating at 1,000K, one of multiple heating-quenching stages adopted to ensure proper condensation of the PMMA system. (c and d) Pullout of pristine SWCNT from PMMA matrix with initial lengths of 100 nm.

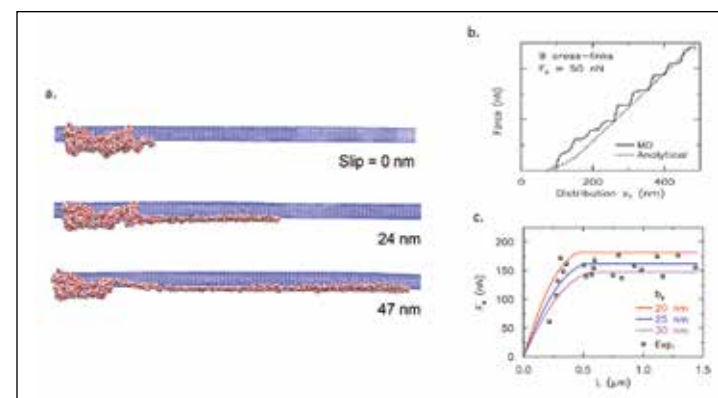



Figure 2: (a) Close-up views of polymer chains attached to crosslinks at various slip-lengths; configurations are filtered to display only the SWCNT and the crosslinked polymer chain. (b) Axial force distribution along nanotube with nine crosslinks. (c) Comparison of model predictions for nanotubes versus experimental pullout data.

50,000 g/mol in experiments. A continuous configurational bias Monte Carlo algorithm was used to randomly assign backbone dihedral angles, which ensured statistically variant representations of monomers in the chain. We equilibrated the composite system until the polymer reached a steady-state density of ~1.02 gm/cc. Quasistatic force-controlled pullout of the nanotube (Fig. 1c, d) was simulated by applying an external energy gradient and then equilibrating the system at 300K for a particular span of time.

RESULTS & IMPACT

We quantified the contribution of nonbonded van der Waals interaction and line (Stone-Wales) or point (vacancy) defects randomly present on the nanotube to the critical pullout force. Results demonstrated that the pullout forces are negligibly small compared to the ones reported by experiments, even in the presence of Stone-Wales and vacancy defects. For example, the measured force of ~3.75 nN for wavy 500-nm nanotubes translates to an average IFSS of ~0.77 MPa, which is much smaller than experimentally reported IFSS values of 3–47 MPa for CNT-PMMA interfaces. This implies the presence of bonded interactions in the form of crosslinks between the CNT and PMMA. By introducing equally spaced covalent C-C bonds/crosslinks between nanotube and PMMA monomers, and by defining a distance-based bond-breaking criteria, we investigated the pullout process and the effect of crosslink density and nanotube length. We found that to initiate nanotube pullout from the PMMA matrix, the applied pullout force must locally induce the pullout of polymer chains attached to the individual crosslinks and mechanically unfold these chains. Subsequent scission of the crosslink bonds can occur once the applied pullout force is sufficiently large. This process allows for substantial relative slip between the nanotube and PMMA matrix, which is not accounted for in classical shear lag models.

We, therefore, incorporated the mechanical behavior of chain deformation (Fig. 2a) to modify the classical shear lag model, which can now take the relative interfacial slip into account. Based on this analytical model, we predicted an evenly distributed load transfer through crosslink bonds, which results in almost linear axial force distribution in the nanotube for lower crosslink densities. This is also confirmed by our MD simulation results

(Fig. 2b). Furthermore, the slip mechanism is able to capture the saturation of forces as tube embedment length increases through a transition from pure slip to slip-shear lag type of load transfer. As an outcome, the model can quantify crosslink bond distances (20nm–30nm in Fig. 2c) for previously performed nanomechanical experiments.

WHY BLUE WATERS

The Blue Waters computational capacities were necessary for several reasons. First, the number of atoms in an all-atom (with no coarse-graining) system of ~100 nm- to 500-nm-long nanotubes embedded in PMMA matrices ranges from ~0.2 to 10 million. Second, the heating-quenching process adopted to build the PMMA configurations required ~3 nanoseconds of simulation time, corresponding to 3 million timesteps for a timestep increment of 1 femtosecond. Additionally, the force-controlled pullout simulations required 0.5 million timesteps each. Such large-scale simulations cannot be performed on existing campus clusters or XSEDE resources.

PUBLICATIONS & DATA SETS

Yi, C., et al., Direct nanomechanical characterization of carbon nanotubes-titanium interfaces. *Carbon*, 132 (2018), pp. 548–555.
 Bagchi, S., A. Harpale, and H.B. Chew, Interfacial load transfer mechanisms in carbon nanotube-polymer nanocomposites. *Proceedings of the Royal Society A*, 474 (2018), DOI:10.1098/rspa.2017.0705.

DI

IMPROVING VIRTUALLY GUIDED PRODUCT CERTIFICATION WITH IMPLICIT FINITE ELEMENT ANALYSIS AT SCALE

Allocation: Director Discretionary/75 Knh

PI: Seid Korić¹

Co-PIs: Robert F. Lucas², Erman Guleryuz³

Collaborators: Rolls-Royce Corporation, Livermore Software Technology Corporation, Cray Inc.

¹University of Illinois at Urbana-Champaign

²University of Southern California

³National Center for Supercomputing Applications

EXECUTIVE SUMMARY

We continue advancing the state-of-the-art in implicit finite element analysis (FEA) by studying real-world models and identifying and removing scaling barriers. We study large-scale representative gas turbine engine models used in virtually guided certification to support physical engine testing, and for post-test validation of modeling and simulation techniques. Blue Waters is an enabling platform where massively parallel solver technology can be tested and advanced. Using it, we were the first to solve, in-core, an implicit LS-DYNA model with over 100 million finite element equations. While some of the challenges were known in the abstract, their details were only discovered by actually running at scale. Given the popularity of implicit finite element methods, the results of this effort will further expand the impact of high-fidelity multiphysics modeling of complex structures.

RESEARCH CHALLENGE

Modeling and simulation have become increasingly important in both product design and support. Model complexity and size have increased accordingly, leading to longer runtimes for simulation software. The increase in runtime is not linear, but exponential. Research and development of modeling and simulation software are needed to improve its parallel scalability so that larger, more sophisticated models can be run with faster time-to-solution, informing and shortening the design cycle. Without improvements in scalability, high-fidelity analysis will remain a postdesign or pretest analysis without being integrated into the design cycle. The effort of scaling up is motivated by three broad industrial goals: improving product design, developing high-fidelity digital twins, and allowing for virtual certification for original equipment manufacturers.

METHODS & CODES

LS-DYNA is an FEA software package used by a wide range of industries, including aerospace, automotive, biomedical, construction, and consumer products. It is used to analyze a diverse set of manufacturing problems such as the simulation of automotive collisions, explosions, manufacturing, and problems with large deformations. Commercially supported multipurpose codes are preferred by industry for product design and analysis.

These codes provide important features, such as a rich library of elements, contact capabilities, nonlinear material constitutive models, mesh adaptivity, and flexible coupling.

Our work follows a measure-analyze-improve cycle using a large-scale real-life model. It is a continuation of the implicit analysis work started on Blue Waters in 2017 and builds on an even earlier examination of explicit analysis with LS-DYNA. Our cross-disciplinary research team involves members from NCSA, Rolls-Royce, Livermore Software Technology Corporation (LSTC), and Cray. Rolls-Royce, NCSA's Industry partner, provides large-scale, whole-engine FEA models (Fig. 1) having as many as 200 million degrees of freedom. These are representative of gas turbine engine problems used in virtually guided certification for risk mitigation, for supporting physical engine testing, and for post-test validation of modeling and simulation techniques.

RESULTS & IMPACT

Multifrontal direct solvers are commonly the default solver technology in implicit FEM codes due to their generality and robustness. Their scaling has been investigated for decades. Solution of a sparse linear system by factoring is generally performed in four stages: reordering to reduce the required storage and number of operations in numeric factorization; symbolic




Figure 1: A cross-section of the jet engine model.

factorization, in which the elimination tree is created; numeric factorization; and forward elimination and backward substitution. Previous work has usually ignored the reordering and symbolic factorization stages, which were often sequential.

Scaling studies on Blue Waters show that the reordering and symbolic factorization stages are sequential bottlenecks, consuming a large and increasing fraction of the wall-clock time as the processor count grows. We have developed a new graph partitioning algorithm (LS-GPart) to replace Metis for sparse matrix reordering. Metis has been the default option for many years and was a sequential bottleneck in both time and memory. LS-GPart uses heuristics that aim to maintain the same quality regardless of the number of MPI ranks. Reordering scales well (Fig. 2a) and will do even better once OpenMP directives are added. The reordering stage has a significant impact on memory requirements and the number of floating-point operations of the factorization. The plot of the memory footprint and computation rate for factorization (Fig. 2b) shows scaling to 16,000 threads with hybrid (MPI and OpenMP) parallelization. We are continuing performance optimizations of the LS-GPart reordering and are also working toward eliminating the symbolic factorization sequential bottleneck.

We also discovered another scaling bottleneck in LS-DYNA that was unknown to its authors. Even a small Blue Waters configuration of 256 nodes exceeds that available to developers at LSTC or similar U.S. independent software vendors. By increasing to 2,048 Blue Waters nodes, the code was stressed in new ways, and a constraint processing function whose timing was heretofore insignificant suddenly emerged as a sequential bottleneck, growing with the number of processors. This too is now being addressed.

WHY BLUE WATERS

Sparse direct solver algorithms in implicit FEM analysis are both computation- and memory-bound, as well as being communication-bound at large scales. While some of these challenges were known in the abstract, their details were only discovered when we tried to scale. Processing and memory bottlenecks revealed themselves as the number of processors increased by an order of magnitude beyond that familiar to today's developers and users. Many other scientific and engineering codes intended to run on high-performance computing platforms will have similar challenges. Blue Waters is an enabling platform where massively parallel sparse solver technology can be tested and advanced by taking full advantage of large amounts of distributed memory, thousands of multicore processors, and the low latencies and increased bandwidth of advanced interconnect technologies.

PUBLICATIONS & DATA SETS

Sabet, F., O. Jin, S. Korić, and I. Jasiuk, Nonlinear micro-CT based FE modeling of trabecular bone—Sensitivity of apparent response to tissue constitutive law and bone volume fraction. *Int J Numer Meth Biomed Engng*, 34:4 (2018), DOI:10.1002/cnm.2941.

Idkaidek, A., S. Korić, and I. Jasiuk, Fracture analysis of multi-osteon cortical bone using XFEM. *Comput Mech*, (2017), pp. 1–14.

Barabash, R., et al., Finite Element Simulation and X-Ray Microdiffraction Study of Strain Partitioning in a Layered Nanocomposite. *Journal of Crystallography*, (2016), DOI:10.1155/2016/4351347.

Vazquez, M., et al., Alya: Multiphysics engineering simulation toward exascale. *J Comput Sci*, 14 (2016), pp. 15–27.

Korić, S., and A. Gupta, Sparse matrix factorization in the implicit finite element method on petascale architecture. *Comput Methods Appl Mech Eng*, 302 (2016), pp. 281–292.

Puzyrev, V., S. Korić, and S. Wilkin, Evaluation of parallel direct sparse linear solvers in electromagnetic geophysical problems. *Computers & Geosciences*, 89 (2016), pp. 79–87.

Kale, S., S. Korić, and M. Ostojca-Starzewski, Stochastic Continuum Damage Mechanics Using Spring Lattice Models. *Applied Mechanics and Materials*, 784 (2015), pp. 350–357.

Kale, S., A. Saharan, S. Korić, and M. Ostojca-Starzewski, Scaling and bounds in thermal conductivity of planar Gaussian correlated microstructures. *J. Appl. Phys.*, 117:10 (2015), DOI:10.1063/1.4914128.

Huang, W., et al., Precision Structural Engineering of Self-Rolled-up 3D Nanomembranes Guided by Transient Quasi-Static FEM Modeling. *Nano Lett.*, 14:11 (2014), pp. 6293–6297.

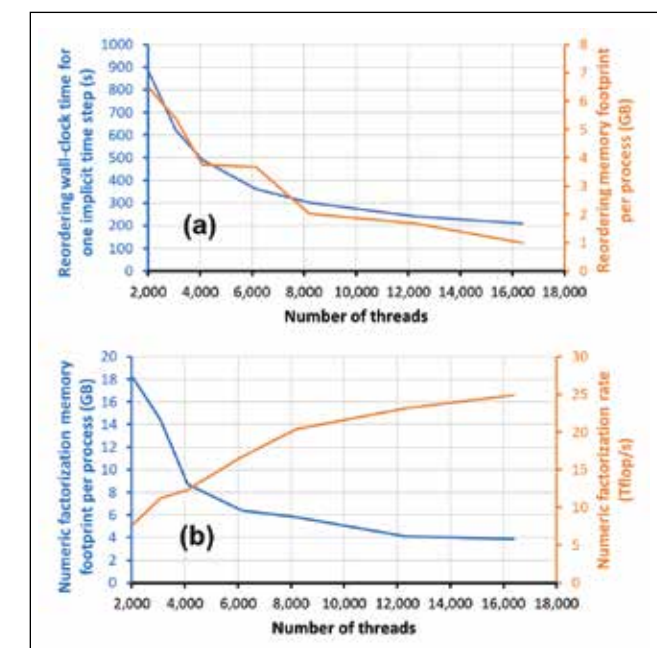


Figure 2: Scaling of sparse matrix reordering—(a) Scaling of numeric factorization, (b) A hybrid (MPI and OpenMP) parallelization model using eight threads per MPI rank. The finite element model has 105 million degrees of freedom.

RUNNING PLASCOM2 ON BLUE WATERS WITH OPENMP OFFLOADING

Allocation: Blue Waters Professor/250 Knh
PI: Daniel Bodony¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Heterogeneous systems consisting of CPU+GPU on a single node (such as Blue Waters' XK nodes) are becoming more common. Programming effectively for such systems is a difficult challenge, especially for applications that were not developed to support them from the ground up. In this context, OpenMP has recently emerged as an interesting solution, with recent versions (≥ 4.5) permitting offloading capabilities in existing code. We added OpenMP offloading to PlasCom2, a multiphysics simulation application, with the HybridOMP framework. HybridOMP enables concurrent execution of application code on CPUs and GPUs, resulting in efficient resource usage and high-performance portability. Performance results on Blue Waters show gains of 2.2x on a single XK7 node.

RESEARCH CHALLENGE

Programming heterogeneous systems is a challenging task, as few programming models support executing code on accelerator devices, leading to the use of specialized solutions such as CUDA, OpenCL, Legion, or Kokkos. Such specialized languages have the advantage of being able to provide the best performance in many cases, as they can often offer support for special device features and offer good code generation for specific device types. However, existing application code often can not be reused and must be rewritten in a new language. Code is often not portable between devices (for example, between CPUs and GPUs), such that various devices may require different implementations, leading to duplicated code and an increased difficulty of code maintenance.




Figure 1: Single-node results of PlasCom2 kernels on the Blue Waters system.




Figure 2: Scaling results of PlasCom2 on the Blue Waters system.

node. Compared to running only on the GPU, performance from heterogeneous execution was improved by 10%. Some of the computationally intensive kernels of PlasCom2 showed a speedup of 5x (Fig. 1). These gains remained similar when performing a strong scaling experiment on multiple nodes (Fig. 2). OpenMP offloading and HybridOMP show that existing codes can be enabled to run on heterogeneous systems with a low number of changes and high performance efficiency.

WHY BLUE WATERS

Blue Waters was essential to the research by providing a stable, high-performance platform with easy access to modern accelerators. HybridOMP and its integration into PlasCom2 could be developed directly on Blue Waters to evaluate and compare different implementation possibilities on a real system.

For PlasCom2, we needed an approach could use the existing code (C++ and FORTRAN) on several types of accelerators without having separate implementations for different device types. Furthermore, we wanted to be able to support different hardware and software environments, and run concurrently on both the host and offloading devices in order to use all available computing resources efficiently.

METHODS & CODES

Based on the offloading support available in recent OpenMP versions (≥ 4.5), we developed HybridOMP, which is a library to support concurrent execution on host and accelerator devices. For PlasCom2, HybridOMP measures the relative performance of the host and accelerator at startup and determines the best work distribution based on these data. During execution, HybridOMP handles the data movement between host and offloading device, as well as the actual code execution. HybridOMP requires an OpenMP 4.5 compiler and runtime but has no other dependencies.

RESULTS & IMPACT

OpenMP offloading enables simple and efficient execution of a single code base on different types of devices, with minimal changes to existing code. HybridOMP builds on top of OpenMP offloading to add support for fully heterogeneous execution; that is, running parts of the problem concurrently on different device types. Using HybridOMP with PlasCom2 resulted in a speedup of 2.2x compared to CPU-only execution on a Blue Waters XK7

QUANTUM MAGNETS AND MODELS

Allocation: Illinois/232.8 Knh

PI: Lucas Wagner¹

Collaborators: Elif Ertekin¹, Gregory MacDougall¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Effective models are critical to understanding and simulating the behavior of systems. Examples of effective models include very simple ones, such as ignoring the internal degrees of freedom within a nucleus when considering materials, to very complex ones, like device-level semiconductor modeling. Much of our thinking about materials uses some simplified effective model to make the computations easier, and many of the big successes in materials physics have been the development of high-quality effective models such as the Sommerfeld model, the Bardeen–Cooper–Schrieffer theory for superconductivity, and so on. Meanwhile, due to the continual improvement of high-performance computing resources such as Blue Waters, it is possible to perform amazingly detailed simulations of electronic behavior within materials from first principles knowing only the identity of the atoms involved. We have used these simulations to rigorously connect the detailed simulations to coarse-grained effective models for magnetism.

RESEARCH CHALLENGE

Thanks to the advance of highly parallel and available high-performance computing such as Blue Waters, it is possible to accurately compute the properties of interacting electronic systems using a quantum Monte Carlo method. This technique has been demonstrated on quantum systems of up to 1,000 electrons. While 1,000 interacting quantum particles are challenging to describe, real materials are made up of 10^{23} interacting quantum particles, and so it is imperative to consider multiscale quantum methods. However, the systematic derivation of quantum models from quantum simulations is still under development.

METHODS & CODES

We have developed a new technique that rigorously maps interacting electronic systems from high-accuracy and high-detail quantum simulations to coarse-grained models. This requires quantum Monte Carlo simulations that can compute the one- and two-particle reduced-density matrices. To perform the quantum Monte Carlo calculations, we used the QWalk package developed at the University of Illinois at Urbana-Champaign. QWalk is particularly optimized to use very little memory and to scale almost perfectly to many thousands of cores, and it can compute the necessary observables.

RESULTS & IMPACT

The goal of this work was to produce a prediction of the singlet excitation in the material MgTi_2O_4 . We have produced a true quantitative prediction in this complex correlated material, which will be checked directly by experiment. In the figure, we show the computed spins, which, when they flip to the same direction, will produce the excitation.

WHY BLUE WATERS

Blue Waters project staff helped optimize the code we used in the calculation. The Blue Waters Symposium (BWS) was an excellent research forum to help us use modern coding practices; for example, we started using Travis CI based on conversations at the BWS.

PUBLICATIONS & DATA SETS

Busemeyer, B., G. MacDougall, and L. Wagner, Singlet-triplet excitation in MgTi_2O_4 computed by quantum Monte Carlo. In preparation (2018).

Schiller, J., A. Munoz, L. Wagner, and E. Ertekin, Benchmarking superexchange parameters in the monoxides using quantum Monte Carlo. In preparation (2018).




Figure 1: The spin density in MgTi_2O_4 . Red indicates a net spin-up in a region, while blue indicates a net spin-down. The excitation predicted flips of both spin regions up. The ground state is a quantum superposition of one region up/the other region down with the exchanged configuration.

COMPUTER SCIENCE & ENGINEERING

ENGINEERING AND FINANCE

OPTIMIZATION

PARALLEL FUNCTIONALITY

SCALABILITY

VISUALIZATION

180 *Accelerating Deep Neural Network Training with Model-Aware Network Scheduling*

182 *Cinematic Scientific Data Visualization for CADENS*

184 *Algorithms for Extreme Scale Systems*

186 *Topology-Aware Distributed Graph Processing for Tightly Coupled Clusters*

187 *Deep Learning Applications in Engineering and Quantitative Finance*

188 *Holistic Monitoring and Diagnosis of Large-Scale Application Performance Degradation*

190 *RADICAL-Cybertools: A Building-Blocks Approach to Harnessing the Power of Many*

192 *Toward Robust and High-Fidelity Multiphysics Modeling on Petascale Architecture*

194 *Parallel Algorithms for Solving Large Multidimensional Assignment Problems with Decomposable Costs*

196 *Scaling Elliptic Solvers via Data Redistribution*

198 *An Efficient Optimization Algorithm for Automatic Tuning of Machine-Learning Models*

ACCELERATING DEEP NEURAL NETWORK TRAINING WITH MODEL-AWARE NETWORK SCHEDULING

Allocation: Exploratory/40 Knh
PI: Roy H Campbell¹
Co-PIs: Sayed Hadi Hashemi¹, William Gropp¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Deep Neural Networks (DNNs) form the crux of advanced solutions in a variety of domains such as computer vision and natural language processing. The increasing complexity of these problem domains necessitates large-scale, distributed training of the associated DNNs. Today, performance and scalability of distributed DNN training are bottlenecked by iterative computations among nodes. In this work, we took a more in-depth look at communications in DNN workloads.

- We proposed a performance model for DNN workloads and used the model to explain the lesser-known performance bottlenecks in these workloads.
- Our benchmark results showed that while high-performance networks such as Gemini substantially reduce the communication time, they may decrease overall performance by reducing the overlap compared to commodity networks.
- Lastly, we developed Caramel, a model-aware approach to take advantage of faster networks while achieving the highest performance.

RESEARCH CHALLENGE

In frameworks such as TensorFlow [1] or PyTorch, the computation and communication involved in training are represented using dataflow graphs, which are directed acyclic graphs (DAGs). The state of the DNN is represented by a vector of parameters. Each iteration involves the computation of parameter updates, followed by their exchange among the participating nodes.

Currently, performance and scalability of distributed DNN training is bottlenecked by this parameter aggregation [2,3]. At a high level, there are two causes of these network inefficiencies during aggregation. First, the most commonly used pattern of parameter aggregation relies on a centralized server, Parameter Server (PS), and an all-to-one/one-to-all communication pattern, causing incast/outcast at the PS and limiting the goodput (sustained performance) of the system. Second, the scheduling of operations is such that training computation blocks on the network-intensive aggregation leaves GPUs idle. The increasing sizes of training data sets and DNN models and, therefore, increasing the number of machines involved in training, is likely to reduce the overall inefficiency.

In this work, our goal was to minimize network bottlenecks in distributed DNN training to reduce the iteration time and increase GPU utilization. Toward this goal, we designed and built Caramel. The intuition behind Caramel is that while some related incast/outcast problems may be found in data analytics and graph-processing systems, DNN training offers a particular opportunity for network optimization since it has a more predictable environment with fixed parameter sizes and nearly identical iterations. Caramel applies this idea by: (1) analyzing a variety of aggregation patterns and choosing the appropriate aggregation pattern for a given environment and model, which we refer to as “network-aware optimization”; and (2) leveraging an understanding of the operations within model training to improve communication/computation overlap, which we refer to as “model-aware optimization.”

METHODS & CODES

To achieve performance improvement through network-aware optimization, we analyzed all-to-one and decentralized data aggregation techniques such as the bucket [4] and halving-doubling (HD) [5] algorithms. The bucket algorithm and HD offer better network load distribution across workers compared with PS and naive all-to-all communication. However, these




Figure 2: Performance comparison of our system (Caramel) versus TensorFlow Default (PS) versus TensorFlow and MPI (Horovod).

techniques incur the cost of multiple stages per transfer and require synchronization across workers during aggregation. With Caramel, we showed that the decentralized aggregation patterns such as the bucket algorithm and HD can work well in DNN applications when coupled with model optimizations that enable better synchronization among workers.

To understand the opportunity for network acceleration through model-aware optimization, we investigated dataflow models associated with 16 DNNs and identified common model characteristics that enable efficient network transfers. We found that the same model can result in network transfers activated (i.e., parameters being ready for aggregation) in different orders across multiple workers. This can prove detrimental to decentralized aggregation where all workers should activate the same parameter before the transfer is initiated. To solve this problem, we enforced ordering in network transfers by adding additional dependencies in the DAG to force all workers to activate network transfers in the same order.

All DNNs we analyzed have a large number of small parameters that incur significant overhead during network transfer. To tackle the small-parameter overhead, we implemented model-aware batching in Caramel, while ensuring that the batched parameters were ready at nearly the same time to avoid waiting. In addition, we identified the opportunity for increasing the window of network transfer during an iteration of DNN training without delaying the training. An iteration has two phases: forward pass and backward pass. Currently, transfers are restricted to the backward pass. We proposed techniques for extending network transfers to the forward pass in Caramel, thereby increasing the overlap of communication and computation.

We also implemented Caramel over TensorFlow and demonstrated that the iteration time can be reduced by up to 3.8 times without changing DNN training functionally, with the communication time reduced by at least two times in all networks.

RESULTS & IMPACT

We analyzed a variety of DNN models on TensorFlow with the commonly used Parameter Server implementation, TensorFlow with Horovod [6], which relies on MPI all reduce implementation of the bucket algorithm and halving-doubling, and our method.

We conducted our experiments on both the Blue Waters network as well as a public cloud (Azure) with commodity networks.

In Fig. 1, we observe that TensorFlow with Parameter Server offers high overlap. However, the communication cost of PS is very high. The Horovod implementation, on the other hand, reduces the communication cost with an efficient aggregation pattern. However, these decentralized patterns suffer from a poor overlap of communication and computation. As a result, PS with a higher communication cost achieves better GPU utilization in this instance. Ideally, we want to be in the bottom right region of this plot to achieve high utilization.

In Fig. 2, we report the iteration time and GPU utilization with 16 workers for Caramel, PS with eight servers co-located with the workers, and Horovod for five representative networks. We observe that Caramel can improve the iteration time by up to 3.84 times (in VGG-16) and GPU utilization by up to 2.46 times (in AlexNet-v2). We also find that performance gains are not equal across networks. The speedup is higher when communication time is much higher than computation time. In compute-intensive networks such as Inception-v3 and Resnet-v2-152, network optimization results in minimal improvement in iteration time and GPU utilization compared with PS.

WHY BLUE WATERS

Blue Waters’ excellent platform makes it easy to conduct large-scale exploration to find potential performance opportunities. Furthermore, the vibrant community of Blue Waters users and staff helped us to get up to speed faster using past experiences on other systems.

PUBLICATIONS & DATA SETS

Hashemi, S.H., S.A. Jyothi, and R. Campbell, On The Importance of Execution Ordering in Graph-Based Distributed Machine Learning Systems. *SysML* (2018).

Hashemi, S.H., S.A. Jyothi, and R. Campbell, Network Efficiency through Model-Awareness in Distributed Machine Learning Systems. *USENIX NSDI* (2018).




Figure 1: Overlap coefficient and communication/computation ratio for different communication methods with GPU utilization contours in the background (using Inception-v3 with eight workers). Our system (Caramel) versus TensorFlow Default (PS) versus TensorFlow and MPI (Horovod).

CINEMATIC SCIENTIFIC DATA VISUALIZATION FOR CADENS

Allocation: Illinois/275 Knh

PI: Donna Cox¹

Co-PIs: Robert Patterson¹, Kalina Borkiewicz¹, A.J. Christensen¹, Stuart Levy¹, Jeffrey Carpenter¹

Collaborators: Thomas Lucas², Melih Sener¹, John Stone¹, Robin Canup³

¹University of Illinois at Urbana-Champaign

²Thomas Lucas Productions

³Southwest Research Institute

EXECUTIVE SUMMARY

The Advanced Visualization Lab (AVL) continues work on the NSF-funded CADENS project (the Centrality of Advanced Digitally Enabled Science, ACI-1445176). AVL co-produced and rendered visualization scenes for the upcoming full-dome planetarium show, *Birth of Planet Earth*. We have used Blue Waters for processing data as well as rendering scenes in 4K monoscopic and full-dome formats.

RESEARCH CHALLENGE

Drawing on data from scientists involved in high-performance computing-based research, the AVL creates visualizations in a cinematic style intended for public outreach, through flat-screen science documentaries, films, and full-dome planetarium shows.




Figure 1: Photosynthetic energy harvesting in the purple bacterium *Rba. sphaeroides* [1]. Adapted from a visualization [2] of the structure [3] of chromatophores—photosynthetic vesicles—illustrating energy conversion steps from the capture of sunlight to the synthesis of ATP. This segment is dedicated to Klaus Schulten, as its content had been a lifetime passion for studies [1–3] in his Theoretical and Computational Biophysics Group until his passing in 2016. (To appear in *Birth of Planet Earth*.)

METHODS & CODES

Our Blue Waters visualization work depends on several externally provided packages: Houdini, a commercial visual effects software package from SideFX; yt, the data analysis and visualization package (yt-project.org) for ingesting and re-gridding some types of data; VMD, the Visual Molecular Dynamics package from John Stone of the Theoretical Biophysics group at the University of Illinois at Urbana-Champaign (www.ks.uiuc.edu); and Python with numpy and scipy for many sorts of data preprocessing from the formats provided by the scientists into formats usable by Houdini. We have developed our own software tools, including Ytini for yt–Houdini integration, and Blurend to organize the Houdini rendering workflow for the Blue Waters environment.

Visualizing energy harvesting in a photosynthetic purple bacterium [1] involved combining structural models from atomic, protein, organelle, and cell scales using multiple software tools [2]. The construction of structural models [3] was performed with VMD [2] and Mathematica, which was also used for the determination of relevant energy conversion steps [1]. Multiple copies of a static chromatophore model [3] were assembled in Houdini to emulate the interior of a low light-adapted purple bacterial cell. Dynamic elements—photons, electronic excitations, protons, quinols/quinones, ATP—were choreographed using Houdini to illustrate energy conversion processes for a lay viewer. Since the timescales for these conversion processes span almost 12 orders of magnitude (femtoseconds to milliseconds) [1], the animations deliberately represent a simplified visual narrative rather than the results of a specific simulation at one timescale. Some visual elements were rendered using Houdini, others with VMD.

We also relied on Blue Waters' capacity to create a high-quality visualization of Robin Canup's Moon-forming collision [4], as explained below.

RESULTS & IMPACT

Blue Waters enabled us to create and refine two data-driven cinematic animations, both appearing in the forthcoming full-dome planetarium show, *Birth of Planet Earth*, to be released in 2019:

- Formation of the Moon [4]—the first 24 hours after the collision that formed Earth's moon,
- Visualizing Energy Harvesting in a Photosynthetic Purple Bacterium [1–3].

Excerpts will be presented at the International Planetarium Society 2018 conference in Toulouse, France. In advance of the show's completion, a work-in-progress trailer is also available at <https://vimeo.com/277190989>, including contributions from AVL and other groups.

To suggest potential public impact for this work, we compare it with another full-dome planetarium show in the CADENS series, *Solar Superstorms*, for which AVL also relied on Blue Waters for data visualization. Since its 2015 release, it has been booked by over 70 planetaria and science museums in 15 countries, and translated into at least ten languages.

Further impact from past Blue Waters work involved two visualizations created during 2016–2017 and described in the *2017 Blue Waters Annual Report*: “First Light in the Renaissance Simulation” [5] and “Milky Way Analogue Isolated Disk Galaxy Visualization.” These were later submitted to the Supercomputing 2017 conference's Scientific Visualization Showcase where they were awarded first and second place, respectively.



In addition to full-length documentaries, two short videos using material created during this year's and last year's allocations appear in the Publications and Data Sets section.

WHY BLUE WATERS

Access to Blue Waters has allowed us to iterate quickly and meet deadlines. Over the course of a single weekend, we were able to render a Houdini scene involving a chromatophore model [3] and its surrounding environment that was made up of six separate render layers and a total of 20,870 image frames. The capability to render a large amount of images in a short period of time allowed us to make several iterations of the scene before finalizing a video to send to the IPS 2018 conference. This would not have been possible on our local cluster.

Furthermore, Blue Waters has made it possible to not sacrifice visual quality for render speed. One ~200-frame segment of our visualization of a planetary collision [4] was taking as many as 20 hours per frame to render. Blue Waters made it possible to render these images despite the long render time. Without Blue Waters, we would have had to either change many render settings and significantly decrease the render quality, change the camera position and lose the dramatic effect of having an arm of disk material pass closely overhead, or spend many days trying to come up with a different data representation of the simulation to make the render more manageable on our local cluster.

PUBLICATIONS & DATA SETS

Birth of Planet Earth, full-dome planetarium show directed by Thomas Lucas, to be distributed by Spitz, Inc. in 2019. A work-in-progress trailer appeared in June 2018: <https://vimeo.com/277190989>.

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).

Collision that Formed the Moon, mini-documentary: <https://www.youtube.com/watch?v=o2lRpiediP8>.

Grand Journey to Understand the Evolution of the Universe, mini-documentary: <https://www.youtube.com/watch?v=YvZs3IBILk>.

Figure 2: Visualizing the first 24 hours after the collision that formed Earth's moon, from a simulation [4] by Robin Canup of SWRI, as a Mars-sized body strikes the early Earth. (To appear in *Birth of Planet Earth*.)

ALGORITHMS FOR EXTREME SCALE SYSTEMS

Allocation: Blue Waters Professor/80 Knh
PI: William Gropp¹
Collaborator: Luke Olson¹
Students: Amanda Bienz¹, Paul Eller¹, Ed Karrels¹

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 this both in the increasing number of cores per processor and the use of large numbers of simpler computing elements in general purpose GPUs. The largest systems must network tens of thousands of nodes together to achieve the performance required for the most challenging computations.

Successfully using these systems requires new algorithms. In the last year, we have further improved a new communication model that better fits the performance of multicore nodes to develop new algorithms for sparse matrix-vector products and better understand the behavior of nonblocking algorithms for the Conjugate Gradient method. We also developed a simple implementation of the MPI Cartesian topology routines that significantly outperforms the available implementations and several parallel I/O libraries.

RESEARCH CHALLENGE

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. These methods depend both on high-performance matrix-vector products, which are

communication intensive, and on collective all-reduce operations, which introduce synchronizations that can limit scalability. Developing and demonstrating a more scalable version of this algorithm would immediately benefit those applications. Our approach begins with developing a performance model that captures the key aspects of the intra- and internode communication costs, and uses that model to inform the development of new algorithms. This approach has also yielded improved parallel I/O routines and a better implementation of a process placement operation that can improve the performance of applications.

METHODS & CODES

To address these challenges, we have developed several performance models that address limitations in off-node communication bandwidth, message matching costs, network contention, and the effect of system “noise.” We have developed benchmarks to test these performance models and conducted experiments with some applications. Some of the codes are open source and freely available.

RESULTS & IMPACT

Over the last year, we made progress in four areas. Amanda Bienz, working with Luke Olson and William Gropp, has looked at the communication cost of irregular point-to-point

communications in sparse matrix operations, with a focus on the operations needed for algebraic multigrid (AMG) methods. They are investigating methods for improving the postal and max-rate models to account for the main costs associated with point-to-point communication. They added node-aware parameters to the max-rate model, distinguishing the large differences in cost among intrasocket, intersocket but intranode, and internode messages. Furthermore, they added a quadratic queue search penalty to accurately model the cost of sending multiple messages and a network contention penalty to account for contention of links in the network when bytes are communicated across a large number of links. Fig. 1 shows the measured versus modeled cost of communication throughout sparse matrix operations on various levels of algebraic multigrid hierarchies. The results show the importance of including queue search times as well as contention effects in the network.

Paul Eller, working with William Gropp, has been using Blue Waters over the last year for an investigation into performance modeling of scalable Krylov solvers for structured grid problems. This includes developing code for measuring and processing parallel runtimes and network performance counters, developing a collection of kernels relevant to structured grid problems, and developing performance models with penalty terms that accurately model parallel performance at scale. Eller has run experiments to determine parameters for the performance models and performed scaling studies for the various parallel communication kernels and scalable conjugate gradient solvers. He is currently designing and running experiments with optimizations designed to improve the performance of these kernels on Blue Waters. He has also done some initial work on studies to investigate the impact of network noise on solver performance and to investigate the performance of scalable Krylov solvers within a quantum chromodynamics application. These experiments have helped to better understand Krylov solver performance at scale, to develop more accurate performance models, and to optimize the solvers to obtain better performance.

Ed Karrels, working with William Gropp, has been testing a variety of input/output access patterns and developing tools for improving input/output performance. These tools include:

- MeshIO—a tool for reading and writing N-dimensional meshes in parallel. This is being used by the XPACC project, sponsored by the U.S. Department of Energy, to accelerate job start-up and shut-down, and is being evaluated by two other science teams.
- Zlines—a tool for reading compressed text data at arbitrary line offsets efficiently. This is being used by bioinformatics researchers on genomic data.
- Zchunk—a tool for reading compressed binary data at arbitrary offsets efficiently. This is not currently being used by other projects.

Karrels also prepared a tutorial on input/output best practices for Blue Waters.




Figure 2: Results for simple halo exchange for a 3D mesh on up to 4,000 processes on Blue Waters. “Cart-8” shows the performance for the Cray MPI for an 8×8×8 process mesh; “Ncart-8” shows the performance of our replacement for the same mesh. From Using Node Information to Implement MPI Cartesian Topologies, submitted to *EuroMPI* (2018).

Finally, William Gropp has developed a new algorithm for implementing process mapping for Cartesian grids, which is needed for many applications that use structured grids. MPI provides a convenient routine for this operation, but few MPI libraries provide a good implementation of this operation. As a result, applications must either forgo the performance or use *ad hoc*, nonportable techniques to achieve a good mapping. Such tools do exist for Blue Waters, but they do not provide the right solution to the problem. Applications should be able to rely on the features in MPI and not need to use nonstandard, nonportable methods.

In addition, by using insight gained from our new performance model, we developed an alternative implementation of MPI_Cart_create that provides a significant performance benefit, as shown in Fig. 2.

WHY BLUE WATERS

Scalability research relies on the ability to run experiments at large scale and requires tens of thousands of nodes and hundreds of thousands of processes and cores. Blue Waters provides one of the few available environments where such large-scale experiments can be run. In addition, Blue Waters provides a highly capable I/O system, which we used in developing improved approaches to extreme-scale I/O.

PUBLICATIONS & DATA SETS

Bienz, A., W. Gropp, and L. Olson, Improving Performance Models for Irregular Point-to-Point Communication. Submitted to *EuroMPI* (2018).

Gropp, W., Using Node Information to Implement MPI Cartesian Topologies. Submitted to *EuroMPI* (2018).

Zchunk and Zlines are available at <https://github.com/oshkosh/bioio>.

MeshIO is available at <https://github.com/oshkosh/meshio>.

The improved implementation of MPI_Cart_create is part of baseenv and is available from wgropp@illinois.edu.




Figure 1: Measured versus modeled costs for communication during sparse matrix-vector multiplication (left) and sparse matrix-matrix multiplication (right) on each level of an algebraic multigrid hierarchy for unstructured linear elasticity. From Improving Performance Models for Irregular Point-to-Point Communication, submitted to *EuroMPI* (2018).

TOPOLOGY-AWARE DISTRIBUTED GRAPH PROCESSING FOR TIGHTLY COUPLED CLUSTERS

Allocation: Innovation and Exploration/10 Knh
PI: Indranil Gupta¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Cloud applications have burgeoned over the last few years, but they are typically written for loosely coupled clusters such as datacenters. In this research, we investigated how one can run cloud applications in tightly coupled clusters and network topologies on supercomputers. Specifically, we looked at a class of distributed machine learning systems called distributed graph processing systems, and ran them on NCSA's Blue Waters. The key to achieving performance in distributed graph processing systems is to partition the graph intelligently across the nodes of the cluster. Through this work, we have presented new partitioning techniques that are topology-aware. Our strategies partition the graph in a manner that reduces both runtime and network traffic in the supercomputer topology. Compared to previously existing work, our new Restricted Oblivious and Grid Centroid partitioning approaches produce 25%–33% improvement in runtime, along with a sizable reduction in network traffic. To help operators select the best graph partitioning technique, we organized our experimental results into a decision tree.

RESEARCH CHALLENGE

Recent years have seen a massive proliferation in the amount of graph data available, from social networks such as Twitter and Facebook, peer-to-peer networks such as Gnutella, Web graphs, and autonomous systems. These data troves provide an opportunity to glean useful insights about the nature of such interconnected systems. However, these graphs are enormous—the Facebook graph alone has over a billion vertices and a trillion edges, and the human brain network has many billions of vertices. The sheer size of the graphs, therefore, requires the use of distributed graph processing frameworks such as Pregel and PowerGraph to analyze them. An important step in the analysis is partitioning the graph across the various workers in the cluster. Our work studies how we can partition these graphs in a network topology-aware way so as to avoid sending data over long routes.

METHODS & CODES

Our first partitioning strategy (Grid Centroid) determines the nodes that should perform state computations for each vertex to minimize network communication cost with other mirrors. Our second partitioning strategy (Restricted Oblivious) uses a greedy heuristic while creating vertex mirrors to ensure they are all placed close to each other and to the node performing the computation.

RESULTS & IMPACT

We developed two partitioning strategies that consider the topology of the underlying network while partitioning an input graph across worker nodes. We also analyzed the performance of our partitioning strategies and developed a decision tree to help operators determine the appropriate strategy to use based on parameters such as the cluster size and the nature of the input graph. In certain cases, our strategies were able to reduce the graph processing time by up to 33% by reducing average vertex master-mirror distance and the mean number of hops travelled by network data. Another observation from our experiments was that certain algorithms such as PageRank send comparatively less data over the interconnect and instead benefit from frequently flushing the network buffers.

These results demonstrate how tightly coupled clusters can be more efficiently used for distributed graph processing.

WHY BLUE WATERS

Blue Waters' 3D Torus interconnect provides a topology with heterogeneous communication costs. It was thus an important part of our experiments.

DEEP LEARNING APPLICATIONS IN ENGINEERING AND QUANTITATIVE FINANCE

Allocation: Illinois/115.25 Knh
PI: Justin Sirignano¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Our research team used the Blue Waters supercomputer to support research into applications of deep learning in engineering and quantitative finance. The first application involved solving high-dimensional partial differential equations using a deep learning algorithm. The second application modeled high-frequency financial data using deep learning.

RESEARCH CHALLENGE

Deep learning has revolutionized fields such as image, text, and speech recognition. Due to this success, there is growing interest in applying deep learning to other fields in science, engineering, medicine, and finance. We used Blue Waters to develop deep learning methods and models for important applications in engineering and quantitative finance. We developed a deep learning method for solving high-dimensional partial differential equations, which have been a longstanding computational challenge. In another project, we developed a deep learning model for high-frequency financial data.

METHODS & CODES

Solving partial differential equations with deep learning. High-dimensional partial differential equations (PDEs) are important in engineering, physics, and quantitative finance. However, they are computationally challenging to solve. Finite difference methods are infeasible in higher dimensions. Therefore, we developed a deep learning algorithm to solve PDEs and implemented it on the Blue Waters supercomputer. The algorithm is mesh-free, which is crucial for solving high-dimensional PDEs. The deep learning algorithm was able to accurately solve several high-dimensional PDEs.

We tested this deep learning algorithm on several challenging PDEs. First, we solved a high-dimensional free-boundary PDE for American options (a financial derivative on a portfolio of stocks) in up to 200 dimensions. We also tested the algorithm for solving a high-dimensional Hamilton–Jacobi–Bellman PDE. Finally, we solved Burgers' equation using the deep learning algorithm, which approximated the general solution to Burgers' equation across a range of different boundary conditions and physical conditions.

Modeling high-frequency data with deep learning. We analyzed a large high-frequency dataset of electronic market quotes and transactions for U.S. equities. Our approach used a large-scale deep learning approach and revealed several interesting insights into the relationship between price formation and order flow,

which is the submission and cancellation of buy and sell orders. We found that the deep learning model is relatively stable across time and is able to provide a “universal model” across a range of different stocks. This implies that the relationship between price formation and order flow is far more stationary and universal across different stocks than previously believed.

RESULTS & IMPACT

Due to the success of deep learning in traditional computer science fields (e.g., image recognition), there is now significant interest in applying deep learning techniques in engineering, science, medicine, and finance. Our projects, supported by Blue Waters, are some of the first attempts at developing deep learning methods for engineering and finance applications.

WHY BLUE WATERS

Deep learning uses multilayer neural networks (i.e., deep neural networks) to build statistical models of data. This training of the deep learning model can be computationally intensive due to both the large number of parameters and the large amounts of data. Graphics processing units (GPUs) can be used to accelerate training of deep learning models. We leveraged Blue Waters' large amount of GPU resources to develop deep learning models for applications in engineering and finance. Blue Waters' technical staff provided invaluable help throughout the project, including solving a number of technical issues related to deep learning computational frameworks such as PyTorch and TensorFlow.

PUBLICATIONS & DATA SETS

Sirignano, J., and K. Spiliopoulos, DGM: A deep learning algorithm for solving high-dimensional partial differential equations. arXiv: 1708.07469 (2017).

Sirignano, J., and R. Cont, Universal features of price formation in financial markets: perspectives from Deep Learning. arXiv: 1803.06917 (2018).

HOLISTIC MONITORING AND DIAGNOSIS OF LARGE-SCALE APPLICATION PERFORMANCE DEGRADATION

Allocation: Exploratory/50 Knh
PI: Ravishankar Iyer¹
Co-PI: Zbigniew Kalbarczyk¹
Collaborators: Saurabh Jha¹, Benjamin Lim Wen Shih¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Significant application performance variation has been attributed to “hot-spots” in the high-performance interconnect network, or localized congestion regions arising from applications contending for the same resources within a computer system. Such attribution has largely relied on limited inferential, simulated, or theoretical data. We present a methodology, implemented as a tool, to provide congestion characterizations at runtime for systems with large-scale interconnect networks and to inform diagnostic investigations. We studied Blue Waters high-speed network congestion data to characterize and diagnose likely congestion causes in applications. Our findings include:

- Continuous presence of highly congested links in the network. From our data, in 95% of the operation time we observed congestion regions with a minimum of 20 links and a maximum of 6,904 links. The average congestion duration is 16 minutes and the 95th percentile is 16 hours.
- Limited efficacy of default congestion-mitigation mechanisms. On average, congestion mitigation mechanisms trigger every seven hours but fail to detect 61% of the high-congestion events.




Figure 1: Characterization and Diagnosis workflow for interconnection-networks. The data set size highlighted in red above each block represents the size of the data set collected from Blue Waters.

RESEARCH CHALLENGE

Significant application performance variation has been attributed to network congestion in localized “hot spots” that arise from application contention. Determination of congestion-related performance impact in large-scale HPC production interconnect networks has been largely inferential, based on messaging rates or counters measured within an instrumented application or benchmark. Such measurements may not expose the hot spots, particularly for architectures where an application may be affected by congestion in components not directly accessible from its resource assignments (e.g., the congestion may occur in network components not under the control, nor visible to, the application). Simulations and theoretical models are limited in dynamic detail and often present “steady state” results that do not enable early detection of congestion or analysis of its evolution.

Production characterizations of congestion manifesting as hot spots (as opposed to fully congested, i.e., underprovisioned, networks) can be difficult, since they require systemwide, coherent data on the state of network performance counters at each component that may be limited in the amount of information exposed that can be used for pinpointing and attributing the root cause of congestion. In addition, data must be collected at fidelities necessary to capture the relevant phenomena. Such characterizations can, however, inform designs and acquisition decisions. Runtime evaluations that identify localized areas of congestion and assess congestion duration can be used to trigger Congestion Effect Mitigating Responses (CEMRs), such as resource scheduling, placement decisions, and dynamic application reconfiguration.

In this work, we present a methodology for characterizing congestion in large-scale high-speed 3D Torus networks.

METHODS & CODES

Data Sources. We have used five months of production monitoring data from Blue Waters. We used system-generated network resiliency logs (via network logs), Light-weight Distributed Metric Service (LDMS) data, and generated network performance metrics and Blue Waters application logs. The sizes of the LDMS logs, network logs, and application logs are about 15 TB, 100 GB, and 7 GB, respectively.

Methodology and Tools. We developed *Monet*, a generic framework for supporting congestion characterization and

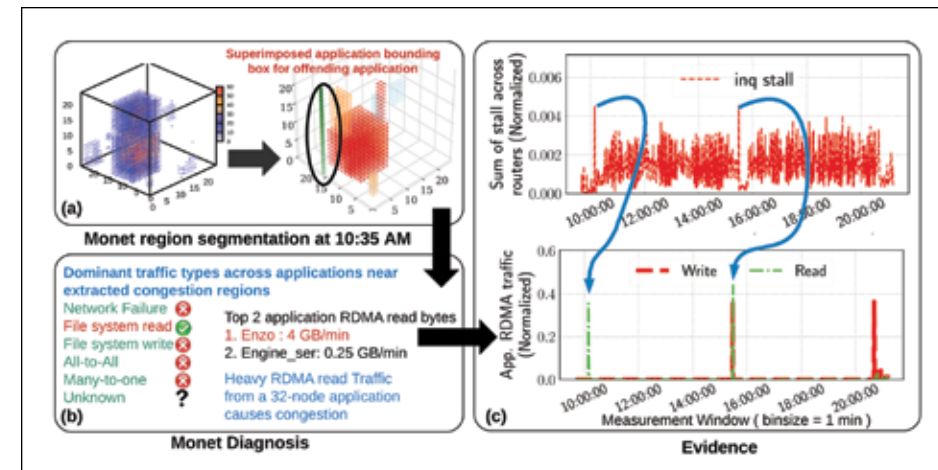


Figure 2: Detection and Diagnosis methodology applied to real scenario: (a) extracting congestion regions (CR), (b) finding anomalous application traffic pattern, and (c) generating evidence.

diagnosis in HPC systems. Fig. 1 shows the data collection and analysis pipeline of the *Monet* framework in the context of Blue Waters. *Monet* aggregates and analyzes the network and I/O, resilience, and workload data. The network stall counters are used in the extraction of congestion regions. A congestion region (CR) is a bounding (cuboid) box over all the links that (a) are close to one another (e.g., within a hop distance of three), and (b) have similar stall values. The identified congestion regions are then combined with other data sets (workload data, network failure data, and network performance data) to enable detection, diagnosis, and characterization of network congestion.

RESULTS & IMPACT

In this section, we show an example use case in which our analysis methodologies and framework, *Monet*, were used to detect and diagnose the congestion for an example scenario obtained from real data for which the ground truth of the cause was available.

Step 1: Extraction of CR. Fig. 2 shows that our analysis indicated widespread high-level congestion across the system (see the left graph in Fig. 2a). An in-depth analysis of the raw data resulted in identification/detection of congestion regions (see the top-right graph Fig. 2a).

Step 2: Congestion diagnosis. First, the tool correlates the CR-data with application-related network traffic (for all applications that overlapped or were near the congestion regions) and network information to generate candidate factors that may have led to congestion. In this example, there were no failures; hence, this analysis generated only application-related candidate factors A_{CR_i} . Next, we identify anomalous application traffic characteristics by using a median-based outlier detection algorithm.

In our example, as indicated in Fig. 2b, the offending application was Enzo running on 32 nodes allocated along the “Z” direction at location $(X, Y, Z) = (0,16,16)$ (indicated by a black circle in Fig. 2a). At the time of detection, Enzo was reading from the file system at an average rate of 4 GB/min (averaged over the past 30 minutes and with a peak rate of over 70 GB/min), which was

16 times greater than the next-highest rate of read traffic by any other application in that time-window. The tool identified RDMA read bytes/min of the Enzo application as the outlier feature. Hence, Enzo was marked as the anomalous factor that led to the congestion.

Our tool generates evidence for system managers and users by producing timeseries data and statistical distributions of stall and traffic characteristics for the implicated application. The two peaks (marked) in this top plot correspond to the increase in read bytes (normalized to total read bytes during the application run) shown in the bottom plot. Although in this example scenario the Cray congestion mitigation mechanism was triggered, it was not successful in alleviating the network congestion, in part because the congestion was due to I/O traffic rather than message traffic. Instead, the CR size grew over time, impacting several applications.

WHY BLUE WATERS

Blue Waters is one of the few open-science capacity systems that provides a testbed for scaling computations to tens or hundreds of thousands of cores on CPUs and GPUs. It also enables the study of failures and degradations of applications in production petascale systems with 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 & DATA SETS

Jha, S., et al., Holistic Measurement Driven System Assessment. *Proceedings of the Workshop on Monitoring and Analysis for High Performance Computing Systems Plus Applications* (Honolulu, Hawaii, September 2017).

Jha, S., et al., Characterizing Supercomputer Traffic Networks Through Link-Level Analysis. Submitted (2018).

Jha, S., et al., Characterizing and Diagnosing Congestion in Large-Scale Networks. Submitted (2018).

RADICAL-CYBERTOOLS: A BUILDING-BLOCKS APPROACH TO HARNESSING THE POWER OF MANY

Allocation: NSF PRAC/3,450 Knh
PI: Shantenu Jha¹

¹Rutgers, the State University of New Jersey

EXECUTIVE SUMMARY

RADICAL-Cybertools (RCT) is a set of software tools that provides extensible building-block capabilities to enable the scalable utilization of production-distributed cyberinfrastructure. We use RCT to address the computational challenges in running ensemble-based methods for free-energy protocols (HTBAC) and enhanced sampling techniques (ExTASY). We are also using RCT to provide deep-learning workflows for classification of high-resolution satellite imagery.

RESEARCH CHALLENGE

RADICAL-Cybertools has been used to develop three toolkits, which have been utilized to support scalable science applications:

- HTBAC (High-Throughput Binding Affinity Calculator)—HTBAC aims to facilitate the running of drug-screening campaigns of molecular simulations for computing binding affinities.
- ExTASY (Extensible Toolkit for Advanced Sampling and analysis)—ExTASY uses enhanced sampling techniques to explore the same conformational space using computationally efficient approaches to study more complex proteins at longer timescales.

- ICEBERG (Imagery Cyberinfrastructure and Extensible Building Blocks to Enhance Research in the Geosciences)—ICEBERG aims to facilitate pan-Arctic and pan-Antarctic science projects requiring manipulation and interpretation of high-resolution satellite imagery.

METHODS & CODES

We have developed the HTBAC toolkit to address the need for scalable and accurate protein–drug binding affinity results in acute precision medicine. In a computational drug campaign that requires the screening of thousands of compounds, accuracy comes at the cost of capturing subtle chemical properties of individual drug candidates. Accurate free-energy calculations require adaptive workflows that intervene during runtime to analyze intermediately generated simulations to optimally place the next set of simulations. In the second citation under Publications and Data Sets (below), we demonstrate that adaptive simulation methodologies can improve binding affinity accuracy given a fixed amount of compute time to make drug campaigns more computationally effective. Our code is open source and provided at <https://github.com/radical-cybertools/htbac>.




Figure 1: Adaptive simulation methodologies reduce both resource consumption and uncertainty errors in free energy, leading to a reduction in time-to-solution and improvement in accuracy. The “optimal” simulation parameters are difficult to determine *a priori*, thus requiring runtime determination and adaptation of the workflow based on intermediately generated data.

The ExTASY project aims to address the fact that the adaptive sampling of proteins is more complex to set up than plain molecular dynamics since it requires periodic iterations of molecular dynamics and analysis steps. By using a workflow system, this complex setup is simplified. The fourth citation under Publications and Data Sets (below) provides results demonstrating scalability and early adaptive execution of ExTASY on Blue Waters. ExTASY is a Python library supporting domain scientists at Rice University. Our code is open source and provided at <https://github.com/radical-collaboration/extasy-grlsd>.

The ICEBERG project aims to address the following challenge: methods of using pixel-based classification models to extract feature representations for low- and medium-resolution images provide limited use of classification for high-resolution imagery. ICEBERG will provide deep-learning workflows for classification of high-resolution satellite imagery, and its integration with other data streams (e.g., Digital Elevation Models). The code of this work may be found at <https://github.com/iceberg-project>.

RESULTS & IMPACT

The first and second citations under Publications and Data Sets (below) provide results demonstrating scalability and early adaptive execution of HTBAC on Blue Waters. In addition, we were recently awarded the 2018 IEEE SCALE Challenge for using HTBAC on Blue Waters to explore simulations that calculate the strength of drug binding in order to reduce costs and improve computational efficiency for drug design and to underpin development in personalized medicine. We were able to extend the previous version of ExTASY to use GPUs and also increased its flexibility with regard to what kind of adaptive sampling algorithms can be used. So far, studies on Blue Waters have helped to pinpoint the most critical parameters to ExTASY performance and to improve sampling techniques. In the coming months, full-scale simulations of Chignolin, using ExTASY, will be completed in attempts to show that these methods can reproduce the work of [1].

WHY BLUE WATERS

Blue Waters is essential for our computational campaign in order to reach the scales needed to make a clinical impact. HTBAC and ExTASY require a scalability to at least 2,000 nodes per use-case. The full use-case of ICEBERG requires scalability up to 1,000 GPU nodes to classify images using deep neural networks.

PUBLICATIONS & DATA SETS

Dakka, J., et al., High-throughput binding affinity calculations at extreme scales. *Computational Approaches for Cancer Workshop at Supercomputing17* (2017).

Dakka, J., et al., Concurrent and Adaptive Extreme Scale Binding Free Energy Calculations. <https://arxiv.org/pdf/1801.01174.pdf> (2018).

Dakka, J., et al., Enabling Trade-off Between Accuracy and Computational Cost: Adaptive Algorithms to Reduce Time to

Clinical Insight. *SCALE Award Winner, IEEE/ACM CCGrid'18* (2018).

Balasubramanian, V., et al., ExTASY: Scalable and flexible coupling of MD simulations and advanced sampling techniques. *Proceedings of the 2016 IEEE 12th International Conference on e-Science*, (March 3, 2017), pp. 361–370.

TOWARD ROBUST AND HIGH-FIDELITY MULTIPHYSICS MODELING ON PETASCALE ARCHITECTURE

Allocation: Exploratory/50 Knh
PI: Seid Korić¹
Co-PI: Guillaume Houzeaux²
Collaborators: Paula Córdoba², Anshul Gupta³

¹University of Illinois at Urbana-Champaign
²Barcelona Supercomputing Center
³IBM–T.J. Watson Research Center

EXECUTIVE SUMMARY

Demonstrating the exceptional level of parallel scalability of the massively parallel multiphysics finite element code Alya from the Barcelona Supercomputing Center (BSC) on the petascale architecture of Blue Waters was one of the major global high-performance computing (HPC) breakthroughs in 2015. The next step in this joint international work is to provide Alya with a fast and efficient direct solver to deal with solid mechanics simulations in which using an iterative solver together with preconditioning techniques is either too expensive or it does not converge to the desired solution. Research teams from the National Center for Supercomputing Applications, BSC, and IBM–Watson have been working jointly integrating the solver from IBM–Watson into Alya’s multiphysics code. A numerically challenging problem of the buckling of thin structures was chosen to test the implementation of solvers from the Watson Sparse Matrix Package (WSMP) into Alya.

RESEARCH CHALLENGE

Complex physical problems for both applied fields and basic research, such as fluid dynamics, heat transfer problems, solid mechanics, or general transport equations, are often represented by partial differential equations that have to be discretized and solved numerically. This takes the continuum formulations of

physics to systems of algebraic equations, and in order to obtain good approximations to the real-life solutions of such problems it is necessary to solve linear systems, such as $Ax=b$, with a large number of unknowns. The resulting matrices obtained from this discretization are very sparse (few nonzero matrix elements), ill-conditioned, and frequently are unsolvable by iterative solution methods.

METHODS & CODES

The WSMP [1] solver, developed by Dr. Anshul Gupta of IBM–Watson, is the first direct solver showing good scalability and robustness in dealing with problem sizes of millions of ill-conditioned equations on tens of thousands of processor cores on Blue Waters. This work, published by Drs. Korić and Gupta [2], was named by *HPCwire* as a top supercomputing achievement [3] at SC17.

Alya is the multiphysics code developed at the Computer Applications in Science and Engineering (CASE) department of the BSC. The code was built from the ground up to run efficiently in parallel, solve many different problems, and maintain programmatic ease. We used Blue Waters to evaluate the parallel performance of Alya for several coupled multiphysics applications such as airflow in the human body, contraction of the heart, and

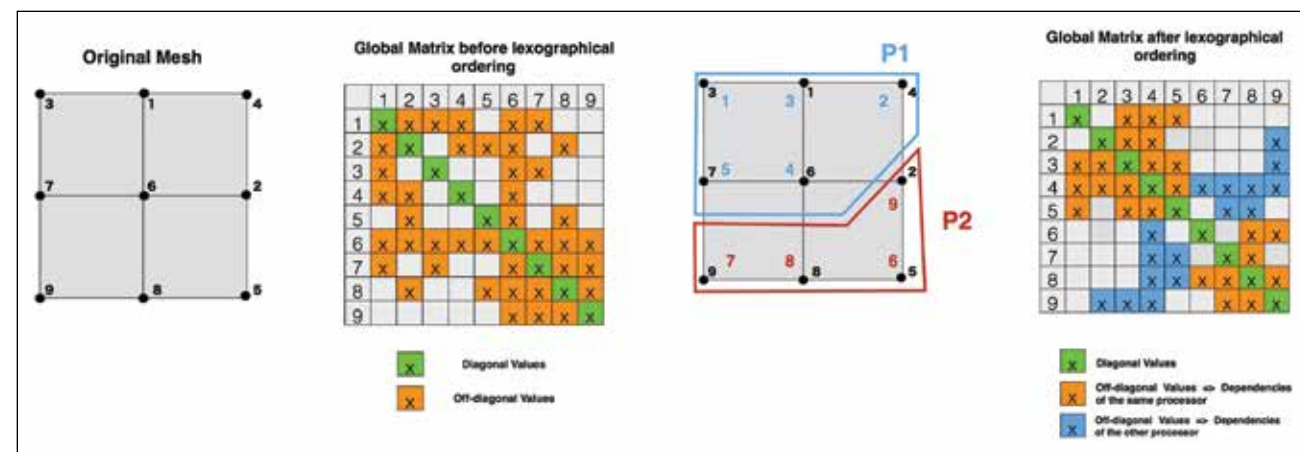


Figure 1: Global matrix structure before and after the lexicographical reordering




Figure 2: Fuselage barrel 4M mesh and results

combustion in a kiln furnace. Alya showed good scalability up to 100,000 cores [4].

Alya did not originally provide full and consecutive rows on interface nodes as required by WSMP. To achieve this, sequential processes must contain consecutive portions of the matrix, which we achieved by a lexicographical ordering. An example is shown in Fig. 1. Then, Alya’s block compressed sparse row matrix format was modified to match WSMP’s format. Finally, we integrated WSMP into Alya’s workflow for symbolic and numerical factorizations as well as solution steps

RESULTS & IMPACT

We have validated the WSMP implementation in Alya while solving a challenging large ill-conditioned structural problem of great practical importance for society inside the Alya’s Solidz module.

One of the main objectives in the structural analysis of thin structures such as fuselages has been to determine a panel’s buckling and postbuckling capacity, trying to assure the structural integrity up to the ultimate load. Postbuckling determination for thin, curved panels in aero-structures becomes a critical factor in certification and generally requires very fine meshes to model the phenomenon. This is often beyond the regular finite element analysis software and HPC hardware resources.

We conducted a buckling and postbuckling analysis of a full fuselage barrel loaded in uniaxial compression (Fig. 2). The problem has been validated within the European project SHERLOC (CS2-AIR-GAM-2014-2015-01). The analysis is highly dynamic and nonlinear. In previous work, we generated a mesh of four million 3D solid elements with a high-aspect ratio to perform

the finite element analysis using Alya [5]. The direct solver in Alya has provided a more efficient and robust solution than the iterative solvers.

WHY BLUE WATERS

Blue Waters is the only resource where massively parallel multiphysics codes such as Alya with the WSMP solver can be tested by taking full advantage of large amounts of distributed memory, hundreds of thousands of computing cores, and the low latencies and increased bandwidth of leading interconnect network technologies.

PUBLICATIONS & DATA SETS

Córdoba, P., et al., Petascale Level Multiphysics Modeling. *Workshop on Materials Computation* (University of Illinois, Urbana-Champaign, Ill., August 14–15, 2017).

Borrell, R., et al., Parallel mesh partitioning based on space filling curves. *Computers & Fluids*, 173 (2018), pp. 264–272.

Lu, Q., et al., Convergence Analyses for Fluid Structure Interaction Simulation in a Thin Hyper-Elastic Pipe. *International Conference on Computational Fluid Dynamics-ICCFD10* (Barcelona, Spain, July 9–10, 2018).

PARALLEL ALGORITHMS FOR SOLVING LARGE MULTIDIMENSIONAL ASSIGNMENT PROBLEMS WITH DECOMPOSABLE COSTS

Allocation: Illinois/210 Knh

PI: Rakesh Nagi¹

Co-PI: Ketan Date¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

The objective of this research is to develop fast and scalable GPU-based parallel algorithms that can solve large instances of the Multidimensional Assignment Problem with Decomposable Costs (MDADC). MDADC has important applications in information fusion, resource planning, and multitarget tracking domains. Typically, these problems are solved using tree search procedures, in which the lower bounds are calculated at each tree node by solving a Lagrangian dual problem, for which we propose fast and scalable GPU-accelerated algorithms. For this project, Blue Waters' computational resources are extremely beneficial due to the availability of a large number of GPU-enabled processors. These can significantly speed up the lower bound calculations at each node, allowing us to solve large problems that have not been attempted to date.

RESEARCH CHALLENGE

The axial Multidimensional Assignment Problem (MAP) is a generalization of the Linear Assignment Problem (LAP) to $K \geq 3$ dimensions, and therefore, it is nothing but a minimum-cost K -partite matching problem. To be more specific, if we have a K -partite graph with N vertices in each partition (or dimension), then the problem is to find N vertex disjoint subsets of size K (one vertex per dimension), such that the sum of the costs of the subsets is minimized. Different variants of MAP can be conceived for different strategies of defining the subset costs in the objective function [1]. The MDADC [2] is one such variant in which each edge connecting a pair of nodes from two different partitions has a certain weight, and the goal is to minimize the sum of the weights of the pairwise edges included in the subsets (basically a "clique" cost).




Figure 1: Speedup factors for MDADC instance with $N=500$ and $K=500$.

MDADC is routinely used to model the Data Association Problem [3], which is a fundamental problem in data science applications involving multiple data sources. Data gathered by these sources may be in different formats, and the goal of data association is to merge these data into a cumulative evidence aggregate that can be used for sense-making tasks or to obtain information about the current state of the real world. MDADC provides a systematic way of modeling and solving these important problems and has tremendous potential in information fusion and multitarget tracking models.

METHODS & CODES

We chose to parallelize the Lagrangian subgradient search heuristic for the MDADC formulation similar to [4], in which we need to solve $O(K^2)$ LAPs of size $N \times N$ and adjust $O(K^3 N^3)$ Lagrange multipliers to obtain strong lower bounds. We designed a parallel Lagrangian heuristic for solving MDADC using hybrid MPI+CUDA architecture. The $O(K^2)$ LAPs are split across multiple GPU nodes and solved using our GPU-accelerated Hungarian algorithm [5], while the $O(K^3 N^3)$ Lagrange multipliers are updated by multiple CUDA threads in parallel. This problem is primarily memory bound due to the presence of a large number of Lagrange multipliers. However, we found that many of these multipliers have a value of zero as a consequence of the corresponding constraint being inactive. This important observation led us to design our innovative algorithm, which identifies and stores only the nonzero multipliers and alleviates the large memory requirement.

RESULTS & IMPACT

With our method, we were able to obtain strong lower bounds on problem instances with $N=500$ and $K=500$. To solve this problem, we needed at least 64 GPUs, but the number of GPUs can be increased to achieve additional parallel speedup. We performed strong scalability studies with up to 128 GPUs. Fig. 1 shows we obtained good speedup in the initial stages. However, as we continued to increase the number of GPUs in the system, we got diminishing returns in the execution times because some of the steps require excessive MPI communication.

Our single-GPU experiments revealed that the proposed algorithm is over 200 times faster than the Map-Reduce implementation of [4] for some of the smaller problem instances. The larger MDADC instance with $N=500$ and $K=500$ contains over 3.1188×10^{10} binary variables and over 7.7657×10^{15} constraints (and corresponding Lagrange multipliers). To the best of our knowledge, problems of this magnitude have not been attempted in the literature, and we are at the frontier of this research domain.

WHY BLUE WATERS

In a typical MDADC instance, we need to solve a large number of LAPs and adjust a large number of Lagrange multipliers in order to obtain a tight lower bound. This problem is memory bound and, therefore, we need a large number of processors that can handle the LAPs and adjust the multipliers in parallel. As a

result, the GPU-accelerated Lagrangian heuristic benefits from the host of powerful GPU-enabled processors available on Blue Waters. We have used over 128 XK compute nodes, 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 & DATA SETS

Natu, S., K. Date, and R. Nagi, GPU-accelerated Lagrangian Heuristic for Multidimensional Assignment Problems with Decomposable Costs. *Parallel Computing*, submitted (2018).

SCALING ELLIPTIC SOLVERS VIA DATA REDISTRIBUTION

Allocation: Exploratory/60 Knh
PI: Luke Olson¹
Collaborators: Amanda Bienz¹, Andrew Reisner¹, Lukas Spies¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Elliptic partial differential equations play a central role in many scientific simulations. Yet, solvers for such methods are limited in efficiency, particularly at high core counts. Structured multilevel solvers are especially effective for this class of problems; however, parallelism is limited with traditional methods. In this work, data are redistributed in order to extend the scaling limits of the method.

RESEARCH CHALLENGE

Structured multilevel methods are a highly effective and common tool in solving elliptic partial differential equations across a range of applications areas. Yet, scalability of these methods is limited, particularly at high core counts. The focus of this work is on the development of a multilevel method that redistributes data on coarse levels in the solver in order to optimally extend the parallel scalability of the solver.

METHODS & CODES

Structured multilevel solvers use a series of successively coarser grids to approximate the error in the original, fine-level problem. As grids coarsen, the amount of local work is significantly reduced, leaving high levels of communication per degree of freedom in the problem. Moreover, the coarsest levels often result in only a few (or one) degree of freedom per core, necessitating a gathering of the problem to one or a subset of processors.

In this work, Blue Waters is used to test a predictive performance model [1] for scaling the solver to 100,000 cores, while the baseline solver is limited to 4,000 cores. Central to the approach is the agglomeration of data as presented in Fig. 1. Here, a method is used to:

- Agglomerate processors into groups of coarse tasks;
- Gather or collect data for processors in each block; and
- Continue cycling redundantly with redistributed data.





Figure 1: A 4 × 4 processor grid is redundantly reduced to a 2 × 2 processor grid.



There are a large number of possibilities for redistributing data. In this work a search algorithm is used to identify the optimal redistribution based on a parallel performance model. As a result, the solver is capable of scaling to very large core counts with little overhead, as shown in Fig. 2. The code used and tested for this work is the Cedar Framework (<https://github.com/cedar-framework/cedar>).

Figure 2: Weak scaling of a 3D elliptic problem on Blue Waters with a per-core size of 5M+ and a total problem size of over 660 billion points at the largest scale.

RESULTS & IMPACT

There are several outcomes of this work. The first is the testing of a new scalable solver outlined in [1]. The method uses predictive performance models to guide the redistribution of data for optimal network traffic, leading to a robust and efficient method for a range of structured elliptic problems. Another key element of this work was the development of an efficient halo exchange suitable for heterogeneous systems. Finally, the modeling aspects of this work led to new insights on parallel performance. Together, these steps have contributed to fundamental advances in robust structured solvers.

WHY BLUE WATERS

Blue Waters was instrumental in developing accurate performance models and in testing the scalability of the methods. The allocation helped to extend the method and the code to new scales.

PUBLICATIONS & DATA SETS

Reisner, A., L. Olson, and D. Moulton, Scaling Structured Multigrid to 500K+ Cores through Coarse-Grid Redistribution. *SIAM Journal on Scientific Computing*, to appear (2018).

AN EFFICIENT OPTIMIZATION ALGORITHM FOR AUTOMATIC TUNING OF MACHINE-LEARNING MODELS

Allocation: Exploratory/48 Knh
 PI: Matthew West¹
 Co-PI: Chenchao Shou¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Hyperparameters are crucial to the performance of a machine-learning algorithm. The difference between poor and good hyperparameters can mean the difference between a useless model and state-of-the-art performance. In this project, we developed an efficient optimization algorithm (Progressive Stochastic Response Surface, or ProSRS) for automatically tuning of hyperparameters. ProSRS exploits multiple cores of a machine by performing the tuning in parallel. We compared ProSRS to popular Bayesian optimization algorithms on a suite of standard benchmark functions and two real hyperparameter-tuning problems. ProSRS not only achieves significantly faster optimization convergence but is also one to three orders of magnitude cheaper in computational cost.

RESEARCH CHALLENGE

Machine learning has emerged in the past decade as one of the most exciting technologies, with widespread applications including object recognition, speech recognition, fraud detection, spam filtering, and recommender systems [1–5]. Many machine-learning algorithms have a set of tunable configuration parameters, known as hyperparameters. These hyperparameters (e.g., regularization constants, learning rates, etc.) generally have a huge impact on

the performance of a machine-learning algorithm. Indeed, the difference between poor and good hyperparameter settings can mean the difference between a useless model and state-of-the-art performance [6].

Recently, interest has grown in developing automatic procedures for tuning hyperparameters of machine-learning algorithms. Among these procedures, Bayesian optimization is a popular method. However, one issue with this method is the high computational cost, which limits the range of the hyperparameter-tuning problems that it can be applied to. Therefore, there is a need to develop a more efficient optimization algorithm for hyperparameter-tuning applications.

METHODS & CODES

We developed an efficient optimization algorithm (ProSRS) for hyperparameter tuning. Unlike Bayesian optimization that uses Gaussian processes [7], our algorithm uses radial basis functions, which are much more efficient computationally. Moreover, we developed a novel “zoom strategy” to further improve the efficiency of the algorithm.

Our codes are structured in a master-worker configuration. For each iteration, the optimization algorithm (master) proposes multiple sets of hyperparameters to a batch of workers, with each worker being assigned to exactly one set of hyperparameters.




Figure 1: Performance of optimization algorithms for hyperparameter-tuning problems. The error bar shows the standard deviation of 20 independent runs. The ProSRS algorithm is best among all the algorithms.

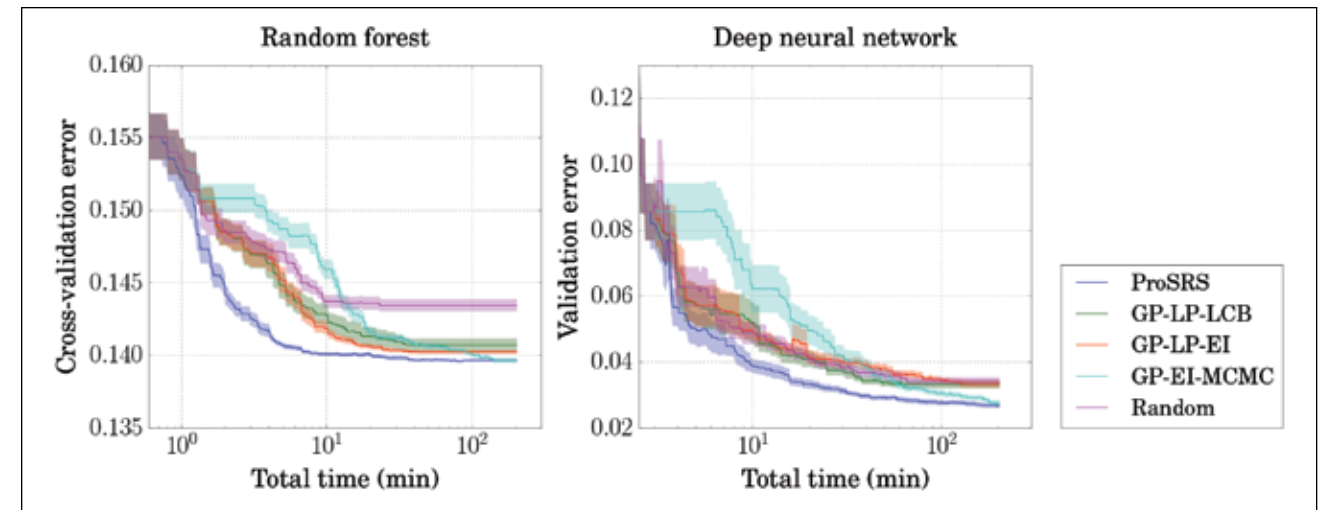


Figure 2: Optimization efficiency of different algorithms for hyperparameter-tuning problems. The total time on the horizontal axis is the actual elapsed time. Our ProSRS algorithm is significantly faster (more efficient) than the other algorithms.

A worker trains a machine-learning model with the assigned hyperparameters and returns the validation error of the model back to the master. The tasks of the workers are performed in parallel using multiple cores of a machine. Our algorithm is implemented in Python with common Python libraries (Numpy, Scipy, Scikit-learn, and Mpi4py).

RESULTS & IMPACT

We compared our algorithm to three state-of-the-art parallel Bayesian optimization algorithms: GP-EI-MCMC [8] and GP-LP [9] with LCB and EI acquisition functions. We tuned five hyperparameters of a random forest and seven hyperparameters of a deep neural network.

Fig. 1 shows the optimization performance versus iteration for different algorithms. We see that our ProSRS algorithm is better than the GP-EI-MCMC algorithm (particularly on the random forest tuning problem), and is much better than the two GP-LP algorithms and the random search algorithm. Fig. 2 shows the overall optimization efficiency that takes into account not only the optimization performance per iteration but also the cost of the algorithm and the cost of training machine-learning models. As we can see, our ProSRS algorithm is best among all the algorithms. Because of the high cost of the GP-EI-MCMC algorithm, the advantage of our algorithm over GP-EI-MCMC becomes even more pronounced compared to that of the iteration-based performance measurement (Fig. 1).

The fact that our algorithm shows superior optimization performance with significantly lower cost means that our algorithm is suitable for a wider range of hyperparameter-tuning problems, not just very expensive tuning problems.

WHY BLUE WATERS

Hyperparameter tuning requires training and evaluating machine-learning models for many iterations. To perform a full numerical experiment, we need to compound this base unit of computation with the number of tested algorithms, the number of hyperparameter-tuning problems, and the number of repeated runs for each algorithm and each problem. Therefore, this project demands large-scale computation that only Blue Waters can provide. The high-quality computing service offered by Blue Waters and the professionalism of the NCSA staff have been key to the success of the project.

BIOLOGY, CHEMISTRY & HEALTH

BIOPHYSICS

CELLULAR

HEALTH

MEDICINE

MOLECULAR

NEUROSCIENCE

PUBLIC HEALTH

- 202** *DNA “Breathing” Captured in All-Atom Detail*
- 204** *Outperforming Nature: Synthetic Enzyme Built from DNA Flips Lipids of Biological Membranes at Record Rates*
- 206** *Detection of Amino Acids with Highly Sensitive MoS₂ Nanopores: Toward Machine Learning-Based Predictive Models*
- 208** *Untangling the Origins of Protein Flexibility and Biological Functions*
- 210** *Molecular Dynamics Simulations Unveil the Mechanism of the Sensitivity to Heat of an Ion Channel*
- 212** *Using Large Ensembles of AMBER Molecular Dynamics Simulations to Reproducibly and Reliably Converge the Conformational Distributions of Nucleic Acids*
- 214** *Improving NWChem Scalability Using the DataSpaces Framework*
- 216** *Structural Basis for Extreme Cold Tolerance in the Eye Lenses of Teleost Fishes*
- 218** *Combining Physics and Supercomputers to Predict Protein Structures from Sequences*
- 220** *Algorithms for Reconstructing the Life History of Tumors from Genomics Data*
- 222** *In Silico Vaccine Design through Empirical Fitness Landscapes and Population Dynamics*
- 224** *Modeling and Simulations of Complex Dynamic Musculoskeletal Architectures*
- 226** *Constructing Large Evolutionary Trees on Supercomputers*
- 228** *Toward Predictive Computational Design of Precision Molecular Optoelectronics*
- 230** *An Efficient Hybrid Stochastic-Deterministic Simulation Technique for Living Cells*
- 232** *Identification of Missing Variants in Alzheimer’s Disease, and the New Standards for Genomic Variant Identification in Large Cohorts*
- 234** *Quantum–Classical Path Integral Simulation of Electronic Transitions*
- 236** *Blood–Artery Interaction and Simulation-Based Analysis of Aortic Aneurysm*
- 238** *Unveiling the Functions of Virus Capsids through the Computational Microscope*
- 240** *The Free Energy Landscapes Governing the Function of Complex Biomolecular Machines*
- 242** *Unraveling the Molecular Magic of Witchweed*
- 244** *Molecular Simulations on Blue Waters Aid the Understanding of How Plants Transport Sugars Between Cells*
- 246** *Data-Driven, Biologically Constrained Computational Model of the Hippocampal Network at Scale*
- 248** *Simulation of Viral Infection Propagation through Air Travel*
- 250** *HPC-Based High-Resolution Quantitative MRI*
- 252** *Automatic Knowledge Base Construction and Hypothesis Generation: Antibiotic Resistance Mechanisms for Escherichia coli*
- 254** *Molecular Mechanism of Lipid and Ion Transport in Phospholipid Scrambling*
- 256** *Structure and Function of Bacterial Chemosensory Arrays*
- 258** *Large-Scale Coarse-Grained Simulations of HIV-1: New Therapeutic Targets*
- 260** *Improving Accuracy and Scalability for Core Bioinformatics Analyses*

DNA “BREATHING” CAPTURED IN ALL-ATOM DETAIL

Allocation: Illinois/225 Knh
PI: Aleksei Aksimentiev¹
Collaborator: David Winogradoff²

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Meters of human DNA are packaged with proteins into nucleosomes and ultimately form chromosomes within the micron-diameter nucleus. This genetic material must be unpackaged and the nucleosomes disrupted in order for essential biological processes such as transcription to occur. Transcription is the process by which the information in a strand of DNA is copied into a new molecule of messenger RNA. Cell stores its genetic material in the nucleus.

Through simulations performed on the Blue Waters supercomputer, we observed in atomistic detail how DNA “breathes” by unwrapping and rewrapping an individual nucleosome. This is of particular interest in medical research

because the misregulation of nucleosome and chromosome dynamics is implicated in a number of diseases such as Coffin–Siris and Rett syndromes and can lead to the development of various types of cancer.

RESEARCH CHALLENGE

It is well-known that DNA provides the code for all living things. Evidence is now mounting that patterns in DNA sequencing as well as and in minor chemical modifications provide a structural roadmap for the global organization of genetic material within the nucleus and for determining how genes are expressed throughout the cell lifecycle. Single-molecule *in vitro* experiments [1] and bioinformatics analysis [2] have identified specific DNA motifs

and global features as important determinants of nucleosome stability. Through simulations performed on the Blue Waters supercomputer, we have observed, in atomistic detail, spontaneous and reversible unraveling of nucleosome particles, also known as “DNA breathing.” Reaching the timescale and statistics needed to characterize this process with confidence presented a significant challenge.

Arrays of nucleosomes form chromatin fibers, which ultimately organize into chromosomes. Our project is important for the understanding of how fundamental biological processes occur, and it is medically relevant as well. The improper occurrence of nucleosome and chromatin unraveling in the cell can lead to several human diseases, including various forms of cancer, Coffin–Siris and Rett syndromes, and alpha thalassemia [3].

METHODS & CODES

We performed explicit-solvent all-atom molecular dynamics (MD) simulations with the latest version of NAMD2 of a set of individual nucleosomes in varying ionic conditions, DNA sequences, and DNA length surrounding the protein core [4,5]. We then confirmed a stepwise mechanism of nucleosomal DNA detachment by performing an additional set of simulations, shifting several identified protein residues away from the DNA and observing enhanced unwrapping.

RESULTS & IMPACT

We observed spontaneous and reversible detachment of the outer stretches of nucleosomal DNA in a set of all-atom molecular dynamics simulations performed at a condition of high ionic strength. The likelihood of observing such DNA breathing events was found to correlate with the coarse-grained (CG) content of the nucleosomal DNA, with higher CG content being associated with more stable nucleosomes. In contrast, the inner stretches of nucleosomal DNA were found to be more stably associated with the histone core by a greater abundance of nonspecific DNA–protein contacts. Analysis of the simulation trajectories revealed the stepwise character of the DNA detachment process orchestrated by the motion of several conserved histone residues.

The sensitivity of unwrapping to nucleosomal DNA CG content we observed may be enhanced when forces or torques are applied to the DNA endpoints *in vivo*. Furthermore, the mutation or chemical modification of specific highly conserved histone residues identified in this study as forming important interactions with DNA could increase the rate of unwrapping. Along with bioinformatics [2] and experimental [6] studies, our results support the possibility that AT (adenine–thymine)-rich segments of DNA form less stable nucleosomes and may signal the start of transcription

WHY BLUE WATERS

Explicit-solvent all-atom MD simulation is needed to examine the fine details of DNA–histone interactions and to accurately characterize the surrounding ionic environment. Because of the

long time scale needed to observe spontaneous detachment, such MD simulations are computationally demanding. The large number of XK nodes on Blue Waters, with graphics processing unit accelerators connected by the fast Gemini interconnect, makes it one of the best publicly available systems for performing simulations studying DNA–protein and DNA–DNA interactions in atomistic detail. Over the past several years, our group has used Blue Waters to carry out a set of landmark simulations in the area of nucleosome and DNA dynamics, bringing high-performance simulations to the forefront of this research field.




Figure 1: A nucleosome particle before and after undergoing spontaneous unwrapping. White and purple highlight the histone protein core and DNA, respectively. DNA’s initial conformation is semitransparent; its final one is opaque.

OUTPERFORMING NATURE: SYNTHETIC ENZYME BUILT FROM DNA FLIPS LIPIDS OF BIOLOGICAL MEMBRANES AT RECORD RATES

Allocation: Blue Waters Professor/240 Knh

PI: Aleksei Aksimentiev¹

Collaborators: Alexander Ohmann², Chen-Yu Li¹, Christopher Maffeo¹, Kareem Al Nahas², Kevin Baumann², Kerstin Göpfrich², Jejoong Yoo¹, and Ulrich F. Keyser²

¹University of Illinois at Urbana-Champaign

²University of Cambridge

EXECUTIVE SUMMARY

Mimicking enzyme function and increasing the performance of naturally evolved proteins is one of the most challenging and intriguing aims of nanoscience [1,2]. Here, we employ DNA nanotechnology to design a synthetic enzyme that substantially outperforms its biological archetypes. Consisting of only eight strands, our DNA nanostructure spontaneously inserts into biological membranes by forming a toroidal pore that connects the membrane's inner and outer leaflets. The membrane insertion catalyzes spontaneous transport of lipid molecules between the bilayer leaflets, rapidly equilibrating the lipid composition. Through a combination of microscopic simulations and fluorescence measurements we found the lipid transport rate catalyzed by the DNA nanostructure to exceed 10^7 molecules per second—three orders of magnitude higher than the rate of lipid transport catalyzed by biological enzymes. Furthermore, we showed that our DNA-based enzyme can control the composition of human cell membranes, which opens new avenues for applications of membrane-interacting DNA systems in medicine.

RESEARCH CHALLENGE

The development of customizable synthetic enzymes will have significant impacts on the fields of biology and medicine but is challenging because the function of an enzyme depends sensitively on its atomic-scale structure and dynamics. This sensitivity to the atomic scale makes all-atom molecular dynamics (MD) the computational method of choice for prototyping designer enzymes. In this work, we simulated a synthetic enzyme made from DNA and demonstrated that it substantially outperforms its biological equivalents.

METHODS & CODES

We used the latest version of NAMD [3,4] to perform explicit-solvent all-atom MD simulations of a synthetic scramblase—an enzyme that transports lipids from one leaflet of a bilayer to the other—that was made from eight DNA strands and was embedded in a lipid bilayer membrane through two covalently attached cholesterol anchors. Microsecond-timescale simulations

allowed for the direct observation of lipid-scrambling activity. Our simulations were complemented by the experimental work of our collaborators in the Keyser Lab in Cambridge, UK.

RESULTS & IMPACT

Through all-atom MD simulations, we have shown that a membrane-spanning single DNA helix decorated with chemical tags forms a toroidal pore that provides a pathway for lipid molecules to cross from one leaflet of the bilayer to the other leaflet. We found a very small ($\sim 2 k_B T$) barrier to lipid crossing at the toroidal pore, indicating that lipids could move at near diffusion-limited rates. The average lipid transfer rate calculated from simulation exceeded 10^7 molecules per second and was found to be in good agreement with experiment. The simulations provided a microscopic description of the mechanism of lipid transport, while experiments demonstrated that the enzyme was also active in cancer cells. Results of this study have been published in *Nature Communications*.

This work exemplifies the potential of DNA nanotechnology for creating synthetic enzymes with biological functionality. The specific outcomes of this work may have direct biomedical applications. As rapid scrambling of a cell membrane's composition can activate programmed cell death, DNA nanostructures decorated with cell surface recognition factors can be developed to insert and scramble the membranes of only the target cells, opening new avenues for development of cancer therapeutics.

WHY BLUE WATERS

Explicit-solvent all-atom MD simulation is the only computational method that can treat objects built with DNA nanotechnology that are enhanced by nonstandard functional groups. It is the only method that can accurately characterize their structural fluctuations and transport properties [5]. Because of the size of the DNA structures, such MD simulations are computationally demanding. The large number of GPU-accelerated nodes and fast Gemini interconnect of Blue Waters make it one of the best publicly-available systems for performing DNA nanotechnology simulations. Over the past several years,




Figure 1: Lipid scrambling by a designer DNA enzyme. The DNA nanostructure (purple) is inserted via hydrophobic cholesterol anchors (yellow) into a lipid bilayer membrane (cut away for clarity). A toroidal pore that connects the upper (blue) and lower (red) leaflets of the bilayer allows lipids to pass between the leaflets.

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 [6–8].

PUBLICATIONS & DATA SETS

[1] Ohmann, A., et al., Outperforming nature: synthetic enzyme built from DNA flips lipids of biological membranes at record rates. *Nature Communications*, 9 (2018), DOI:10.1038/s41467-018-04821-5.

DETECTION OF AMINO ACIDS WITH HIGHLY SENSITIVE MoS_2 NANOPORES: TOWARD MACHINE LEARNING-BASED PREDICTIVE MODELS

Allocation: Illinois/135 Knh

PI: Narayana Aluru¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Many human diseases such as cancer, diabetes, and digestive disorders can be caused by malfunctioning ribosomes that synthesize defective proteins. In this regard, protein sequencing, or the precise identification of the exact sequence of amino acids that comprise a protein, promises to enable breakthrough advances in the early diagnosis of such diseases. Using the Blue Waters supercomputer, we performed extensive molecular dynamics simulations and obtained 66 microseconds (μs) of atomic trajectories that demonstrate that a single-layer molybdenum disulfide (MoS_2) nanopore can be used to detect and accurately identify individual amino acids in a polypeptide chain. With the aid of machine-learning techniques, we identified key features and clustered both the ionic current and the time each of the 20 standard amino acids spends in the nanopore and identified unique fingerprints of the signals. Using advanced machine-learning classification techniques, we predicted the amino acid type of over 2.8 million hypothetical sensors.

RESEARCH CHALLENGE

DNA sequencing using nanopore technology has significantly evolved over the last few years. Oxford Nanopore Technologies, for example, is currently fabricating a USB device called MinION that can sequence DNA in a matter of just a few hours. Both biological and synthetic nanopores have already been employed for label-free, high-resolution sequencing of DNA. Sequencing of proteins is another active area of research that promises to bring even more advances to personalized human healthcare. The main challenges associated with identification of biological molecules using nanopores are the low signal-to-noise ratio, pore degradation, unique identification of individual molecular units in real time, and the high speed of molecular movement through a nanopore [1,2]. Designing biological and synthetic nanopores with predefined properties for molecular transport is one of the most challenging problems in biotechnology. In this study, we showed that a nanopore in an ultrathin MoS_2 membrane can be used to detect and identify all of the 20 standard amino acids in the translocating proteins.

METHODS & CODES

We performed molecular dynamics simulations using the Large-scale Atomic/Molecular Massively Parallel Simulator (LAMMPS), an open-source classical molecular dynamics code for simulation of matter in liquid, solid, and gas phases. In our simulations we employed three different interatomic potentials: Tersoff, Lennard-Jones, and long-range Coulombic potential. Each simulation box comprised of about 32,000 atoms contained a monolayer membrane of MoS_2 , an amino acid chain, water molecules, and ions. The amino acid chain was pulled through a nanopore using an external force. Fig. 1 shows a proline chain translocating through the nanopore in the MoS_2 membrane.

RESULTS & IMPACT

In this study, we showed that a nanopore in a single-layer membrane of MoS_2 can be used to detect and accurately identify individual amino acids in a polypeptide chain. Using extensive molecular dynamics simulations with a total simulation time of 66 μs , we identified key features and clustered the ionic current and the time each amino acid spends in the nanopore. The amino acids were then clustered into different groups based on their physical properties (e.g., size, polarity, and hydrophobicity). Using machine

learning, we classified amino acids for any future ionic current and residence time readouts. We found that Logistic Regression, Nearest Neighbor, and Random Forest machine learning classifiers result in predictions of amino acid types with accuracies of 72.45%, 94.55%, and 99.6%, respectively.

Protein sequencing is an active area of research that promises to enable the early detection of cancer and other diseases. In fact, proteomic fingerprinting is believed to be as crucial for determining the health status of a human body as genomic sequencing. The proposed high-precision, single-base resolution and fast biomolecular sequencing using nanopore technology can lead to the fabrication of inexpensive personal healthcare devices that will help provide targeted healthcare. This will enable the quickly emerging medical fields of predictive and personalized medicine and will mark a significant leap forward for clinical genomics and proteomics.

WHY BLUE WATERS

We performed 4,103 molecular dynamics (MD) simulations of systems of up to 50,000 atoms and obtained 66 μs of molecular trajectories. Such expensive computations would not have been possible without a petascale supercomputer such as Blue Waters. LAMMPS, the MD package we use in our simulations, scales almost linearly with the number of cores on Blue Waters.

PUBLICATIONS & DATA SETS

Barati Farimani, A., M. Heiranian, and N.R. Aluru, Identification of Amino Acids with Sensitive Nanoporous MoS_2 : Towards Machine Learning-Based Prediction. *npj | 2D Materials and Applications*, DOI:10.1038/s41699-018-0060-8.




Figure 1: A snapshot of proline polypeptide translocation through the MoS_2 nanopore.

UNTANGLING THE ORIGINS OF PROTEIN FLEXIBILITY AND BIOLOGICAL FUNCTIONS

Allocation: Illinois/200 Knh

PI: Gustavo Caetano-Anollés¹

Co-PIs: Frauke Graeter², Edmond Lau³

Collaborators: Fizza Mughal¹

¹University of Illinois at Urbana-Champaign

²Heidelberg Institute for Theoretical Studies

³Lawrence Livermore National Laboratory


EXECUTIVE SUMMARY

Protein loops are flexible elements of a macromolecular structure that are responsible for the diverse repertoire of biological functions of a proteome, or the complete assortment of proteins expressed by an organism. The flexibility of protein loops is an evolutionarily conserved molecular property that is critical for the dynamics of proteins. Employing dynamic networks to summarize the atomic trajectories of loops of metabolic enzymes obtained from 300-nanosecond-long molecular dynamics (MD) simulations, we study how network structure changes with molecular functions and protein history. We first scanned ~2,100 representative proteomes with Hidden Markov Model profiles of domain structure. We then used this census to generate evolutionary timelines of domains with which to annotate the networks. Finally, we analyzed the structure of the networks and their associated functions. We uncovered both robustness and strong biases in the dynamic trajectories of our simulations over billions of years of deep evolutionary time.

RESEARCH CHALLENGE

Protein loops are promising candidates for uncovering the evolutionary relationship between protein function and dynamics, an important topic that is still poorly understood. Loops are irregular secondary structures that account for the bulk of the molecular flexibility of the three-dimensional structure of proteins. They can be considered critical parts of the dynamic personality of proteins [1]. Protein dynamics is intricately related to protein structure. It has been hypothesized that dynamics “preexists” and shapes the evolution of proteins as they adapt to carry out specific sets of motions [2]. Additionally, flexibility has been found to be conserved in protein evolution [3]. Our study tries to make the link between dynamics and evolution by: (1) exploring whether form indeed follows function and (2) uncovering the evolutionary drivers responsible for shaping the dynamics of proteins. In our previous studies, we looked at this problem from a biophysical standpoint. Here, we uniquely coupled the nanosecond dynamics of molecules to a historical study of the function–dynamics relationship of proteins spanning billions of years of evolution.

Figure 1: The diameter (y-axis) of dynamic networks across the evolutionary timeline (x-axis) of structural domains spanning 3.8 billion years of evolution. Evolutionary age ranges from the origin of proteins (0) to the present (1). The mean of distributions suggests network size is maintained during protein evolution.



METHODS & CODES

Molecular dynamics simulations. We studied 116 candidate loops that belong to the protein domains present in meta-consensus enzymes [4]. Our data set represents the entire set of seven broad categories of functional classification of structural domains. The MD simulations employed an isobaric–isothermal ensemble, TIP3P water model, harmonic restraints of 2.1 kcal/mol Å² applied to the bracing secondary structure of the protein, and a 100-mmol concentration of sodium and chloride ions. We performed the 50–70 nanosecond (ns) production runs preceded by 1 ns minimization runs using NAMD with the CHARMM36 force field.

Comparative genomic methods. We used the RefSeq database [5] to shortlist proteomes belonging to archaea, bacteria, eukarya, and viruses on the basis of the following two criteria: (1) organisms were classified as part of either the “representative” or “reference” RefSeq categories; and (2) the genome assembly was referred to as either “complete” or “chromosome.”

Unclassified and misclassified organisms were removed from the resulting data set. In addition, we excluded organisms that have an exclusively “obligate” lifestyle (such as endosymbionts or phytoplasm). Such organisms tend to have small genomes and thus possess a limited set of protein domains, which distort phylogenetic relationships [6]. We applied this criterion to organisms from the three superkingdoms but not to viruses in order to have representation of viral domains in our data set. In cases when multiple subspecies or strains were present in the data set, we chose one organism only.

We scanned the resulting ~2,100 proteomes with protein domain HMM profiles using HMMER [7]. The results from HMMER scans are necessary to place loop-domain distribution across the superkingdoms in the historical context by mapping the loop classifications [8] against phylogenomic timelines developed in our lab [9]. We also studied the features of these protein domains in conjunction with loop dynamics. For this purpose, we used a nonredundant set of ~14,000 representative proteins [10]. The domain features were calculated in the form of protein blocks (conformational prototypes) [11], while the loop features were determined by the extent of its dynamicity, i.e., whether the loops are “static,” “slow,” or “fast” [12].

Data mining. To analyze the MD simulations of protein loops in single-domain meta-consensus enzymes, we generated dynamic networks of positive and negative correlations of motions based on these simulations, which we term “dynamic networks,” and calculated important network metrics that measure cohesion and centralities. We then constructed a dynamics morphospace based on network metrics as well as principal component analyses and structural properties such as radius of gyration and Root Mean Square Deviation values.




Figure 2: Distribution of clustering coefficients (cc, y-axis) of dynamic networks indexed according to functional categories of superfamilies (x-axis). The mean of the clustering coefficient belonging to the “information” and “not annotated” categories is significantly higher than that of the other categories.

RESULTS & IMPACT

A global study of the dynamic networks observed in our MD simulations suggests remarkable leads for more detailed exploration. The average values of the diameter and length of the dynamic networks along the evolutionary timeline of structural domains remained constant throughout the entire 3.8 billion years of evolution (Fig. 1). This finding strongly suggests that dynamics is an entrenched physical property of proteins. When studying the clustering coefficient of dynamic networks that were annotated according to Vogel’s functional classification of structural domains [13], the mean of the clustering coefficient distribution of structural domains belonging to the “information” and “not annotated” general functional categories was significantly higher than those of the rest of the categories (Fig. 2). The “information” category encompasses domains that play a role in the upkeep and storage of the genetic material. Additionally, they are involved in information flow (transcription and translation) as well as replication and repair processes of the nucleic acids. Higher clustering coefficients are believed to correspond to higher levels of modularity in network structure. That, in turn, suggests informational processes in biology are compartmentalizing molecular dynamics into distinct but cohesive behaviors. Our analyses suggest that robustness in the structure of dynamic networks is tempered by the emergence of modules of dynamic behavior in specific functions of the molecular systems.

WHY BLUE WATERS

Blue Waters enabled the completion of MD simulations of loop behavior in 300 molecules and the scanning of a vast number of proteomes with advanced HMMs of structural recognition. Without access to Blue Waters, this computationally intensive study would not have been possible to achieve in a reasonable timeframe. We commend the Blue Waters support staff. They have been extremely helpful with prompt resolution of computational and other logistical issues during the execution of this project.

MOLECULAR DYNAMICS SIMULATIONS UNVEIL THE MECHANISM OF THE SENSITIVITY TO HEAT OF AN ION CHANNEL

Allocation: NSF PRAC/6,510 Knh

PI: Vincenzo Carnevale¹

Co-PIs: Michael Klein¹, Giacomo Fiorin¹

Collaborators: Tibor Rohacs²

¹Temple University

²Rutgers New Jersey Medical School

EXECUTIVE SUMMARY

The structure of TRPV1, the ion channel that makes peripheral nerves sensitive to heat and capsaicin, a bioactive component of chili peppers, has been experimentally determined in both the closed and open states. Very little is known about its activation mechanism however. Our research team has discovered a molecular mechanism of activation that involves the rotation of a conserved asparagine in one of the pore lining helices in and out of the pore. This rotation is correlated with the dehydration of four peripheral cavities. In light of this mechanism, we performed bioinformatics analyses of other evolutionarily related ion channels, analyzed newly available structures, and reexamined previously reported water accessibility and mutagenesis experiments. Overall, we have provided several independent lines of evidence that support the newly discovered mechanism.

RESEARCH CHALLENGE

Chronic pain is a diffuse medical condition that affects millions of Americans [1] and is associated with a range of diseases and disorders including diabetic neuropathy, peripheral neuropathy, low back pain, post-therapeutic neuralgia, fibromyalgia, neurological disorders, and arthritis [2,3]. Chronic pain can significantly impact quality of life [4], health, and productivity, with more than \$100 billion lost annually in the United States alone [1,5]. In many instances however chronic pain is ineffectively managed [1,2].

Further, currently used therapeutics can have significant side effects including addiction (opioids); limited pain relief (NSAIDs); cognitive impairment or sleepiness (sedatives, antidepressants, anticonvulsants, muscle relaxants); acclimation (opioids); tissue damage (NSAIDs); and gastrointestinal problems (NSAIDs, opioids); as well as limited compliance (capsaicin) [2,5–8]. As a result, a number of efforts are underway in academia, government research institutions, and the pharmaceutical industry to address the unmet medical need of controlling chronic pain.

The transient receptor potential cation vanilloid type 1 (TRPV1) ion channel transduces noxious stimuli into electrical signals and is present in peripheral sensory neurons. Inflammation causes overexpression and sensitization of TRPV1, resulting in increased responsiveness to painful stimuli (allodynia). Modulating the channel activity for pain control is safer for TRPV1 than for

other ion channels involved in propagating painful signals. This is because, despite being relatively widely expressed, TRPV1 channels do not play a crucial role in the heart or central nervous system [9]. This makes TRPV1 an extremely promising target for treating chronic pain [2, 9–11].

METHODS & CODES

The structure of the TRPV1 capsaicin-bound (CAP-bound) state was taken from the Protein Data Bank (PDB): the PDB code is 3j5r [12]. We refined the structure and modeled the missing residues using Rosetta software [13]. Four capsaicin molecules were docked following the protocol described in [14]. The protein with the ligands was embedded in a hydrated 1-palmitoyl-2-oleoylphosphatidylcholine (POPC) bilayer and surrounded by 150 mM of NaCl solution. The overall size of the system was approximately $170 \times 170 \times 160 \text{ \AA}^3$; the total number of atoms was approximately 400,000. We generated two MD trajectories with the peripheral cavities (PCs) either empty or hydrated. We used the CHARMM36 force field [15] to describe the protein and the POPC lipids. For capsaicin, we used the parameters derived in [16]. The TIP3P model was used to describe water [17]. An analogous setup was used to simulate the TRPV1 apo state (PDB code 3j5p [18]). We performed the equilibration of the systems (three in total: the CAP-bound state with empty and hydrated PCs, and the apo state) using NAMD 2.10 software [19] in several steps. Simulations were performed at constant temperature and pressure (1 atm) using the Langevin piston approach. For the vdW interactions, we used a cutoff of 11 \AA with a switching function between 8 and 11 \AA . We calculated the long-range component of electrostatic interactions using the Particle Mesh Ewald approach [20] with a cutoff of 11 \AA for the short-range component. The equations of motion were integrated using a multiple timestep algorithm, with a timestep of 2 femtoseconds (fs) 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 [21], which relate the error to the width, height, and deposition rate of the hills. Metadynamics simulations

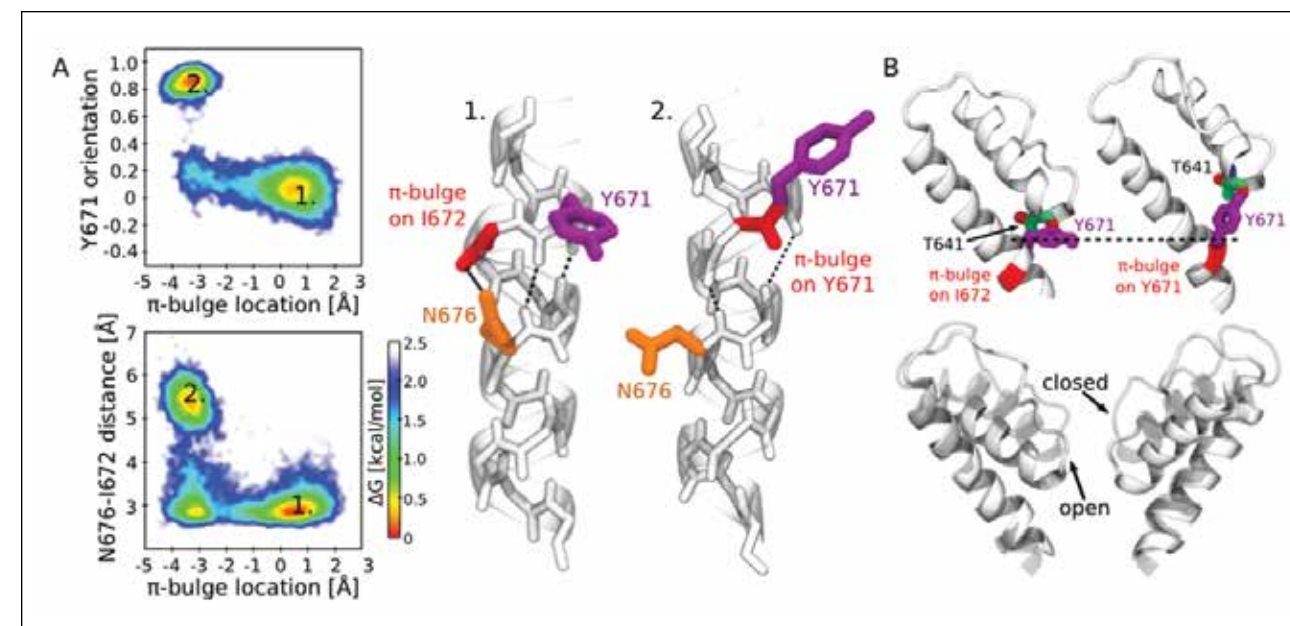


Figure 1: (A) Coupling between N676 and Y671. Top left panel: Y671 orientation and the π -bulge position. Only two conformations are observed: open state (1) and closed state (2). The bottom left panel shows the distance between N676 carboxamide carbon and I672 carbonyl oxygen and the π -bulge position. (B) A change of Y671 orientation results in a displacement of the pore helix.

were performed using the collective variable module implemented in NAMD 2.10 [22] at three temperatures: 300K, 280K, and 340K.

RESULTS & IMPACT

We found that TRPV1 activation involves the rotation of an evolutionarily conserved amino acid located in the middle of the pore lining helix S6 (N676), which projects its side chain toward either the pore or the S4–S5 domain (Fig. 1). Hydration of the pore and thus conduction of ions is only possible in the former case. We found that the conformational switch of N676 is correlated with a wet-to-dry transition in four so-far-unreported peripheral cavities. The presence of the peripheral cavities reconciled seemingly contradictory observations about the accessibility of S6 residues to solvent. We tested our model by introducing mutations in residues lining these cavities in S6 [24,25]. We were thus able to confirm our microscopic mechanism of activation [25].

Given their crucial relevance for activation, these peripheral cavities offer a completely new opportunity to modulate the activation of TRPV1 through small molecule binding. The observation that several drugs bind to the corresponding region in structurally homologous channels such as TRPML, TRPA1, and NaChBac [26–28] lends confidence to the possible druggability of these pockets.

Many attempts at developing TRPV1 antagonists have failed due to side effects and lack of efficacy. The six most advanced candidates in clinical trials have only completed Phase II since 2011; some of these have not progressed despite the passage of considerable time, likely because their development could not take advantage of the newly available structural information. By

providing a microscopic picture of the activation mechanism, our study will allow researchers to expand the chemical space and select optimal compounds.

WHY BLUE WATERS

We investigated a system of approximately 400,000 atoms using molecular dynamics. The time scales involved in the activation process of TRPV1 dictated trajectory lengths on the microsecond range. This extensive simulation was possible thanks to a massively parallel calculation enabled by the computational capabilities of Blue Waters.

USING LARGE ENSEMBLES OF AMBER MOLECULAR DYNAMICS SIMULATIONS TO REPRODUCIBLY AND RELIABLY CONVERGE THE CONFORMATIONAL DISTRIBUTIONS OF NUCLEIC ACIDS

Allocation: NSF PRAC/12,000 Knh

PI: Thomas Cheatham¹

Co-PIs: Adrian Roitberg², Carlos Simmerling³, David Case⁴

Collaborators: Darrin York⁴, Rodrigo Galindo Murillo¹, Daniel Roe⁵, Christina Bergonzo⁶, Niel Henriksen⁷, Hamed Hayatshahi⁸

¹University of Utah

²University of Florida

³Stony Brook University

⁴Rutgers University

⁵National Institutes of Health

⁶National Institute of Standards and Technology

⁷Atomwise, Inc.

⁸University of North Texas Health Science Center

EXECUTIVE SUMMARY

The AMBER simulation codes and force fields allow reliable and accurate modeling of the atomistic structure and dynamics of biomolecular systems. The molecular dynamics codes within AMBER have been highly optimized on GPUs. Taking advantage of the GPUs and using large ensembles of independent but coupled molecular dynamics simulations, we have demonstrated the ability to reliably and reproducibly converge and sample the conformational ensembles of biomolecules including DNA helices and RNA dinucleotides, tetranucleotides, and tetraloops. More efficient or enhanced sampling is enabled by applying multidimensional replica exchange methodologies. The ability to fully sample or converge the conformational ensemble allows detailed validation and assessment of enhanced sampling approaches and of the biomolecular force fields, providing detailed insight into biomolecular structure, dynamics, interactions, and function. However, even more computational power is required to continue this work and push toward larger and more complex biomolecular assemblies on longer timescales.

RESEARCH CHALLENGE

Molecular dynamics simulations are a widely applied methodology for elucidating the structure, dynamics, and interactions among biomolecules. This can help us understand biomolecular function, biomolecular assemblies, and also to probe ligand–receptor interactions and biomolecular folding. Key issues include being able to sample conformational space efficiently and effectively and also to accurately model the molecular interactions with appropriate biomolecular force fields. Using the Blue Waters petascale supercomputer, we have demonstrated the ability to overcome sampling limitations to fully elucidate the conformational ensemble of DNA duplexes and RNA dinucleotides, tetranucleotides, and tetraloops.




Figure 1: The image shows an overlay of six average structures of the Drew–Dickerson DNA dodecamer (CGCGAATTCGCG)₂ from a 15-microsecond simulation (using the 5 to 10 μ s region, all-atom RMS fit) using six different salt combinations: LiCl, LiBr, NaCl, NaBr, KCl, and KBr, each of 200mM concentration.

METHODS & CODES

As the team continued to develop and optimize the AMBER suite of molecular dynamics (MD) and free energy set-up, simulation, and analysis codes and force fields, these codes were applied on Blue Waters to understand the structure, dynamics, and interactions of biomolecules in their native environment. The AMBER molecular dynamics engine PMEMD is well-optimized in CUDA for high performance on GPUs. Application of enhanced sampling methodologies, including multidimensional replica exchange molecular dynamics (M–REMD) with ensembles of independent MD simulations that exchange information periodically—including temperature and Hamiltonian changes—provides a very efficient means of reproducibly converging the conformational ensembles of various biomolecular systems. As the ensembles of independent MD simulations generate big data in the form of time series of atomic positions or trajectories that must be sorted and analyzed, we parallelized the CPPTRAJ trajectory analysis code to enable efficient processing of the large data on Blue Waters' parallel file system.

RESULTS & IMPACT

Using M–REMD methods, we have demonstrated the ability to reliably and reproducibly converge the conformational ensembles of DNA helices and RNA dinucleotides, tetranucleotides, and tetraloops. This means that we can elucidate the thermally accessible set of conformations for a given biomolecule at a given temperature. The AMBER biomolecular force fields, particularly the parmbsc1 and X_{OL3} (contained within the ff14SB AMBER force field designation), perform incredibly well for nucleic acid helices. Fig. 1 shows the remarkable agreement observed in the central base pairs with RMS deviations of less than 0.1 Angstroms between all six conformations of 5 μ s average structures of the DNA duplex d(CGCGAATTCGCG)₂ from MD simulations in different salts. The agreement to the average NMR structures from the highest resolution structure of this system (1NAJ) is less than 0.5 Angstroms with both of the high-performing AMBER force fields (parmbsc1, ff14SB) in OPC water. Note that some disruption of structure is observed with the first and second base pairs (and even into the third base pair) at both ends of the DNA, which shows a deviation of ~1 Angstrom due to fraying and base-flipping events that occur on the microsecond timescale. This observation emphasizes the importance of convergence for these particular systems and the lack of structural impact on the DNA regardless of the salt used to neutralize the system. The research enabled by Blue Waters also demonstrated convergence of the dynamic conformational ensemble of various RNA molecules; multiple conformations are populated over the course of the M–REMD simulations. However, some of these are incompatible with experiment, suggesting that the AMBER force fields are over-populating anomalous conformations. These are likely due to very subtle misbalances in the RNA force field in terms of nonbonded interactions (van der Waals and hydrogen bonding)

that still require optimization to improve the force fields. This work is underway currently.

WHY BLUE WATERS

The Blue Waters petascale resource allowed our team to converge the conformational ensembles of various nucleic acid systems for the first time, and this was enabled by the developments in both the AMBER codes and force fields, and by the exceptional performance of the AMBER PMEMD engine on GPUs. The thousands of GPUs and well-performing Lustre parallel file system facilitated our simulation and analysis workflow. The Blue Waters team helped us overcome problems and facilitated some CUDA, OpenMP, and MPI optimizations of the CPPTRAJ analysis code.

PUBLICATIONS & DATA SETS

The AMBER WWW pages: <http://ambermd.org>.

Data sets from some runs on Blue Waters: <http://amber.utah.edu>.

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

Galindo-Murillo, R., D.R. Roe, and T.E. Cheatham III, Convergence and reproducibility in molecular dynamics simulations of the DNA duplex d(GCACGAACGAACGAACGC). *Biochimica Biophys. Acta*, 1850 (2015), pp. 1041–1058.

Cheatham, T.E., III, and D.R. Roe, The impact of heterogeneous computing on workflows for biomolecular simulation and analysis. *Computing in Science and Engineering*, 17:2 (2015), pp. 30–39.

Galindo-Murillo, R., et al., Intercalation processes of copper complexes in DNA. *Nuc. Acids Res.*, 43 (2015), pp. 5364–5376.

Bergonzo, C., et al., Highly sampled tetranucleotide and tetraloop motifs enable evaluation of common RNA force fields. *RNA*, 29 (2015), pp. 1578–1590.

Bergonzo, C., and T.E. Cheatham III, Improved force field parameters lead to a better description of RNA structure. *J. Chem. Theory Comp.*, 11 (2015), pp. 3969–3972.

Bergonzo, C., K.B. Hall, and T.E. Cheatham III, Stem-loop V of Varkud satellite RNA exhibits characteristics of the Mg²⁺ bound structure in the presence of monovalent ions. *J. Phys. Chem. B*, 119 (2015), pp. 12355–12364.

Robertson, J.C., and T.E. Cheatham III, DNA backbone BI/BII distribution and dynamics in E2 protein-bound environment determined by molecular dynamics simulation. *J. Phys. Chem. B*, 119 (2015), pp. 14111–14119.

Galindo-Murillo, R., D.R. Davis, and T.E. Cheatham III, Probing the influence of hypermodified residues within the tRNA_{3^{Lys}} anticodon stem loop interacting with the A-loop primer sequence from HIV-1. *Biochimica Biophys. Acta*, 1860 (2016), pp. 607–617.

Waters, J.T., et al., Transitions of double-stranded DNA between the A- and B- forms. *J. Phys. Chem. B*, 120 (2016), pp. 8449–8456.

Bergonzo, C., K.B. Hall, and T.E. Cheatham III, Divalent Ion Dependent Conformational Changes in an RNA Stem-Loop

Continued on page 215

IMPROVING NWCHEM SCALABILITY USING THE DATASPACE FRAMEWORK

Allocation: Innovation and Exploration/525 Knh

PI: Gregory Bauer¹

Collaborators: Victor Anisimov¹, Manish Parashar², Melissa Romanus²

¹National Center for Supercomputing Applications

²Rutgers University

EXECUTIVE SUMMARY

Molecular Dynamics (MD) is the major computational technique to bridge the gap between experiment and theory in materials science, engineering, and biomedical research. However, the predictive ability of MD simulations strongly depends on the quality of underlying parameters. The purpose of the present phase of this project is to develop a parameter optimization tool for the AMBER classical force field for DNA, with the potential for extending the underlining methodology to optimization of other popular force fields. The project employs the previously generated data set of experimental quality base–base interaction energies prepared by conducting high-level quantum-mechanics CCSD(T) computations in NWChem on molecular clusters extracted from experimental crystallographic data for DNA bases. The parameter optimization procedure performs a grid-based scan on a set of parameters by trying all their possible combinations, computing the interaction energy for each grid point using the Amber force field, and comparing the result to the reference interaction energy. The computation runs in parallel and returns a set of parameters that best reproduces the target data.

RESEARCH CHALLENGE

The present computational challenge is handling a combinatorial explosion in problem size that accompanies the task of identifying the global minimum in parameter space. The natural solution to this challenge is to employ parallelization. However, the use of a large number of parallel processes in the grid search exacerbates the I/O load on the filesystem since each process frequently performs read and write operations. That places a limit on the number of parallel tasks that can be practically used in the computation without overloading the filesystem. When the grid search is done, each process carries a multidimensional array of results that associates the point in the parameter space with the optimization function. The aggregate distributed array holds billions of records for an average parameter optimization problem. Sorting an array of such size to determine the promising parameter sets for further analysis represents a practically unsolvable problem that requires a creative solution.

METHODS & CODES

Generation of target data for parameter optimization uses a set of in-house scripts to extract molecular clusters from the experimental crystallographic data of DNA bases. Quantum mechanical computations, which follow, determine the base–base interaction energy in the molecular clusters. The computation employs the previously optimized CCSD(T) method [1] in the NWChem package [2]. 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 16,384 XE nodes using 32 cores per node resulting in the use of 524,288 processing units on Blue Waters. The optimization tool generates an adjustable number of alternative parameter sets of comparable quality for further testing in molecular dynamics simulations.

RESULTS & IMPACT

This project introduces a procedure for systematic improvement of classical force field by determining the global minimum in the parameter space for an expandable set of the training data. The beneficiary of the optimized parameter set is the entire molecular dynamics community. As the number and quality of the training data increase with time, rerunning the parameter optimization tool will deliver the improved parameter set. The developed fractional parallel sorting procedure drastically reduces time spent in sorting as well as the required RAM per node. The use of RAM disk for read / write operations on compute nodes eliminates the filesystem overhead and makes the code applicable to compute systems beyond Blue Waters' size.

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.

Continued from page 213

Observed by Molecular Dynamics. *J. Chem. Theory Comp.*, 12 (2016), pp. 3382–3389.

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.

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.

Hao, Y., et al., Molecular basis of broad-substrate selectivity of a peptide prenyltransferase. *PNAS*, 113 (2016), pp. 14037–14042.

Hayatshahi, H.S., et al., Computational Assessment of Potassium and Magnesium Ion Binding to a Buried Pocket in the GTPase-Associating Center RNA. *J. Phys. Chem. B*, 121 (2017), pp. 451–462.

Zgarbova, M., et al., Influence of BII Backbone Substates on DNA Twist: A Unified View and Comparison of Simulation and Experiment for all 136 Distinct Tetranucleotide Sequences. *J. Chem. Info. Model.*, 57 (2017), pp. 275–287.

Wang, Y., et al., Application of thiol-yne/thiol-ene reactions for peptide and protein macrocyclizations. *Chemistry*, 23 (2017), pp. 7087–7092.

Hayatshahi, H.S., C. Bergonzo, and T.E. Cheatham III, Investigating the ion dependence of the first unfolding step of GTPase-associating center ribosomal RNA. *J. Biomol. Struct. Dyn.*, 1:11 (2017), pp. 243–253.

Bergonzo, C., and T.E. Cheatham III, Mg²⁺ binding promotes SLV as a scaffold in Varkud Satellite Ribozyme SLI-SLV kissing loop junction. *Biophys. J.*, 113 (2017), pp. 313–320.

Galindo-Murillo, R., and T.E. Cheatham III, Computational DNA binding studies of (-)-epigallocatechin-3-gallate. *J. Biomol. Struct. Dyn.*, 3 (2017), pp. 1–13 (2017).

Hayatshahi, H.S., N.M. Henriksen, and T.E. Cheatham III, Consensus conformations of dinucleotide monophosphates described with well-converged molecular dynamics simulations. *J. Chem. Theory Comp.*, 14 (2018), pp. 1456–1470.

Cornillie, S.P., et al., Computational modeling of stapled peptides toward a treatment strategy for CML and broader implications in the design of lengthy peptide therapeutics. *J. Phys. Chem. B*, 122 (2018), pp. 3864–3875.

Galindo-Murillo, R., T.E. Cheatham III, and R.C. Hopkins, Exploring potentially alternative non-canonical DNA duplex structures through simulation. *J. Biomol. Struct. Dyn.*, (2018) DOI:10.1080/07391102.2018.1483839.

STRUCTURAL BASIS FOR EXTREME COLD TOLERANCE IN THE EYE LENSES OF TELEOST FISHES

Allocation: Illinois/113.5 Knh
PI: Chi-Hing 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 are responsible for maintaining proper density of the lens and 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 that is 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 at temperatures much

colder than that faced by mammalian lenses. 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 isoforms.

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 six to seven 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, likely to maintain a refractive index gradient to correct for spherical aberration, with increasing concentration from the cortex to the nucleus. If the propensity of phase separation is due to enhanced flexibility driven by density, we expect elevated flexibility among lens crystallins that predominate in the dense lens nucleus compared to crystallins found in the less concentrated lens cortex.

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.

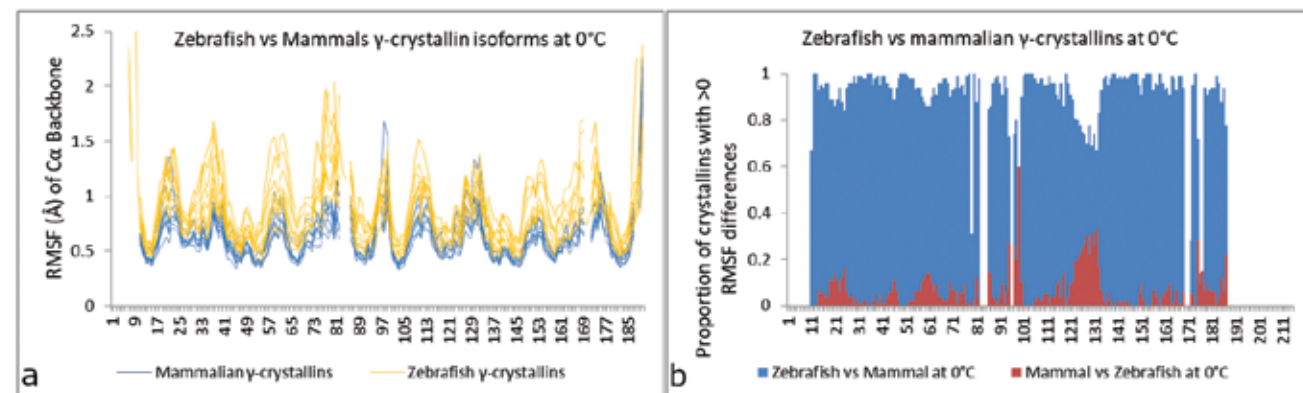


Figure 1: (a) RMSF of eight mammalian γ -crystallin isoforms (blue) and 12 zebrafish isoforms (orange) at 0°C. (b) Proportion of γ -crystallins with RMSF differences among zebrafish γ -crystallins (blue) and mammalian γ -crystallins (orange) that are greater than zero at 0°C.

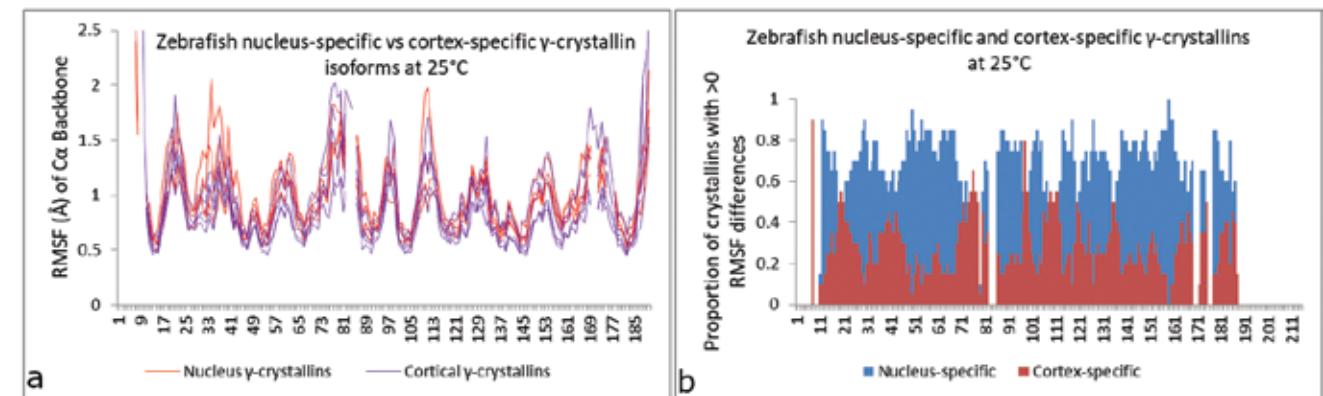


Figure 2: (a) RMSF of zebrafish γ -crystallin isoforms in the dense nucleus (red) and zebrafish γ -crystallin isoforms in the less concentrated cortex (purple) at 25°C. (b) Proportion of zebrafish nucleus γ -crystallins (red) with RMSF differences among zebrafish cortical γ -crystallins (purple) and mammalian γ -crystallins that are greater than zero at 0°C.

METHODS & CODES

We ran molecular dynamics (MD) simulations on twelve zebrafish and eight mammalian isoforms at a cold temperature (0°C), and at the normal body temperature (25°C and 37°C respectively). Three replicates of each γ -crystallin isoform were simulated for 50 nanoseconds (ns) in NAMD [7] 2.12 using CHARMM27 force field parameters. Each of the five mammalian γ -crystallin isoforms were simulated using solved structures obtained from the Protein Data Bank (PDB). Simulation of zebrafish γ -crystallin isoforms used one known structure, the γ M7-crystallin, and remaining 11 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 α atoms per residue of a protein over the duration of the simulation. Average RMSF values of the last 30 ns were taken for each isoform, then formatted based on amino acid sequence alignment generated by MUSCLE 3.8.31 for comparison. To gauge the proportion of zebrafish γ -crystallins that are more flexible than mammalian γ -crystallins at 0°C, RMSF differences were calculated by each individual crystallin per site, and the sum of each instance was divided by the total number of comparisons. This method was also used to determine the proportion of zebrafish nucleus and cortex-specific γ -crystallins at 25°C.

RESULTS & IMPACT

At 0°C, it is evident that zebrafish γ -crystallins are more flexible than the mammalian isoforms across all sites (Fig. 1a). Zebrafish γ -crystallins demonstrate that nearly all isoforms are more flexible than mammalian crystallins (Fig. 1b). In Fig. 2a, zebrafish γ -crystallins predominately in the nucleus are slightly more flexible than those in the cortex at a majority of sites at 25°C (Fig. 2b). A distinct disparity in flexibility is present among zebrafish and mammalian γ -crystallins, and the elevated flexibility of teleost γ -crystallins may account for their remarkably cold-tolerant lenses. Density may have played a large role toward

selection for more flexible crystallins in teleost fishes by preventing interactions among neighboring crystallins that would result in loss of transparency in the lens. In turn, the elevated flexibility may increase the thermodynamic barrier to phase separation at cold temperatures well below thermal habitats occupied by teleost fishes.

WHY BLUE WATERS

Our work requires simulating three trials of 49 proteins at two temperatures, and over a long timecourse of 50 ns to detect meaningful molecular behavior. This work is at the core of a PhD 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.

COMBINING PHYSICS AND SUPERCOMPUTERS TO PREDICT PROTEIN STRUCTURES FROM SEQUENCES

Allocation: NSF PRAC/5,750 Knh

PI: Ken Dill¹

Co-PI: Alberto Perez¹

¹State University of New York at Stony Brook

EXECUTIVE SUMMARY

The team uses physics and supercomputers to reduce the time and cost of determining what proteins look like, atom by atom. This has been a grand challenge in computational biology for the last 50 years. We have developed the MELD (Modeling Employing Limited Data) method to run on GPUs and to combine physics with additional information using Bayesian inference. This reduces the computer time needed to fold proteins and increases the accuracy of the results.

Our MELD approach is the only atomistic physics-based approach to predict a protein's structure from its sequence that is fast enough to compete in a worldwide protein structure prediction competition (CASP). This competition has very strict deadlines, and so most methods used by participants are bioinformatics-based. We have shown that our method can be competitive and that it can bring new insights to the field. And, because it is physics-based, it gives us the option to know alternate states and pathways.

RESEARCH CHALLENGE

Proteins consist of a sequence of amino acids that folds into a 3D shape. This shape allows the protein to do its work inside our bodies. Proteins are drug targets, so developing new drugs requires knowing what these proteins look like. However, determining the structures experimentally is time-consuming, expensive, and not always possible. Bioinformatics tools are good at detecting proteins whose sequence is similar to known structures. Thus, these methods use the known structure for drug design—but there are many proteins that cannot be tackled with this approach.

Physics-based computer simulations mimicking the folding process are an alternative methodology. They used to require years of simulation time and returned incorrect structures. With new force fields [1] and our new MELD [2,3] computer method that leverages external information, we are in a great position to tackle this problem. It is very timely, as well, to collaborate with experimental labs that can produce proteins with complicated sequence information.




Figure 1: Past CASP results produced with MELD and Blue Waters. On top are three protein structures predicted blindly (blue ribbons), along with the comparison with the experimental structure (red ribbons). Below are the results of all groups in CASP (grey lines) and our prediction (blue line). All the results were calculated by the CASP prediction center and posted online (<http://www.predictioncenter.org>). For T0816, MELD was the only one to provide high-accuracy structures. (Figure adapted with permission from [5].)




Figure 2: MELD allows us to identify states to seed multiple unbiased simulations and recover pathways. In this example, we show all states and pathways visited in exploring a helix-to-helix transition in which the system transitions from a left-handed helix conformation to a right-handed one. (Figure adapted with permission from [6].)

METHODS & CODES

We developed a plugin (MELD) to the Molecular Dynamics (MD) package OpenMM. MELD consists of a Hamiltonian and Temperature replica exchange MD protocol in which the Hamiltonian varies according to external information coming from experiment, general knowledge, or bioinformatics. What is unique about MELD is that the information is expected to be unreliable. Hence, rather than enforcing all of it, we impose only a fraction. The part to be enforced changes at every timestep and is chosen in a deterministic way.

RESULTS & IMPACT

With Blue Waters and MELD we are able to compete in a time-sensitive protein structure prediction competition (CASP [4]) that would otherwise be impossible. The competition occurs from May to August; each weekday several sequences are released, along with a three-week deadline to submit the structure predictions. With standard computer resources we might be able to target a few proteins during this time period but not the 100 or more independent protein systems that we target with the help of Blue Waters.

With Blue Waters, we have shown the ability to predict structures that fail with all other state-of-the-art methods (see Fig.1) [5]. These proteins, known as being “unthreadable,” come from orphan genes and have no homology to other known proteins. The abundance of unthreadable proteins encouraged us to specifically target these systems. We have already successfully folded a handful of them and look forward to folding more.

In addition, we are using the states predicted by MELD to predict the pathways that proteins fold to or use to transition between different states. We have named this set of methodologies “MELD-path” and recently published the first work produced with this methodology (see Fig. 2) [6]. We derived a set of nearly 300 states from MELD simulations and then started ~13,800 unbiased simulations using Amber on GPUs. For this system, each trajectory

had to run for less than an hour. With infinite GPU resources it would have taken just one hour to run all the simulations. Using the backfilling queue on Blue Waters, we managed to collect all the data in under 17 days. This impressive feat would have taken us around 255 days to accomplish on a single GPU. Using Markov State Modeling theory we were able to recover the most relevant pathways for a helix-to-helix transition. The pathways clearly showed a preference for unfolding or refolding starting from either end of the helix and proceeding sequentially, rather than starting in random places on the helix.

WHY BLUE WATERS

Blue Waters is the only system in the United States that has enough GPUs for us to compete in CASP, and the only one that allows many jobs requiring a relatively low number of GPUs (30 each) to run for up to 48 hours. Members of the Blue Waters staff provided invaluable support in the compilation of both the Amber and OpenMM/MELD packages, which required nontrivial effort, especially during the deployment of the new Python site libraries. Furthermore, our conversations with the staff have been instrumental in improving the efficiency of running jobs during the CASP competition.

PUBLICATIONS & DATA SETS

Perez, A., et al., MELD-Path Efficiently Computes Conformational Transitions, Including Multiple and Diverse Paths. *J. Chem. Theory Comput.*, 14 (2018), pp. 2109–2116.

ALGORITHMS FOR RECONSTRUCTING THE LIFE HISTORY OF TUMORS FROM GENOMICS DATA

Allocation: Director Discretionary/50 Knh
PI: Mohammed El-Kebir¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Cancer is a genetic disease characterized by intratumor heterogeneity as well as the presence of multiple cellular populations with different sets of mutations. Cellular heterogeneity gives cancer the ability to resist treatment, and quantifying its extent is key to improving our understanding of tumorigenesis. In this project, we developed and employed novel phylogenetic (evolutionary tree-based) techniques to reconstruct the evolutionary histories of individual tumors from DNA sequencing data. Typically, this data is obtained from shotgun sequencing of tumor biopsies using either single-cell or bulk-sequencing technology. Highlights from our research outcomes include: (1) the creation of SPhyR, an algorithm that employs the k-Dollo evolutionary model to reconstruct phylogenetic trees from single-cell DNA sequencing data; and (2) a detailed analysis using simulated data on the nonuniqueness of solutions to the phylogeny estimation problem from bulk DNA sequencing data. One paper based on this work has been accepted for publication and another has been submitted.

RESEARCH CHALLENGE

This research considered two current challenges in tumor phylogenetics. First, the phylogeny estimation problem from single-cell sequencing data is a variant of the classic phylogeny estimation problem with incorrect and missing data due to the sequencing technology. Current methods aim to simultaneously construct a phylogenetic tree and correct these measurement errors using either too stringent or too permissive evolutionary models. There is a need for methods that employ appropriate evolutionary models that strike a balance between being realistic and yet are sufficiently constrained.

Second, the problem of reconstructing a phylogenetic tree given bulk sequencing data from a tumor is more complicated than the classic phylogeny estimation problem, where one is given the leaves of the phylogenetic tree as input. In contrast, by using bulk sequencing data, which forms the majority of current cancer sequencing studies, we do not observe the leaves but rather are given mutation frequencies that result from mixtures of the unknown leaves of the underlying phylogenetic tree. The majority of current tumor phylogeny estimation methods that aim to infer a phylogenetic tree from mutation frequencies employ the perfect phylogeny evolutionary model. In this model, a mutation at a specific genomic site occurs only once throughout the evolutionary history of the tumor and is never subsequently lost. The underlying perfect phylogeny mixture problem is NP-complete and, importantly, from the same input data multiple distinct solutions can be inferred. This nonuniqueness has important consequences for downstream analyses by cancer biologists and clinicians, whose starting point is a single phylogenetic tree. While nonuniqueness of solutions to this computational problem has been recognized in the field, a rigorous analysis of its extent and consequences has been missing.

METHODS & CODES

We used extensive computer simulations to quantify the extent of nonuniqueness of solutions to the phylogeny estimation problem from bulk sequencing data. To solve the phylogeny estimation problem from single-cell sequencing data, we used techniques from combinatorial optimization. More specifically, we formulated the problem as an integer linear program (ILP) and designed and implemented a custom column generation

and cutting plane approach for the ILP formulation. We used CPLEX as the underlying ILP solver. To facilitate reproducibility of the results, the data and open source codes (Python, Bash, and C++) have been deposited in public repositories. In addition, these repositories contain Jupyter notebooks that generate the plots of the respective papers.

RESULTS & IMPACT

Phylogeny estimation algorithms that employ appropriate evolutionary models are key to understanding the evolutionary mechanisms behind intratumor heterogeneity. One of the outcomes of this project is Single-cell Phylogeny Reconstruction (SPhyR), a method for tumor phylogeny estimation from single-cell sequencing data that employs the k-Dollo parsimony model, a relaxation of the perfect phylogeny model where a mutation may only be gained once but lost k times. In light of frequent loss of point mutations in cancer due to copy-number aberrations, the k-Dollo model is more appropriate than the evolutionary models utilized by previous methods. This project resulted in a novel combinatorial characterization of solutions to the underlying computational problem as constrained integer matrix completions, which formed the basis for the efficient integer linear programming approach utilized by SPhyR. Using simulated data, we found that SPhyR outperformed existing methods that are either based on the infinite-sites or the finite-sites evolutionary model, in terms of solution quality and runtime.

We studied the problem of counting and sampling solutions in instances of the phylogeny estimation problem from bulk sequencing data. To avoid any bias in downstream analyses it is important to know the number of solutions and to be able to sample uniformly from the solution space. As part of this project, we proved an upper bound on the number of solutions that can be computed in polynomial time. In addition, we introduced a uniform sampling algorithm based on rejection sampling that

works for small problem instances. Using extensive simulations, we showed that the number of solutions increased with an increasing number of mutations, but decreased with increasing number of bulk samples from the same tumor. We observed similar trends in terms of the quality of the solutions in the solution space. Moreover, we showed that additional constraints from single-cell and long-read sequencing technology significantly reduced the number of solutions. Finally, we demonstrated that current methods are unable to sample uniformly from the solution space, often overlooking solutions. This leads to significant biases that propagate to downstream analyses.

WHY BLUE WATERS

Blue Waters was essential to the two research outcomes of this project, as they both involved extensive benchmarking and validation using simulated data. The computational resources of Blue Waters allowed us to perform these experiments at scale, enabling us to study the performance of our algorithms and the underlying problem statements in many different experimental settings. This is something that would not have been possible on other platforms.

PUBLICATIONS & DATA SETS

El-Kebir, M., SPhyR: Tumor Phylogeny Estimation from Single-Cell Sequencing Data under Loss and Error. *Bioinformatics/ECCB 2018*, accepted (2018).

Pradhan, D., and M. El-Kebir, On the Non-uniqueness of Solutions to the Perfect Phylogeny Mixture problem. Submitted (2018).

OncoLib: <https://github.com/elkebir-group/OncoLib>
 SPhyR: <https://github.com/elkebir-group/SPhyR>
 PPM-NonUniq: <https://github.com/elkebir-group/PPM-NonUniq>




Figure 2: *In silico* simulations of tumor evolution and bulk sequencing show that the number of solutions to the phylogeny estimation problem increased with an increasing number of mutations but decreased with an increasing number of bulk samples from the same tumor.

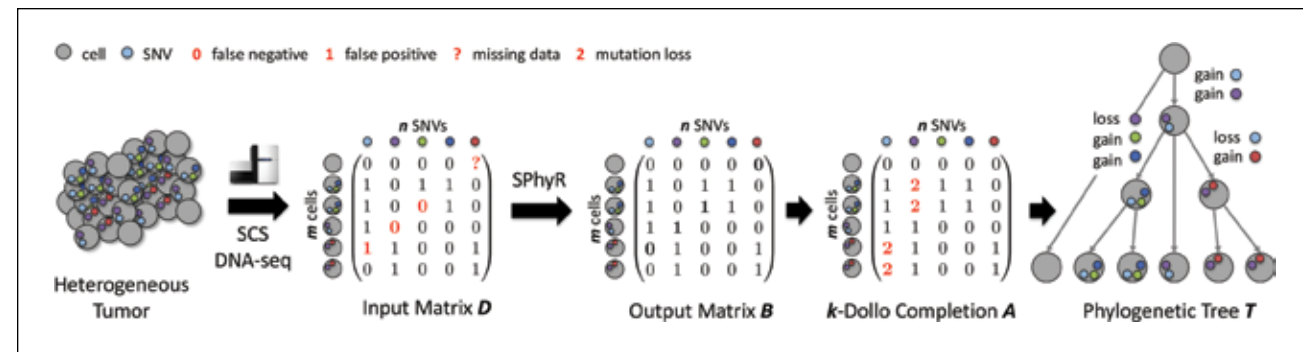


Figure 1: Tumors are composed of cellular populations with distinct sets of mutations. Single-cell sequencing of a tumor yields an input matrix D with incorrect and missing entries. SPhyR aims to simultaneously correct errors in matrix D and infer the evolutionary history of the cells.

IN SILICO VACCINE DESIGN THROUGH EMPIRICAL FITNESS LANDSCAPES AND POPULATION DYNAMICS

Allocation: Exploratory/50 Knh

PI: Andrew Ferguson¹

Collaborator: Greg Hart¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

The hepatitis C virus (HCV) affects 170 million people worldwide, and kills 700,000 annually. Vaccination provides the most realistic and cost-effective hope of controlling this epidemic, but no vaccine is available. Computational models can offer rational precepts to inform and accelerate vaccine design. We have developed a computational tool to translate databases of viral sequences into “fitness landscapes,” mapping the replicative capacity of the virus as a function of its genome. These landscapes represent the mutational playing field over which the virus evolves. By integrating these landscapes with agent-based models of viral mutation and host immune response, we have explicitly modeled the host–pathogen dynamics over its empirically defined fitness landscape. (Agent-based models simulate the actions and interactions of autonomous agents to assess their effects on the system as a whole.) Using this simulator, we have employed the hardware resources of Blue Waters to perform computational screening of candidate vaccine components to identify those best able to cripple viral fitness and block immune escape. These findings can inform next-generation HCV vaccine design.

RESEARCH CHALLENGE

HCV continues to pose a global threat to public health. Despite the availability of efficacious drug treatments in the developed world, the high cost of these therapies make them effectively unavailable in the developing world where the preponderance of infections occur. A prophylactic vaccine represents the most cost-effective and realistic strategy to combat the epidemic, but despite 25 years of research, a vaccine is still not available. A challenge in vaccine design is the identification of promising targets within the virus that can be targeted by a vaccine that simultaneously cripple viral fitness and are not subject to facile mutational escape, whereby a microorganism defends itself from host immune responses by making mutations in its genotype and phenotype. Computational models of viral infection and the host immune response can systematically identify promising targets that may be translated into rational precepts for experimental development and testing of HCV vaccines.

METHODS & CODES

The simulations of the viral mutational evolution over our viral fitness landscapes is implemented via an agent-based model comprising 50,000 distinct viral sequences. The host immune response is described by a set of ordinary differential equations modeling the dynamics of the host T-cells as they recognize the virus, activate, mature, proliferate, and die. The coupling to the viral dynamics occurs through a term imposing a penalty on the fitness of viral strains that are recognized and attacked by particular members of the T-cell population, and through a recognition term in which T-cells that recognize particular viral strains are primed to activate and proliferate. The relatively small T-cell populations within our control volume mean that fluctuations are important, and we implement a stochastic integration protocol via Gillespie dynamics to explicitly capture these effects.

RESULTS & IMPACT

We considered two representative hosts in detail and used our simulator to predict the efficacy of the ensemble of all possible vaccine candidates consistent with the immunological genotypes of the hosts. In each case, we identified a number of promising vaccine candidates that led to strong and durable responses by priming T-cells that imposed strong fitness penalties upon the viral population for long periods of time. Interestingly, we also found vaccine candidates that led to poorer immune responses compared to no vaccination by priming the “wrong” T-cell responses to attack regions of the virus from which mutational escape is facile. These results lay the foundation for large-scale simulations of vaccine candidates for all representative hosts in the North American population for the particular protein considered, and for extending this work to 10 HCV proteins.

WHY BLUE WATERS

The scale and parallelism available within Blue Waters support the computational intensity of each single simulation and also enable the range of simulations to evaluate large numbers of vaccine candidates in a variety of hosts. Furthermore, the volume of data generated is also significant as the viral sequences present at each time point must be written to disc.

PUBLICATIONS & DATA SETS

Hart, G.R., and A.L. Ferguson, Computational design of hepatitis C virus immunogens from host-pathogen dynamics over empirical viral fitness landscapes. To be submitted (2018).




Figure 1: The black crosses characterize the immune response of the inoculated host relative to no vaccination (green circle). Red circles indicate the Pareto optimal vaccine candidates. Blue circles indicate previously identified vaccine candidate. Candidates in the upper-left quadrant provide superior strength and length of control relative to no vaccination.

MODELING AND SIMULATIONS OF COMPLEX DYNAMIC MUSCULOSKELETAL ARCHITECTURES

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 between its senses, body, and the surrounding water. More generally, during millions of years of evolution, animals have refined sophisticated design solutions to master their complex interplay with 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 continuous, distributed biological actuation and control, we have developed a novel mathematical formulation based on assemblies of Cosserat rods, rods with deformable cross-sections, for the modeling of complex soft musculoskeletal architectures interacting with surrounding environments. These techniques are embedded in an automated inverse design cycle based on evolutionary optimization. This allows us to set a desired trait (speed, dexterity, energy efficiency) and to identify optimal designs and muscle actuation sequences in response to sensory cues.

RESEARCH CHALLENGE

All animals 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 engineering designs and work in ways we still do not fully understand. These solutions 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 engineering 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 manmade solutions obtained through inverse design based on an automated optimization process may outperform pure biomimicry.

METHODS & CODES

The characterization of biopropulsion from an optimality standpoint 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 penalization techniques, to capture the effects of the body on the fluid [1,2]. This methodology is coupled with a musculoskeletal solver developed to capture the compliant dynamics of musculoskeletal systems made of bones, tendons, and muscles [3].

RESULTS & IMPACT

We have developed *Elastica* [3], a software able to capture the dynamic response of complex musculoskeletal structures. We validated our solver on a number of benchmark problems with known analytic solutions and against experimental investigations involving (self-)contact, anisotropic surface friction, and highly viscous and inertial fluids. We also verified the solver experimentally in the context of artificial muscles, linking topological changes to mechanical work, and extended it to simulate arbitrary complex musculoskeletal layouts from wings to human elbow joints (Fig. 1).

We then employed *Elastica* to computationally design, simulate, and optimize for the first time the structure of a biohybrid walking bot (Fig. 2). In collaboration with experimentalists at the Micro and Nanotechnology Laboratory, our design was fabricated and tested, confirming our predictive capacity and leading to the largest, fastest locomotive bot to date [4].

WHY BLUE WATERS

Blue Waters' sheer size and cutting-edge technology enable optimization processes that entail thousands of simulations. This allows the design of unprecedented biological architectures, bringing novel high-impact applications within reach, from soft robotics and biomedicine to precision manipulation and fabrication.

PUBLICATIONS & DATA SETS

Gazzola, M., L. Dudte, A. McCormick, and L. Mahadevan. Forward and inverse problems in the mechanics of soft filaments. Accepted in *Royal Society Open Source*, (2018), arXiv:1607.00430v2.

Pagan-Diaz, G.J., et al., Simulation and fabrication of stronger, larger and faster walking biohybrid machines. *Advanced Functional Materials*, (2018), DOI:10.1002/adfm.201801145.



Figure 2: Top view of three thumb-sized biohybrid machines in a culture dish. Modeling and simulations were employed to design and optimize function and performance. These biohybrid machines walk on the dish surface at about one millimeter per second.

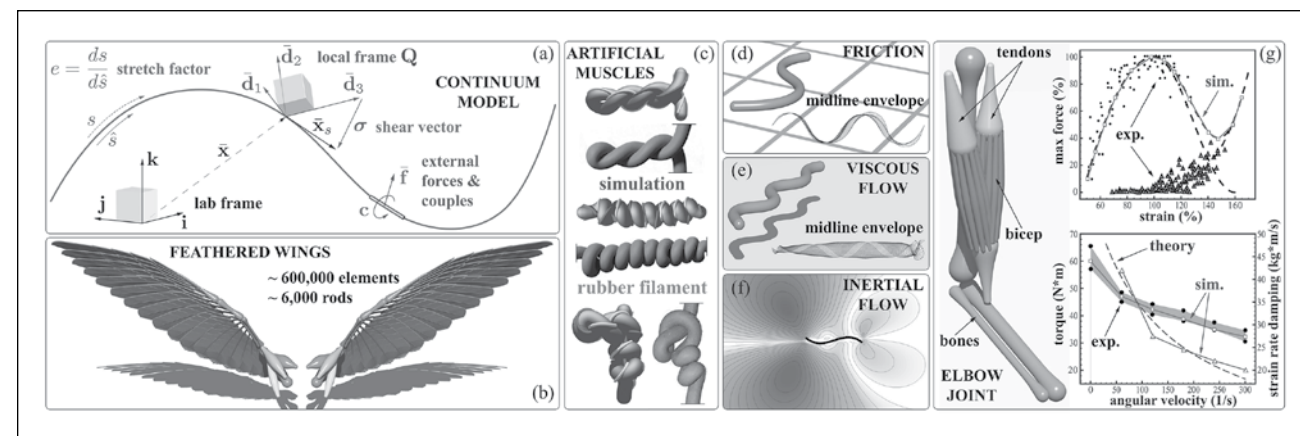


Figure 1: *Elastica* capabilities. (a) Continuum model. Applications to (b) feathered wings, (c) topology and mechanics in artificial muscles, (d) slithering snake, (e) swimming spermatozoa, (f) flexible foil in inertial flow, (g) elbow joint.

CONSTRUCTING LARGE EVOLUTIONARY TREES ON SUPERCOMPUTERS

Allocation hours: Exploratory/50 Knh
PI: William Gropp¹
Co-PIs: Erin Molloy¹, Tandy Warnow¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

The organization of molecular sequences into evolutionary trees, called phylogenies, enables researches to study the evolution of bacteria and viruses that cause disease outbreaks and to identify previously unrecognized microbial organisms found in environmental samples, such as the soil or the human gut. The current and leading approaches to phylogenetic inference require two steps: first, a multiple sequence alignment is estimated and then a tree is estimated from the alignment. This two-phase approach is not scalable. We used Blue Waters to design and test a novel approach that bypasses (1) alignment estimation on the full dataset, (2) maximum likelihood tree estimation on the full dataset, and (3) supertree estimation. This work is a major advancement towards constructing the Tree of Life using supercomputers.

RESEARCH CHALLENGE

The organization of molecular sequences into evolutionary trees, called phylogenies, enables researches to study the evolution of bacteria and viruses that cause disease outbreaks and identify previously unrecognized microbial organisms found in the human

gut [1]. The current and leading approaches to phylogenetic inference require two steps: first, a multiple sequence alignment is estimated and then a tree is estimated from the alignment. A multiple sequence alignment is an $n \times l$ matrix, where n is the number of sequences and l is the alignment length. Building alignments on large numbers of heterogeneous sequences is computationally intensive, and the leading methods produce very long alignments that require large amounts of storage [2]. Maximum likelihood methods are widely accepted as the gold standard for phylogenetic inference on single gene data sets. Building maximum likelihood trees on large numbers of sequences is also computationally challenging given that the number of possible tree topologies increases exponentially with the number of sequences and many numerical parameters must be optimized for each candidate tree topology. Most parallel codes for estimating phylogenetic trees are implemented to handle long alignments but not large numbers of sequences [3, 4]. DACTAL [5], a divide-and-conquer method, handles large numbers of sequences by dividing them into overlapping subsets, constructing trees on subsets, and then merges the trees into a “supertree” using heuristics for

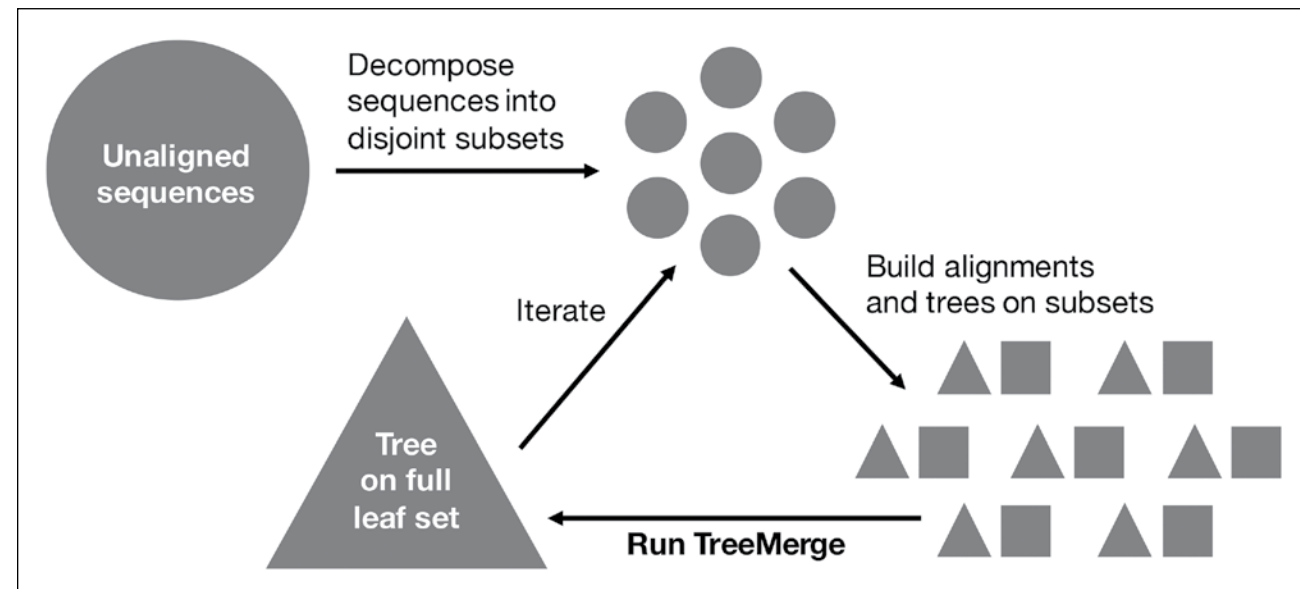


Figure 1: Our approach (TERADACTAL) divides unaligned sequences into disjoint subsets (circles), builds alignments (squares) and trees (triangles) on each subset, and then merges the subset trees together in polynomial time with a highly accurate technique called *TreeMerge*. This process can iterate by decomposing the tree on the full dataset into subsets.




Figure 2: We compared TERADACTAL to methods that compute an alignment (using MAFFT [4] or PASTA [2]) and then estimate a tree (using RapidNJ [5] or RAxML [6]). TERADACTAL achieved an error rate within 1–3% of the leading two-phase method for the two model conditions. Bars are averages over ten replicate datasets.

NP-hard optimization problems. These heuristic methods for constructing supertrees do not scale to ultra-large datasets. Hence, we worked to design, prototype, and test a parallel algorithm for constructing evolutionary trees on large numbers of sequences.

METHODS & CODES

Our approach, called “TERADACTAL,” divides unaligned sequences into disjoint (rather than overlapping) subsets, builds alignments and trees on each subset, and merges the subset trees together using a highly accurate and polynomial time technique we developed called *TreeMerge*. *TreeMerge* builds a minimum spanning tree on the subsets, and the edges in the minimum spanning tree indicate pairs of subset trees to be merged together using constrained Neighbor-Joining. Because the topologies of these merged trees do not conflict with the original subset trees, all of subset trees can be merged together into a tree on the full dataset using the minimum spanning tree.

TERADACTAL can iterate using the tree on the full dataset to create the next subset decomposition. We prototyped TERADACTAL in Python, and it is freely available on Github (<https://github.com/ekmolloy/teradactal-prototype>).

RESULTS & IMPACT

We used the Blue Waters supercomputer to perform a simulation study comparing TERADACTAL to several two-phase methods, including three alignment methods (only two shown: MAFFT [4] or PASTA [2]) and four tree estimation methods (only two shown: RapidNJ [9] and RAxML [8]). RapidNJ is distance-based method

that uses polynomial time to construct a tree from a distance matrix, whereas RAxML [4] is one of the most (if not the most) widely used maximum likelihood codes for phylogenetic inference with over 11,000 citations. We found that TERADACTAL achieves similar error rates (within 1-3%) of the best methods tested. Hence, TERADACTAL achieves similar error rates to the leading two-phase methods but is highly parallel and can handle large number of sequences by avoiding (1) alignment estimation on the full dataset, (2) maximum likelihood tree estimation on the full dataset, and (3) supertree estimation.

This work is a major advancement toward constructing the Tree of Life using supercomputers.

WHY BLUE WATERS

We used Blue Waters to demonstrate that existing parallel codes (e.g., PASTA and RAxML) could not run on datasets with one million sequences on Blue Waters. We also used Blue Waters to extensively test the TERADACTAL prototype and compare the TERADACTAL prototype to the leading and popular two-phase methods. Specifically, we performed a large simulation study requiring over 36,000 node hours. These analyses were completed in under a month but would have required over a year to run on our laboratory’s four nodes on the campus cluster.

PUBLICATIONS & DATA SETS

Molloy, E.K., and T. Warnow, TERADACTAL: A scalable Divide-And-Conquer approach for constructing large phylogenetic Trees (almost) without Alignments. In progress.

TOWARD PREDICTIVE COMPUTATIONAL DESIGN OF PRECISION MOLECULAR OPTOELECTRONICS

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 the key electronic parameters (electron-detachment/attachment energies) of large conjugated molecules and solids—potential components of precision organic optoelectronic devices such as solar cells, light-emitting diodes, field-effect transistors, smart windows, and the like. We transformed the usual, nonscalable sum-of-products expressions of many-body Green's function theories in the complete-basis-set limit into a few high-dimensional integrals, which were then evaluated by a highly scalable Metropolis Monte Carlo algorithm. They efficiently compute energy differences (including quasiparticle energy bands) directly without a sign problem, one of the major unsolved problems in the physics of many-particle systems, on many CPUs or many GPUs, easily achieving an unprecedented speedup by a factor of 31,000 (on 256 GPUs) relative to a serial calculation.




Figure 1: Imaginary-time evolution of electron-pair walkers in the benzene dimer.

RESEARCH CHALLENGE

We are entering an exciting new era of chemical technology in which synthetic chemists can now fabricate complex solid-state materials made of organic components with precise dimensions, so that they display predicted/designed functions and performance as optoelectronic devices. The revolutionary impact of such devices requires no explanation, miniaturizing components for computing, light emission, memory, etc., by many orders of magnitude. Such a research program needs to be assisted by computational methods that can predict the key optoelectronic parameters of the component conjugated organic molecules such as electron-attachment/detachment energies, charge mobility, exciton binding, air stability, etc. A systematically accurate series of approximations for these properties in a molecule and solid exists as many-body Green's function (MBGF) theory [1]. However, as compared with other *ab initio* theories, MBGF is less understood and developed, both theoretically and algorithmically. We aim to address this issue with the aid of Blue Waters.

METHODS & CODES

We mathematically transformed the usual sum-of-products expressions of second-order MBGF (GF2) theory and its complete-basis-set (CBS) correction by explicitly correlated (F12) ansätze into single high-dimensional integrals by a Laplace transform. These integrals (over 12-dimensional coordinates of two coupled electron pairs and one-dimensional imaginary time coordinate; see Fig. 1) are then evaluated by a Metropolis Monte Carlo method with judiciously chosen weight functions. The resulting stochastic methods—Monte Carlo GF2 (MC-GF2) [2] and Monte Carlo explicitly correlated GF2 (MC-GF2-F12) [3]—can compute energy differences (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 scales linearly per MC step and cubic-to-quartic scaling to achieve relative accuracy with negligible memory cost) [4]. The methods can also calculate quasiparticle energy bands of a solid for the entire Brillouin zone as nearly continuous curves of a wave vector [5], and have been extended to third-order MBGF [6] using an expedient interpretation of Brueckner–Goldstone diagrams as well as a convergence-acceleration scheme (“redundant-walker algorithm”) [7].




Figure 2: Molecular components of precision organic electronics devices studied.

RESULTS & IMPACT

We have developed the MC-GF2-F12 method that efficiently executes on both XE and XK nodes, enabling an exact (CBS-limit) GF2 calculation of electron-detachment energies of a wide range of large conjugated organic molecules (Fig. 2), many of which are used in an organic optoelectronic device. Note that the usual divide-and-conquer approach is powerless in treating such a delocalized property of a delocalized electronic structure. The largest calculation was performed on C_{70} with 1,610 basis functions using 128 GPUs for the GF2 portion and 896 CPUs for the F12 part. The result underscores the significance of including both electron-correlation and basis-set-extension effects because the Hartree–Fock calculation with a small basis set gives a deceptively accurate result by error cancellation.

The calculations were accelerated by the redundant-walker algorithm [7], 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. The MC-GF2 algorithm employs a two-level parallelism, in which dense matrix multiplications for many walkers are fine-grained on a GPU, and a Monte Carlo integration is coarse-grained

across multiple CPU-GPUs, thereby enhancing the performance of the redundant-walker algorithm on GPUs beyond the degree that is possible by merely running on many CPUs [4].

WHY BLUE WATERS

The stability and ease-of-use of Blue Waters, 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 & DATA SETS

Johnson, C.M., A.E. Doran, S.L. Ten-no, and S. Hirata, Monte Carlo explicitly correlated second-order many-body Green's function theory. To be submitted (2018).

Johnson, C.M., S. Hirata, S., and S. Ten-no, Explicit correlation factors. *Chem. Phys. Lett.*, 683 (2017), pp. 247–252.

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.

AN EFFICIENT HYBRID STOCHASTIC-DETERMINISTIC SIMULATION TECHNIQUE FOR LIVING CELLS

Allocation: Illinois/169 Knh

PI: Zaida Luthy-Schulten¹

Collaborators: David Bianchi¹, Tyler Earnest¹, Joseph Peterson¹, Michael Hallock¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Stochasticity in gene expression is an important source of noise that can have profound effects on the fate of a living cell. The reactions for gene expression, feedback loops, and transport occurring within cells are typically described by Chemical Master Equations (CME). Sampling the CME using the Stochastic Simulation Algorithm (SSA) results in large computational costs as each reaction event is evaluated explicitly. To improve the computational efficiency of cell simulations involving high-particle-number systems, the authors have implemented a hybrid stochastic–deterministic (CME/ODE) method into the publicly available, GPU-based lattice microbes (LM) software suite, providing a convenient way to simulate complex cellular systems and interface with high-performance CME/RDME/ODE solvers. As a test of the implementation, the authors apply the hybrid CME–ODE method to the galactose switch in *Saccharomyces cerevisiae*, gaining a 10–50× speedup.

RESEARCH CHALLENGE

Many processes within living cells, especially gene expression, are characterized by low particle numbers and a high degree of randomness. The Chemical Master Equation (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 Stochastic Simulation Algorithm (SSA) [1] provides an effective method for obtaining unbiased realizations of these Markov processes. This algorithm is limited by the fact that reaction events are accounted for explicitly by the SSA however, 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 [2–4] in which the highly reactive parts of the system are described by ordinary differential equations (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.

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 (see Fig. 1). At the beginning of each timestep, the LSODA differential equation solver is updated with the species counts obtained from the stochastic regime (transcription, translation) simulated via the SSA, and 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. The hybrid algorithm also communicates updated species counts generated from reactions in the CME regime to the ODE regime at this time. The optimal communication times between the stochastic and deterministic descriptions, as well as the timesteps for each method, need to be assessed 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. The partitioning typically improves the speed of the numerical simulations by a factor of 25–50, 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 in this work. We have already used this hybrid approach to perform a spatially resolved RDME–ODE study (geometry presented in Fig. 2) of the galactose switch system, experiencing similar speedup to that seen in the CME implementation. The geometry of the yeast cell was derived from experimental cryo-electron tomography data through a fitting process and can be easily input using LM [7].




Figure 2: Experiment to LM—A spatially resolved yeast cell to be used in RDME simulation. The plane extending from the cell is the cryo-electron tomogram slice and segmentation data from which the simulation geometry was constructed [7].

WHY BLUE WATERS

Blue Waters was essential to generate thousands of 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 often requires approximately 40 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 a cell.

PUBLICATIONS & DATA SETS

Bianchi, D., et al., Hybrid CME-ODE Method for efficient simulation of the galactose switch in yeast. *IET Systems Biology*, in press (2018), DOI:10.1049/iet-syb.2017.0070.

Earnest, T., et al., Challenges of Integrating Stochastic Dynamics and Cryo-electron Tomograms in Whole-Cell Simulations. *J Phys. Chem. B*, 121:15 (2017), pp. 3871–3881.




Figure 1: Schematic model of the galactose switch in yeast. The reactions depicted in the boxed area are simulated deterministically via an ODE solver, while those outside this region are simulated stochastically using the SSA.

IDENTIFICATION OF MISSING VARIANTS IN ALZHEIMER'S DISEASE, AND THE NEW STANDARDS FOR GENOMIC VARIANT IDENTIFICATION IN LARGE COHORTS

Allocation: Illinois/350 Knh
PI: Liudmila Mainzer¹
Co-PIs: Yan Asmann², Matthew Hudson¹

¹University of Illinois at Urbana Champaign
²Mayo Clinic

EXECUTIVE SUMMARY

Alzheimer's disease (AD) is the sixth-leading cause of death in the United States. While AD is genetically determined, no cure exists and onset remains difficult to predict, as the disease is influenced by a combination of rare and common genomic variants. Correct and complete identification of such variants is essential for advancing research. We have formulated new variant-calling procedures to recover variants previously missed by the community, and have applied this "integration" approach to reanalysis of 10,000 whole human exomes (protein-coding genes) from patients afflicted by AD. By combining two read aligners and several variant callers, we were able to recover 50% of variants that were missed by the standard protocol. This project has delivered a set of newly discovered AD variants for submission into public repositories, as well as new standards for scaling variant calling to cohorts containing tens or hundreds of thousands of samples.

RESEARCH CHALLENGE

Alzheimer's disease (AD) is a neurodegenerative dementia that affects more than five million Americans and more than 35 million people worldwide. It poses an increasing burden to healthcare due to the progressive aging of society. It is hypothesized that AD is shaped by genomic mutations that are highly diverse among afflicted individuals; we have shown that common analytic practice misses a substantial percentage of good-quality genomic variants. Getting the complete variant call set, especially the rare variants, is the critical prerequisite to successful identification of disease predisposition markers and druggable targets in human disease. Our project will deliver novel genomic variants that have remained hitherto undetected by the standard workflow in AD. These variants will be posted into public databases for use by researchers and clinicians worldwide to improve our understanding of the genomic underpinnings of AD, as well as drug development and treatment outcome prediction.

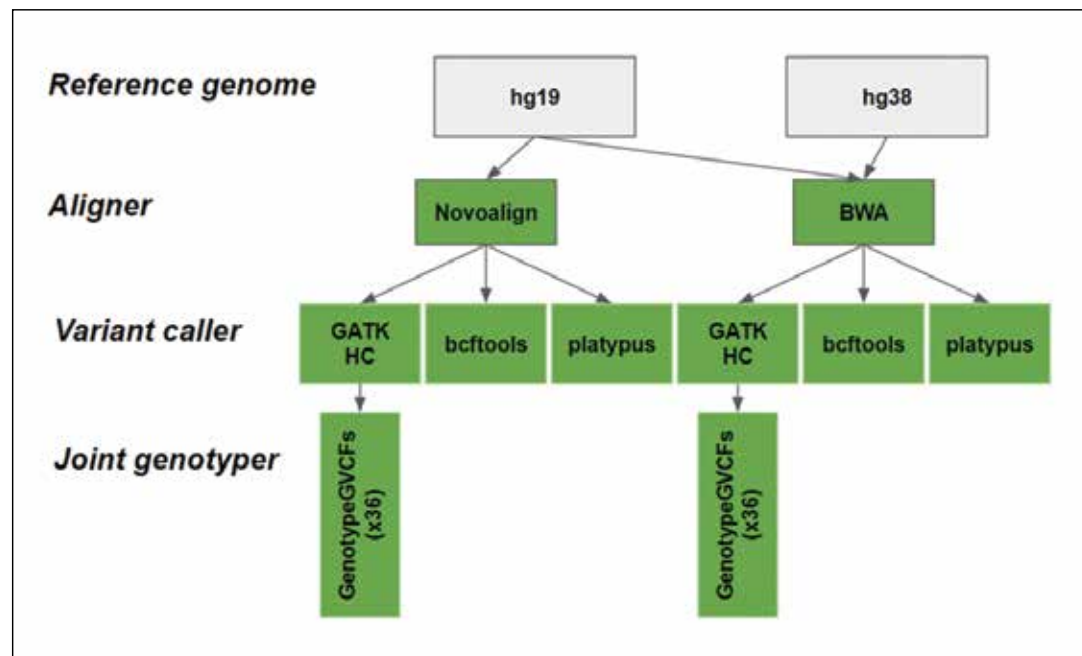


Figure 1: Workflow configurations utilized in this study. Two reference genomes, two aligners, and three variant callers have been included in various permutations to recover as many variants as possible. Joint genotyping was performed where possible.

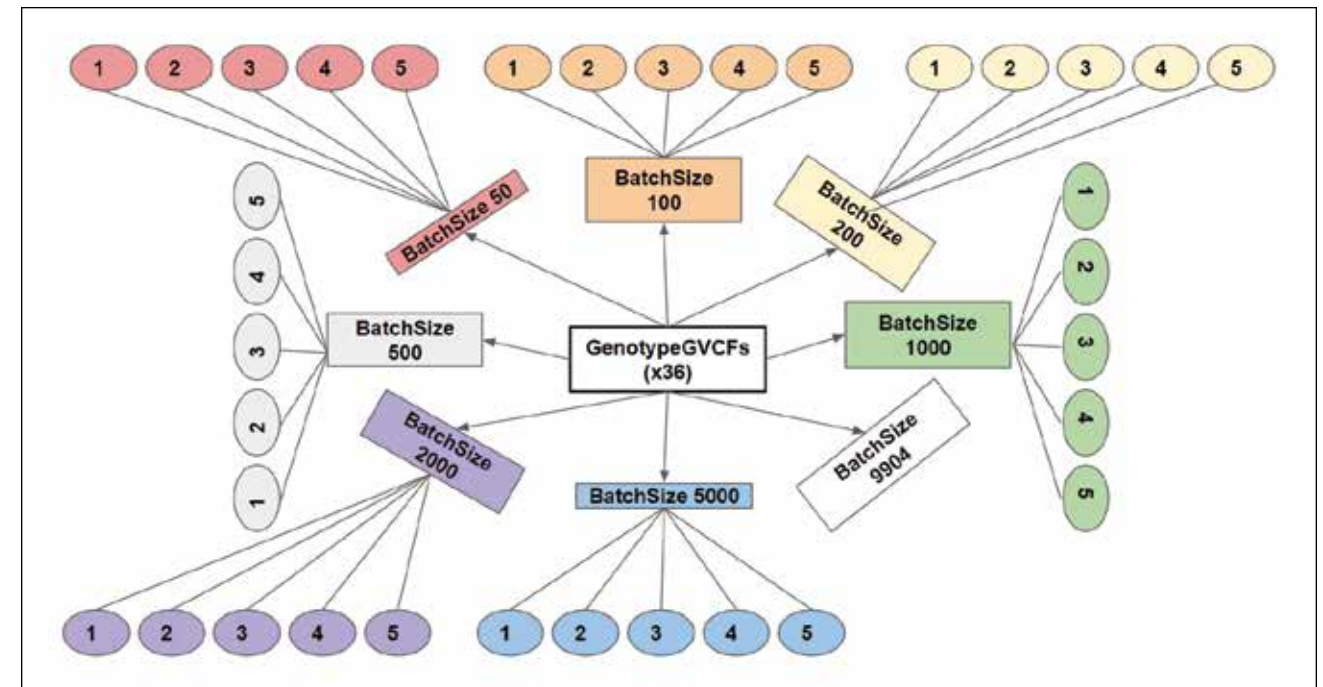


Figure 2: To study the effects of sample size, we performed joint genotyping on variable-size batches subsampled from 10,000 whole exomes. We subsampled each batch size five times, and ran joint genotyping to see how many variants could be recovered. See https://github.com/dpwickland/JointGenotyping_BW.

METHODS & CODES

We tested a select number of steps and parameters in the variant detection pipeline in the context of sample sizes. There were no published works to prioritize the set of steps or parameters to test. We grouped the Alzheimer Disease Sequencing Project (ADSP) samples into different sample sizes of 50, 500, 1,000, 2,000, 5,000, and 10,000. For each sample size, we tested two different aligners, three different variant callers, multi- vs. single-sample variant calling, and five different parameter settings in the variant calling and quality control process. The goal was to build a new set of guidelines for variant discovery based on project size.

RESULTS & IMPACT

We constructed a flexible workflow using the Swift/T workflow management language that allowed easy swapping of tools, quality control, sample analysis parallelization, and job dependencies. By combining two read aligners and several variant callers into our workflow, we were able to recover 50% of the variants in the ADSP data that were missed by the standard protocol. Importantly, recovered variants had higher proportions of low-frequency variants, which are of most interest. We further annotated SNPs, or genetic variations in a single DNA building block, as synonymous or nonsynonymous and assessed the proportion of alternate alleles between cases and controls. We found 47 SNPs within 41 genes that resulted in a switch to a less-frequent codon and showed a greater percentage of alternate alleles in the cases rather than in the controls. 14 of these SNPs, and seven of the top 10 most

significant, lay within genes previously reported to interact with Alzheimer's-related proteins or to function in the brain.

WHY BLUE WATERS

Our study utilizes data from the Alzheimer's Disease Sequencing Project, consisting of over 10,000 whole exome sequencing samples. For each sample, we tested multiple combinations of steps and parameters in the variant detection pipeline. Due to the need to test many different parameter combinations, this study requires a petascale resource. This work is unprecedented for human diseases and traits because the "workflow integration" approach necessary to overcome the inadequacy of individual variant-calling procedures is computationally prohibitive outside a petascale resource like Blue Waters. The total amount of time that would be required to complete this project on a single server would be 109 years. On Blue Waters, we were able to run a single workflow on the entire set of 10,000 AD samples by parallelizing across thousands of nodes. Further, we integrated results across all runs using the cleaned BAMs produced by multiple aligners, because Blue Waters is one of the very few systems that allows users to keep hundreds of terabytes of data in active storage for simultaneous processing.

PUBLICATIONS & DATA SETS

Ren, Y., et al., Identification of missing variants by combining multiple analytic pipelines. *BMC* 19:139 (2018), DOI:10.1186/s12859-018-2151-0.

QUANTUM–CLASSICAL PATH INTEGRAL SIMULATION OF ELECTRONIC TRANSITIONS

Allocation: Illinois/75 Knh
PI: Nancy Makri¹

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 of the ultrafast ferrocene–ferrocenium electron transfer reaction in liquid hexane on Blue Waters yielded results of unprecedented accuracy and enabled the first quantitative demonstration of Gaussian behavior for a complex molecular solvent. Simulation of slow processes in sluggish solvents remains challenging however. The present phase of the work focuses on the blip decomposition of the QCPI code, which leads to a dramatic acceleration of the calculation. In addition, current work incorporates an efficient classical trajectory-based method for including zero-point energy effects in the solvent.

RESEARCH CHALLENGE

Classical molecular dynamics methods are inadequate for describing charge transfer and (more generally) electronic transitions. On the other hand, quantum mechanical simulations of dynamical processes in the condensed phase continue to be extremely challenging because quantum mechanics is a nonlocal theory requiring computational cost that scales exponentially 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 degrees of freedom (e.g., the coordinates of a proton or a small number of electronic states). 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

This work focuses on the further development and implementation of a rigorous quantum–classical methodology

based on Feynman's path integral formulation of quantum mechanics [1]. The quantum–classical path integral (QCPI) methodology [2–6] takes advantage of the local nature of the Feynman paths to treat the interaction of a quantum mechanical subsystem with a classical environment without any *ad hoc* assumptions.

The QCPI expression appears impractical, however, because it contains an astronomical number of terms. Several advances in the understanding of interference and decoherence [7] 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 applied 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.

Current work implements a new acceleration of the code, which is based on the blip decomposition of the path integral [8–10]. This decomposition eliminates the vast majority of terms required, leading to a dramatic reduction of CPU time. Further, the QCPI code is augmented to allow the treatment of zero-point energy in the degrees of freedom treated via classical trajectories. This is important for molecular solvents, which contain high-frequency vibrations, as well as for processes in crystalline solids, in nanoscale structures such as carbon nanotubes, or in quantum fluids. The inclusion of zero-point energy is achieved by quantizing the thermal-phase space distribution using classical trajectories [11] or a new path integral-based approach [12] that does not suffer from a severe “sign problem” and thus converges rapidly.

RESULTS & IMPACT

The QCPI methodology enables the simulation of charge-transfer reactions in solution with unprecedented accuracy [13]. The ability to perform all-atom calculations with potential interactions treated in full detail leads to results of unparalleled




Figure 1: Schematic illustration of a classical trajectory with initial conditions sampled from a path integral-quantized phase space density and hopping between the donor and acceptor potential surfaces corresponding to the two charge states of an electron transfer pair according to instantaneous quantum states.

precision. These calculations shed light on the complex interplay among molecular/solvent timescales, electronic couplings, and reorganization energy [14], 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 the Gaussian response model. They also demonstrate how the interference among quantum mechanical phases leads to decoherence, and reveal substantial quantum delocalization of the otherwise classical solvent as a result of its interaction with the quantum electron transfer pair [13].

WHY BLUE WATERS

The QCPI formulation is well suited to a decomposition based on multilevel 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 relatively short, it is possible to assign a single trajectory to each core within a given processor while maintaining computational efficiency. This multilevel 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 interprocessor communication should be much faster than if the information were more widely distributed.

BLOOD-ARTERY INTERACTION AND SIMULATION-BASED ANALYSIS OF AORTIC ANEURYSM

Allocation: Illinois/50 Knh

PI: Arif Masud¹

Collaborators: Soonpil Kang¹, Elizabeth R. Livingston¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Fluid-Structure Interaction (FSI) techniques in biofluid dynamics integrate blood-flow and tissue models in a unified manner and provide a tool for comprehensive analysis of arterial diseases. High-fidelity simulations can enhance our understanding of the underlying biomechanical processes involved in the progression of the disease and also help with cardiac surgical planning. In this work, we have advanced our numerical methods and coupled, finitely deforming, nonlinear viscoelastic tissue models for the artery with non-Newtonian models for blood. The method is applied to a patient-specific geometry with progressive abdominal aortic aneurysm. The algorithm and code are optimized for XE nodes of Blue Waters.

RESEARCH CHALLENGE

Aortic aneurysm is an arterial disease that involves abnormal bulging of the walls of the aortic vessels. Aneurysm is characterized by a low stiffness of the artery walls, which leads to local ballooning effects that give rise to a localized swirling flow of blood. The challenges in numerical simulations of this physiological phenomenon include the complexity of constitutive models of blood as well as the large motion of the confining arterial walls. Our research has focused on developing scalable algorithms for computational Fluid-Structure Interaction regarding the cardiovascular system.




Figure 1: Growth of the abdominal aortic aneurysm (colored by the deformation of the artery wall).

METHODS & CODES

The scope of the research ranges from building a patient-specific model to developing the FSI formulation, a scalable algorithm, and a computer code.

Developing a Patient-Specific Model from CT-Scan Images: We obtained Computerized Tomography (CT) scan images from our collaborative institution, the Carle Foundation. The patient-specific geometric model of the aortic and femoral arteries was constructed from the CT images. A computational grid was constructed using 10-node tetrahedral elements while accounting for the thickness of the arterial walls, with two tetrahedral elements through the thickness of the wall.

Computational Fluid-Structure Interaction: The coupled FIS problem was solved on XE nodes of Blue Waters. Blood was modeled as a non-Newtonian fluid, while the deforming artery was modeled as a hyperelastic solid with fiber reinforcement in the circumferential direction. A key component in FSI methods is embedding of the kinematic and kinetic coupling between fluid and solid to the framework. The FSI techniques we developed enable the simulation of swelling behavior of the artery walls, or aneurysm, due to the blood pressure.

Parallel Implementation: The code was parallelized for distributed memory computations, adapting the ghost node approach and using the standard Message-Passing Interface (MPI). Various linear system solvers available in the PETSc [1] libraries were employed for the iterative solution in the Newton method.

RESULTS & IMPACT

We have simulated the evolution of the aneurysm, as shown in Fig. 1. The first few cycles correspond to the progression of the disease, which occurs over months or years. This is shown via the second and third image in Fig. 2, which show instantaneous snapshots of the system. The last three cycles correspond to flow through the progressed aneurysm, and the last image in Fig. 2 presents an instantaneous snapshot of the deformed geometry at the peak of systole. We have attempted to show via this test case that numerical simulations can be employed for accelerated modeling of disease progression, and at the same time they can be used to decipher the physiological features of the advanced stage of the disease. The computed results in Fig. 2 show the instantaneous snapshot of the velocity field of blood projected on the instantaneous displacement fields of the blood vessel. Also




Figure 2: Numerical results of the blood-artery interaction: (a) velocity streamlines and viscosity contour in blood projected onto an instantaneous snapshot of deformation of the artery wall; (b) wall shear stress on the inner surface of the wall that constitutes the interface between blood and artery.

projected is the isosurface of blood viscosity, which is an internal parameter in these calculations.

This test case enables us to identify the high-viscosity regions where blood coagulation can potentially take place. We have also projected the arterial Wall Shear Stress (WSS) on the arterial walls. It is important to note that WSS is one of the most significant factors for the progression of arterial disease. We further want to note that WSS data is very difficult to obtain via *in vivo* experiments.

To summarize, we have shown that advanced FSI simulations that can provide important data for the diagnosis and treatment of arterial diseases (such as an aneurysm) are critical in planning and developing strategies to cure these diseases.

WHY BLUE WATERS

The algorithms that we have developed for coupled FSI analysis are amenable to efficient parallelization because major portions of the computations are carried out at the local-element level. These algorithms are well-suited for distributed memory parallelism, and Blue Waters provides an ideal platform for implementation of such algorithms.

PUBLICATIONS & DATA SETS

Kwack, J., and A. Masud, A Three-field Formulation for Incompressible Viscoelastic Fluids. *International Journal of Engineering Science*, 48 (2010), pp. 1413–1432.

Masud, A., and R. Calderer, A Variational Multiscale Method for Incompressible Turbulent Flows: Bubble functions and Fine Scale Fields. *Computer Methods in Applied Mechanics and Engineering*, 200 (2011), pp. 2577–2593.

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.

Hussain, M., et al., A Parallel Implementation of ALE Moving Mesh Technique for FSI Problems using OpenMP. *International Journal of Parallel Programming*, 39 (2011), pp. 717–745.

Masud, A., and T. Truster, A Framework for Residual-Based Stabilization of Incompressible Finite Elasticity: Stabilized Formulations and F-bar Methods for Linear Triangles and Tetrahedra. *Computer Methods in Applied Mechanics and Engineering*, 267 (2013), pp. 359–399.

Anand, M., J. Kwack, and A. Masud, A new generalized Oldroyd-B model for blood flow in complex geometries. *International Journal of Engineering Science*, 72 (2013), pp. 78–88.

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 (2014), pp. 751–776.

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 (2017), pp. 704–734.

UNVEILING THE FUNCTIONS OF VIRUS CAPSIDS THROUGH THE COMPUTATIONAL MICROSCOPE

Allocation: Illinois/582 Knh

PI: Juan R. Perilla¹

Co-PI: Jodi A. Hadden¹

Collaborators: Adam Zlotnick², Louis M. Mansky³, Peijun Zhang⁴, Tatyana Polenova¹, Angela M. Gronenborn⁵, Christopher R. Aiken⁶, Tuli Mukhopadhyay²

¹University of Delaware

²Indiana University

³University of Minnesota

⁴University of Oxford

⁵University of Pittsburgh School of Medicine

⁶Vanderbilt University School of Medicine

EXECUTIVE SUMMARY

Virus capsids play critical roles in viral infection and are, therefore, promising molecular targets for antiviral drugs. Blue Waters enables all-atom modeling and molecular dynamics (MD) simulations of these remarkable molecular machines, exposing details underlying their biological functions and potential vulnerabilities to drug compounds. This project specifically studied the viruses that afflict humans with hepatitis B liver disease and T-cell leukemia.

RESEARCH CHALLENGE

Viral pathogens are a major risk to public health; millions of people die annually due to a lack of effective antiviral 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, 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]. Our simulations employed NAMD [3], 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

HBV. Hepatitis B virus (HBV, Fig. 1) is a leading cause of liver disease worldwide, including cancer. The capsid of HBV is icosahedral. To reduce computational expense and increase resolution in experimental structure determination, studies of icosahedral capsids commonly constrain them to be symmetric. We previously utilized Blue Waters to perform the first simulations of the HBV capsid without symmetry constraints [2]. At present,

we have leveraged Blue Waters to investigate the unconstrained capsid on the microsecond timescale, representing the most extensive unbiased simulations for the largest icosahedral virus capsid achieved to date.

Our simulations reveal that the HBV capsid is highly flexible and capable of asymmetric distortion, even under equilibrium conditions. Experiments suggest that the capsid's ability to flex and distort may be important during reverse transcription or nuclear import. Our simulations also reveal that triangular pores in the capsid surface selectively filter the passage of ionic species. This observation implicates the triangular pores as the openings through which the capsid exposes charged peptides that signal nuclear localization; this observation further suggests that the capsid controls signal exposure based on experimentally determined changes in peptide charge. Drugs designed to rigidify the capsid to prevent distortion, or to block triangular pores to prevent signal exposure, could inhibit key stages of viral infection.

We also used capsid conformers sampled during simulations on Blue Waters to perform a theoretical single-particle image reconstruction, mimicking structure determination by cryo-electron microscopy (EM). Our analysis indicated that capsid flexibility and asymmetry significantly lowered the resolution of the image reconstruction. Thus, even in this era of advancing cryo-EM technology, our simulations reveal that protein flexibility likely remains a major limiting factor to achieving true atomic resolution for virus capsids and other large biological structures.

HTLV-1. Human T-cell leukemia virus type 1 (HTLV-1, Fig. 2), which causes cancer, is a cousin of the HIV-1 virus, which causes AIDS. HTLV-1 assembles as an immature particle. In the immature state, the proteins destined to become the mature capsid are incorporated into a polyprotein known as Gag, which packs into a spherical lattice. Leveraging our previous experience using Blue Waters to determine the structure of the mature HIV-1 capsid [3] and immature Rous sarcoma virus capsid lattice [4], we have now utilized Blue Waters to construct a model of the immature HTLV-1 capsid lattice.

Mutagenesis experiments indicated a series of residues critical to HTLV-1 assembly, and our simulations reveal that these residues mediate key interactions at dimer and trimer interfaces between proteins within the immature capsid lattice. Our results also demonstrate that the HTLV-1 lattice is more structurally similar to that of HIV-1 than Rous sarcoma virus, although all three are retroviruses with similar capsid morphologies. Drugs designed to interfere with the key protein interactions identified by our study could disrupt the capsid structure or assembly of HTLV-1, representing a novel treatment approach.

WHY BLUE WATERS

Due to their formidable computational expense, simulations of virus capsids are only feasible on a petascale machine like Blue Waters. Capsid systems encompass millions of atoms, and computing the interactions among such large numbers of particles over microsecond timescales can take months, even on thousands




Figure 2: HTLV-1 assembles as an immature particle, with the polyprotein Gag packed into a spherical lattice with the viral genome at its center. We used Blue Waters to determine the structure of the capsid component of HTLV-1 Gag.

of processors. Further, analysis of the colossal data sets generated by capsid simulations is enabled 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 antiviral treatments, and demonstrate that access to leadership-class computing facilities holds the potential for significant impact on public health.

PUBLICATIONS & DATA SETS

Hadden, J.A., et al., All-atom molecular dynamics of the HBV capsid reveals insights into biological function and cryo-EM resolution limits. *eLife*, 7 (2018), DOI:10.7554/eLife.32478.

Martin, J.L., et al., Critical role of the HTLV-1 capsid N-terminal domain for Gag–Gag interactions and virus particle assembly. *Journal of Virology*, accepted (2018), DOI:10.1128/JVI.00333-18.

Wang, M., Quenching protein dynamics interferes with HIV capsid maturation. *Nature Communications*, 8:1 (2017), DOI:10.1038/s41467-017-01856-y.

Rayaprolu, V., et al., Length of encapsidated cargo impacts stability and structure of in vitro assembled alphavirus core-like particles. *Journal of Physics: Condensed Matter*, 29:48 (2017), DOI:10.1088/1361-648X/aa90d0.

Alvarez, F.J.D., et al., CryoEM structure of MxB reveals a novel oligomerization interface critical for HIV restriction. *Science Advances*, 3:9 (2017), DOI:10.1126/sciadv.1701264.

Hadden, J.A., and J.R. Perilla, Chapter 13: Molecular Dynamics Simulations of Protein–Drug Complexes: A Computational Protocol for Investigating the Interactions of Small-Molecule Therapeutics with Biological Targets and Biosensors. In *Computational Drug Discovery and Design—Methods in Molecular Biology*, 1762 (2018), DOI:10.1007/978-1-4939-7756-7_13.




Figure 1: The HBV capsid is an icosahedral protein shell composed of 240 identical proteins. Although experimentally determined structures suggest that the capsid is rigid and symmetric, our simulations on Blue Waters reveal that the capsid is highly flexible and can distort asymmetrically.

THE FREE ENERGY LANDSCAPES GOVERNING THE FUNCTION OF COMPLEX BIOMOLECULAR MACHINES

Allocation: NSF PRAC/3,720 Knh
PI: Benoit Roux¹

¹University of Chicago

EXECUTIVE SUMMARY

Proteins such as ion channels, transporters, and pumps play an essential role in controlling the bidirectional flow of material and information across the cell membrane, enabling the cell to accomplish complex tasks. In the present research, we investigated the selective ion-binding processes at the heart of the molecular mechanism of an ATP-driven ionic pump: the Na⁺/K⁺-ATPase. (ATP—adenosine triphosphate—is the immediate source of energy for most cellular work, in this case the transport of sodium and potassium ions across a membrane.) We also studied the key factors controlling the stability of the C-type inactivated state of the KcsA K⁺ channel, a nonconducting state of K⁺ channels with great physiological implications. Our project aims to gain a deep mechanistic perspective of protein function, linking structure to dynamics, by characterizing the multidimensional free energy “landscapes” that govern key functional aspects of strong electrostatic interactions associated with either ion binding or charged moieties. A free energy landscape presents a powerful and rich concept to help understand key functional processes in these systems.

RESEARCH CHALLENGE

We study complex biomolecular assemblies that consume energy in order to perform specific biological functions. The concerted action of these “molecular machines” underlies all activities of the living cell. The proteins associated with biological membranes are particularly remarkable in their control of the two-way flow of material and information across the membrane. An important question is whether one can truly design an effective and scalable computational strategy that is well adapted for a leadership-class computer such as Blue Waters to tackle processes that occur over long timescales. Unbiased molecular dynamics (MD) trajectories, while very valuable, can be limited. Advanced free energy methodologies can help overcome these limitations [1–4].

The solution that we proposed to this problem is to break down the task by calculating the free energy landscape that governs the key functional motions within a subspace of predetermined order parameters. A free energy landscape (or potential of mean force) presents a powerful and unifying concept to aid in understanding these systems. By studying a small set of biologically significant

but experimentally well-characterized systems of increasing size and complexity within a unified computational perspective provided by free energy landscapes, our overarching goal is to push the envelope and advance the theory–modeling–simulation technology and open this virtual route to address fundamental biological questions. It is our hope that the present study will serve as a road map for elucidating key thermodynamic features affecting the function of important molecular nanomachines by using an extremely scalable computational strategy.

METHODS & CODES

The solution that we proposed to overcome the timescale limitation is to break down the task by calculating the free energy landscape that governs the key functional motions within a subspace of predetermined order parameters using umbrella sampling (US) Hamiltonian-tempering replica exchange MD (US/H-REMD) simulations [3,4]. Based on the knowledge of a specific system, it is possible to manage efficient US/H-REMD simulation in any multidimensional subspace of collective variables [3,5,6]. US/H-REMD simulations are extremely scalable to thousands of CPUs on high-performance computers. These computations employ NAMD extended to treat multiple-copy algorithms. (NAMD is a parallel molecular dynamics code designed for high-performance simulation of large biomolecular systems.)

RESULTS & IMPACT

C-type inactivation in K⁺ channels is thought to be due to constriction of the selectivity filter. Our recent study on KcsA showed that rapid constriction occurs within 1–2 microseconds when the intracellular activation gate is fully open, but not when the gate is closed or partially open. These results imply that the observed kinetics underlying activation/inactivation gating reflect a rapid conductive-to-constricted transition of the selectivity filter that is allosterically controlled by the slow opening of the intracellular gate.

Our simulations of the SERCA (sarcoplasmic reticulum Ca²⁺-ATPase) pump using different protonation configurations at the binding sites reveal how deprotonation events affect the opening of the cytoplasmic gate. The results show that there is a strong coupling among the chronological order of deprotonation, the entry of water molecules into the transmembrane (TM) region, and the opening of the cytoplasmic gate. From a functional point of view, the chronology of the deprotonation process described here helps explain how the transporter maintains strict coupling at physiological pH by controlling solvent accessibility to the binding sites.

The computational framework we developed is likely to impact a vast array of problems in structural biology. Our goal is to advance and expand the fundamental knowledge in the theoretical and computational methodologies that are used to characterize complex macromolecular biological systems. Technological advances in computational biophysics are expected to have a significant long-term fundamental impact on human




Figure 2: Main conformational changes occurring in the transmembrane (TM) region and cytoplasmic domains during the transition between major functional states of the SERCA pump. The figure shows the dominant motions in the TM (top left), the P domain (top right), the A domain (bottom left), and the N domain (bottom right).

health by furthering the development of reliable technologies for screening large biomolecular systems to databases of compounds and identifying potential lead drug molecules. All the research is done in an academic environment where education and training of students and postdocs is an intrinsic component of our activities.

WHY BLUE WATERS

Umbrella sampling Hamiltonian replica-exchange molecular dynamics simulations are extremely scalable to thousands of CPUs on high-performance computers; multidimensional calculations on systems that range in size from 66,104 atoms to 291,148 atoms require thousands of nodes. Thus, these computations are at the forefront of what is possible now thanks to petascale leadership computers such as Blue Waters.

PUBLICATIONS & DATA SETS

Li, J., et al., Chemical substitutions in the selectivity filter of potassium channels do not rule out constricted-like conformations for C-type inactivation. *Proceedings of the National Academy of Sciences U.S.A.*, 114:42 (2017), pp. 11145–11150.

Li, J., et al., Rapid constriction of the selectivity filter underlies C-type inactivation in the KcsA potassium channel. *Journal of General Physiology*, in press (2018).

Rui, H., Proton countertransport and coupled gating in the sarcoplasmic reticulum calcium pump. *Journal of Molecular Biology*, in review (2018).

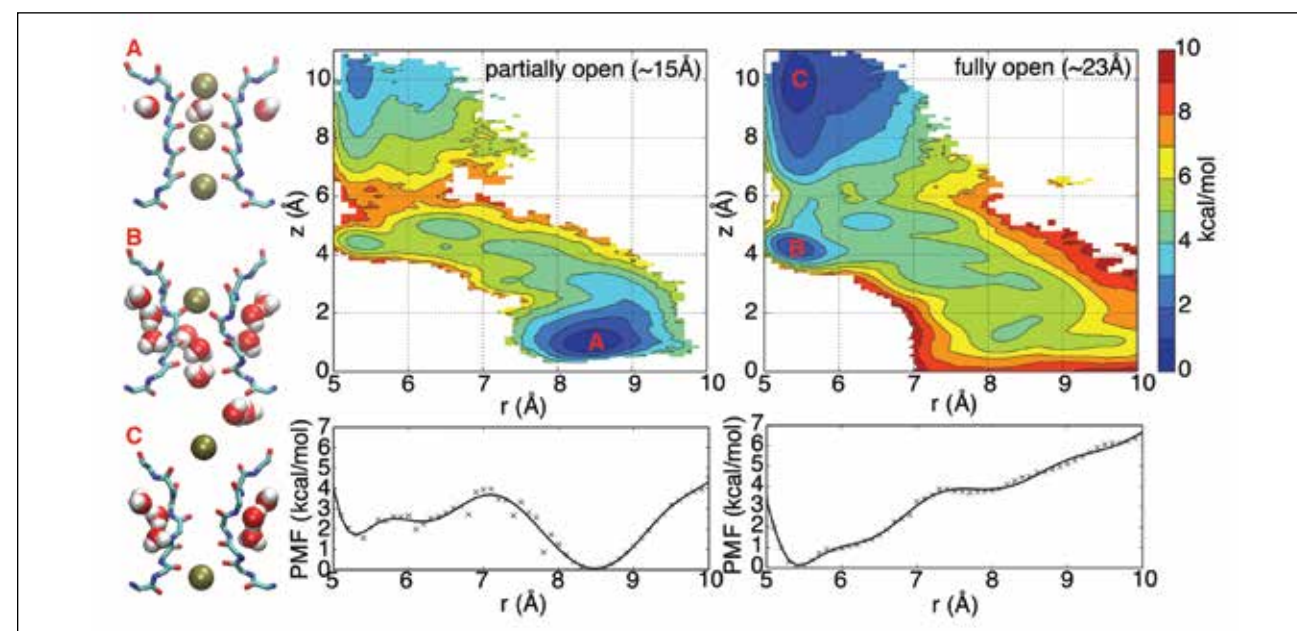


Figure 1: Two-dimensional PMF (potential of mean force) to assess the conformational preferences of the selectivity filter with partially and fully open intracellular gate in K⁺ channel KcsA. The horizontal and vertical reaction coordinates, respectively, describe the width of the selectivity filter (r) and the position of the external K⁺ ion along Z axis (z).

UNRAVELING THE MOLECULAR MAGIC OF WITCHWEED

Allocation: Illinois/300 Knh
PI: Diwakar Shukla¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Parasitic weeds of the genus *Striga*, commonly called witchweed, are considered the most damaging agricultural agents in the developing world. An essential step in parasitic seed germination is the sensing of a family of plant hormones called strigolactones, which are released by the host plants. Despite the economic importance of strigolactones, little fundamental information is known about this plant hormone. The recently obtained crystal structures of strigolactone receptors provided a unique opportunity to explore the functional mechanism of strigolactone signaling in plants. Using computational time on Blue Waters, we investigated the *apo* and *holo* strigolactone receptors in the *Arabidopsis thaliana* plant through molecular dynamics simulations. We were able to identify multiple intermediate states as well as significant differences in the activation mechanisms of *apo* and *holo* receptors. Our findings pave the way toward molecular design for the chemical control of witchweed infestations.

RESEARCH CHALLENGE

Witchweed is a root parasite that is considered a serious agricultural pest affecting crops such as sorghum, maize (corn), rice, millet and cowpea, among other crops. Witchweed seeds remain in the soil for decades until favorable germination conditions are

provided by a host plant. These seeds sense host-exuded stimulant molecules—a plant hormone called strigolactones (SLs)—and start growing rapidly by attaching themselves to the root of the host plant and competing for nutrients with the crops. The parasite continues to grow beneath the soil undetected. The emergence of flowering shoots occurs months later, by which time the crop has been completely destroyed, leading to huge economic damage. Currently, there is a need for a control technique to combat the outbreak of this menacing parasite. In *Arabidopsis thaliana*, strigolactones are identified by the receptor protein D14. Upon identification/binding to the receptor, the SLs undergo hydrolysis.

Several experimental groups have reported potential activation mechanisms for the D14 protein. The key challenge involves assessing the effect of the ligand produced after the hydrolysis reaction on the activation process of the receptor protein. The available crystal structures [1,2] provide only some of the snapshots of the stable protein conformations that will enable the community to understand the conformational change from the inactive (open) to the active (closed) state (Fig. 1). To enable the discovery of the various intermediates involved in the activation process, we performed extensive molecular dynamics (MD) simulations of the *apo* and covalently linked intermediate molecule (CLIM) *holo* AtD14 starting from the available crystal structures, for




Figure 1: Conformational change in the strigolactone receptor D14 leading to an open and closed state.

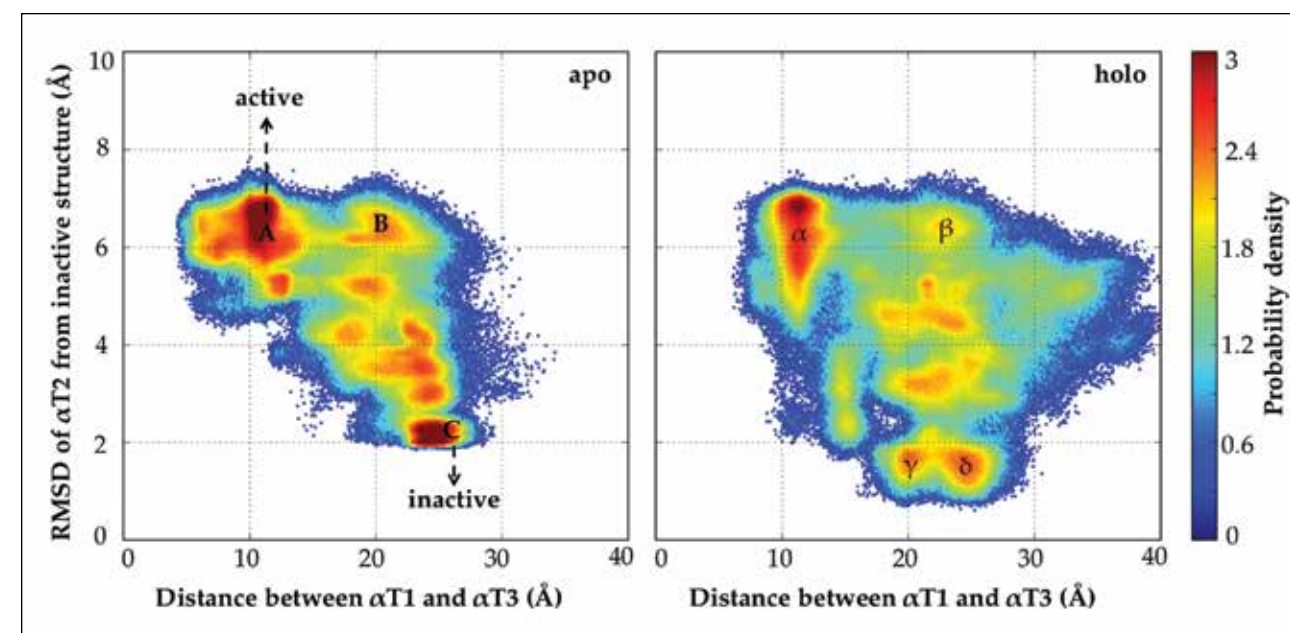


Figure 2: Probability density maps of *apo* and *holo* D14 strigolactone receptor based on simulation data.

approximately 190 and 120 microseconds, respectively. The conformational changes associated with the unfolding of helix $\alpha T2$ and the closing in of $\alpha T1$ toward $\alpha T3$ are expected to occur over long timescales. Therefore, large amounts of unbiased simulation times have to be accessed in order to characterize the activation mechanism.

METHODS & CODES

The MD systems were neutralized by adding a 0.15M concentration of sodium and chloride ions. All MD systems were subjected to minimization and equilibration for 2 nanoseconds on CPUs. The production runs were performed on one Blue Waters GPU node at 300 K and 1 bar. To find more informative starting points for the second round of simulations that can speed up the exploration of the conformational landscape, we used the adaptive sampling approach in the study.

RESULTS & IMPACT

The lack of molecular-level insights regarding SLs and SL-induced activation of the D14 protein, as well as the lack of any computational studies, has prevented the development of agrochemicals for controlling the spread and growth of this parasite. However, these fundamental questions can be answered using atomistic MD simulations. The present MD simulation study on Blue Waters is necessary for developing a chemical control mechanism for the witchweed parasite, and to prevent a billion dollars' worth of annual agricultural losses.

The simulation results mapped into a two-dimensional conformational landscape of the $\alpha T1$ – $\alpha T3$ distance versus the root mean squared deviation of the $\alpha T2$ helix with respect to the

inactive structure (PDB ID: 4IH41) show that the closed (active) state is more accessible in *holo*; hence, multiple activation pathways may be possible. We found that the *apo* protein can also exhibit active-like conformations (region A in Fig. 2) for which no previous crystal structure is known. However, the active state is extremely difficult to escape in *apo*, whereas in *holo*, there are multiple minima (regions indicated by γ and δ in Fig. 2). We used the simulation data to build Markov state models of conformational dynamics by clustering the structurally similar conformations into states and obtaining the interconversion rates among these states from the simulated trajectories.

Our future work will focus on the activation mechanism of the SL receptor in the presence of multiple hydrolysis reaction products that have been reported in the literature. Since the conformational change study will require microseconds-long simulations, we will use Blue Waters to run thousands of MD simulations of the receptor protein. Our goal is to characterize the behavior of the protein at the molecular level and pinpoint the difference in activation in the host and the parasitic plant species' protein.

WHY BLUE WATERS

Understanding the slow conformational transitions in proteins requires hundreds of microsecond-long simulations. Blue Waters provides the state-of-the-art computer architecture needed to perform such studies. We employ large-scale adaptive sampling protocols, which can be efficiently performed on Blue Waters' GPU and CPU framework. The current work would not be possible without Blue Waters.

MOLECULAR SIMULATIONS ON BLUE WATERS AID THE UNDERSTANDING OF HOW PLANTS TRANSPORT SUGARS BETWEEN CELLS

Allocation: Illinois/600 Knh
PI: Diwakar Shukla¹
Collaborator: Li-Qing Chen¹

¹University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

SWEETs are a new family of sugar transporter proteins in plants that play a crucial role in various fundamental processes such as nectar production, pollen development, and plant-microbe interaction. SWEETs function is based on a rocker-switch mechanism; that is, an outward-facing (OF) to inward-facing (IF) transition to transport substrate molecules across the cell membrane. The OsSWEET2b crystal structure in rice was recently obtained in the IF state and provided the first glimpse of structural information on this class of proteins. However, these transport proteins are very flexible in nature and it is difficult to understand the structural changes based on a single, static X-ray crystal structure. In this study, we performed all-atom unbiased molecular dynamics (MD) simulations to investigate the conformational dynamics of the OsSWEET2b transporter. For the first time, we characterized the complete sugar-transport cycle of a plant transporter, and determined the critical residues that mediate the transport of sugar.

RESEARCH CHALLENGE

Global climate change and increasing world population pose a great threat to the current agricultural economy. Although genetic engineering has emerged as a successful approach that enhances crop productivity under optimum environmental conditions, it still fails to meet the current food demand. An alternative solution to increase crop yield may be to change the phenotype of plants. (The phenotype is the set of observable characteristics of an individual plant resulting from the interaction of its genotype with the environment.) Two years ago, researchers produced the largest pumpkin, setting a world record of 1,190 kg [1]. This

experimental study suggested that these plants store sugar in the phloem (the vascular tissue of plants) and transport it from the leaves to various organs [2]. The sugars (glucose) are produced in the leaf during photosynthesis and transported to the phloem via the SWEETs' sugar transporters [3].

In our research study, we performed hundreds of microsecond-long simulations to understand the functional dynamics of SWEETs and characterized how the sugars are exported and transported to the cell. Our study provides novel insights into the molecular mechanism of sugar transport.

METHODS & CODES

MD simulations rely on numerical integration of Newton's equation of motion for the interacting atoms. This results in time-dependent trajectories for all atoms of the system, which together provide a simulation of the biomolecule's dynamical motion. Our simulations were performed in AMBERv14 [4]. AMBER is highly parallelized to massively accelerate complex molecular simulations and enhance sampling efficiency. Our simulations generated several terabytes of data that cannot be analyzed visually. CPPTRAJ [5] is the processing and analytical tool written in C++ that is available with the AMBER suite for trajectory processing. In addition, we used Python modules such as PyTraj [6] and MDTraj [7] for data analysis and processing. The MD data we obtained were featurized to biologically relevant reaction coordinates and converted to Python array NumPy files for efficient processing and analysis. We used Markov State Models (MSM) to cluster the data based on relevant kinetics [8]. From the clusters, MSM constructed a transition probability matrix to find the rate of transition from one state to another. Using transition path theory

Figure 1: The numbers 1–6 in Fig. 1A and 1B represent the inward-open (1A, 5B); occluded (2A, 4B); intermediate states (3A, 4A, 6A, 2B, and 3B); and outward-open (5A, 1B). The pore channel opening and closure at the periplasmic and cytoplasmic side are obtained by measuring the distances between the gating residues. The crystal structure is shown as a black star. The dominant paths 1A (grey line), 2A (dotted black line), 3A (thin black line), 1B (dotted grey line), and 2B (black line) are shown as arrows.


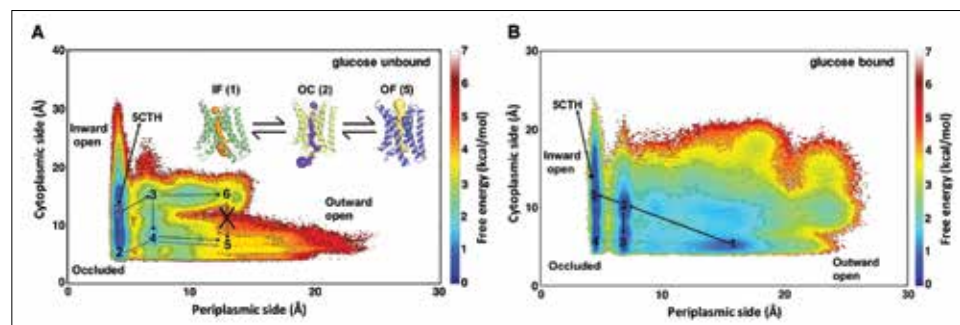


Figure 2: The glucose-binding mechanism and the glucose-driven conformational changes are shown in Fig. 2A to 2H. The distances between gating residues [Arg70 (CZ)–Asp190 (CG) and Phe165 (Ca)–Phe43 (Ca)] are shown in magenta. The intracellular hydrophobic gates Phe24–Met146–Met177 are shown in cyan. The glucose molecule is shown in green.

(TPT), the intermediate states between the source and sink states were identified [9]. The structures were extracted and visualized in VMD [10] and Pymol [11].

RESULTS & IMPACT

Conformational dynamics of OsSWEET2b: The IF state OsSWEET2b (PDB ID: 5CTH) crystal structure was used as a starting structure for MD simulations [12]. We obtained the complete transition from IF to occluded (OC) and OC to OF over a period of ~145 microseconds (μ s) (Fig. 1A). The simulation data in high-dimension space were transformed to the three slowest processes and clustered to kinetically relevant states. MSM were constructed, with the final model containing 900 microstates. MSM-weighted free-energy plots were obtained by projecting the data on extracellular and intracellular gating distances (Fig. 1A and 1B).

The free-energy barrier for one complete cycle of an *apo* (ground state) transporter from IF to OC and OC to OF was estimated to be ~4 kcal/mol. Using TPT, the dominant pathways of transition were determined and Path-1A was found to be the lowest-energy pathway between IF and OF states. Path-2A was identified as an alternative high-energy pathway. The transition along Path-3A was not feasible as the transporters are wide open at both ends.

We extended the simulation from the OF state to investigate the glucose recognition, binding, and transport mechanism over a duration of ~68 μ s (Fig. 1B). The conformational landscape plots show that glucose decreases the barrier among various states, and the transition between them is easily accessible compared to *apo* dynamics. We obtained two major pathways, namely Path-1B and Path-2B, using TPT. The free-energy barrier was 2–3 kcal/mol for transition of glucose from OF to IF via OC or extended OC–OF states.

Glucose translocation in OsSWEET2b: Glucose is recognized in the OF state by residues such as Arg70, Arg189, and Asp190 at the extracellular surface (Fig. 2). Glucose then diffuses to the pore channel and establishes stable contacts with binding-site residues Asn57, Asn77, and Asn197, which drive the transporter to the OC

state. At this juncture, the intracellular hydrophobic gates Phe24–Met146–Met177 act as a barrier and restrict further movement of glucose. The increase in strength of the ionic contacts at the extracellular surface results in the displacement of hydrophobic gating residues; thus, it results in the entry to the intracellular side of the transporter. The downward movement of glucose leads to the IF state and forms polar contacts with Gln84, Tyr48, and Asp44.

Our study reveals atomistic-level detail of the conformational dynamics of OsSWEET2b and the free-energy barrier associated with the transport process. In this work, we have shown how the glucose regulates the SWEET function and key residues involved in the conformational switches. As a continuation of this work, we will focus on reducing the energy barrier of conformational switches between the intermediate states, thus allowing for more sugar in a plant's phloem in order to increase fruit size and crop yield.

WHY BLUE WATERS

Our work is computationally demanding and requires multiple nodes to run large numbers of simulations. Parallel computing reduces the cost and time effectively. Historically, parallel computing has been the “high end of computing,” and has been used to model difficult problems in many areas of science and engineering. Blue Waters is a perfect resource that has the massive architecture to run parallel jobs. The timescales for simulation of these biological problems range from several microseconds to milliseconds. These computations are not possible to perform in a reasonable time without Blue Waters' petascale computing capability.

PUBLICATIONS & DATA SETS

Selvam, B., Y.-C. Yu, L.-Q. Chen, and D. Shukla, Molecular Basis of Glucose Transport Mechanism in Plants. In review (2018).

Cheng, K.J., B. Selvam, Y.-C. Yu, L.-Q. Chen, and D. Shukla, SWEET vs SemiSWEET: A Biophysical Investigation of their Similarities and Differences. In review (2018).

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

Allocation: NSF PRAC/707.6 Knh

PI: Ivan Soltesz¹

Collaborators: Aaron Milstein¹, Grace Ng¹, Darian Hadjiabadi¹

¹Stanford University

EXECUTIVE SUMMARY

Our work is the first attempt to use a detailed computational model for each neuron in the hippocampus, a brain area important for learning and memory, to fully understand how memories are formed in the brain. The computational capacity of Blue Waters allows us to rapidly conduct simulations of brain function, to observe and record the behavior of millions of neuron models, and to compare the results with experimental data. We have developed methodological advancements to our models to allow unprecedented biological realism and increased computational efficiency, including software tools to manage large data sets on supercomputers and a novel method for model optimization, the process used to fit mathematical models of neurons to experimental data.

RESEARCH CHALLENGE

Our research aims to elucidate the mechanisms of sharp-wave ripples (SWRs), which are oscillatory events in the hippocampus that are required for memory consolidation and subsequent recall. To support this goal, our computational projects aim to construct full-scale, biophysically detailed computational models of the three major neuronal circuits within the mammalian hippocampus: the dentate gyrus (DG), CA3, and CA1. These models will be used to provide insight into the dynamical properties of hippocampal networks that produce the SWR-specific oscillatory patterns.

Furthermore, we propose to utilize our full-scale models to study the mechanisms of abnormal dynamics that emerge in epilepsy or through radiation damage.

METHODS & CODES

The principal simulation environment we use is NEURON [1], which describes neurons in terms of membrane properties and geometric structure [2], and networks in terms of connections between neurons [3]. The biophysical dynamics of the neuronal membrane are described by differential-algebraic equations solved by an implicit integrator optimized for branched structures [1]. NEURON can be fully parallelized via MPI with near-linear scaling [3].

The representation of the geometric structure of neurons and their connectivity requires hundreds of gigabytes for each of our models, which has necessitated a parallel computational infrastructure for data management. Thanks to the Petascale Application Improvement Discovery program, we have developed a parallel I/O software substrate based on the HDF5 file format that allows the rapid generation and analysis of neuronal morphology and connectivity data according to user-specified rules about neuronal structure and distribution of connectivity in a 3D volume.

An important step in the development of biological neuronal network models is optimizing the neuronal cell models to fit experimentally observed biophysics. The corresponding objective

functions are highly nonconvex, resulting in poor performance by gradient-based approaches. We have developed a new optimization algorithm, “population annealing,” that combines aspects of simulated annealing and genetic algorithms to achieve multiobjective optimization in parallel with fewer overall function evaluations (roughly 150 iterations with around 200 models per generation as compared to more than 500 iterations in previous genetic algorithm-based approaches).

RESULTS & IMPACT

We have recently completed the first of its kind, full-scale, anatomically and physiologically constrained model of the CA1 subregion of the hippocampus [4], which contains 300,000 pyramidal cells and 30,000 inhibitory interneurons. We are in the process of tuning the physiological properties of our full-scale model of the dentate gyrus (DG) subregion that is likewise anatomically and physiologically constrained. The DG model contains one million granule cells with unique generated realistic dendritic morphologies depending on its position in the DG volume [5], 30,000 mossy cells, and approximately 20,000 inhibitory interneurons. Simulations of the DG network model thus far can reproduce an important phenomenon in the hippocampus—the formation of “place cells,” which are neurons that encode the physical location of an animal during spatial navigation [6] (Fig. 1).

A necessary major technical innovation has been the development of synapse models that can capture the characteristic process of activity-dependent synaptic saturation, which prevents uncontrolled increase of synaptic conductance and excessive neural firing and plays a crucial role in information transfer between neurons. Previous models have implemented saturation by maintaining unique state variables for each connection, which results in excessive memory use. We have implemented an efficient scheme that tracks the activity of each connection, yet permits use of shared state variables between connections, which results in an approximately three-fold reduction of simulation runtime of our models on Blue Waters (Fig. 2).

Our approach to modeling neuronal networks with realistic cellular diversity and spatial topography provides a detailed understanding of how information propagates in the brain, and how the diverse oscillatory dynamics of brain networks emerge. Thus far, it has not been possible to investigate these parameters with reduced-scale models, where network dynamics are potentially strongly distorted [7]. Furthermore, particular types of high-frequency oscillations are the hallmark of focal epilepsy and are hypothesized to have a causal role in the initiation of seizures, yet these oscillatory patterns have proven difficult to replicate in prior computational models. Altogether, our software infrastructure for large-scale neuronal modeling offers a pioneering methodology for rapidly developing, validating, and measuring the information-processing capabilities of realistic, biophysically detailed networks of millions of neurons.

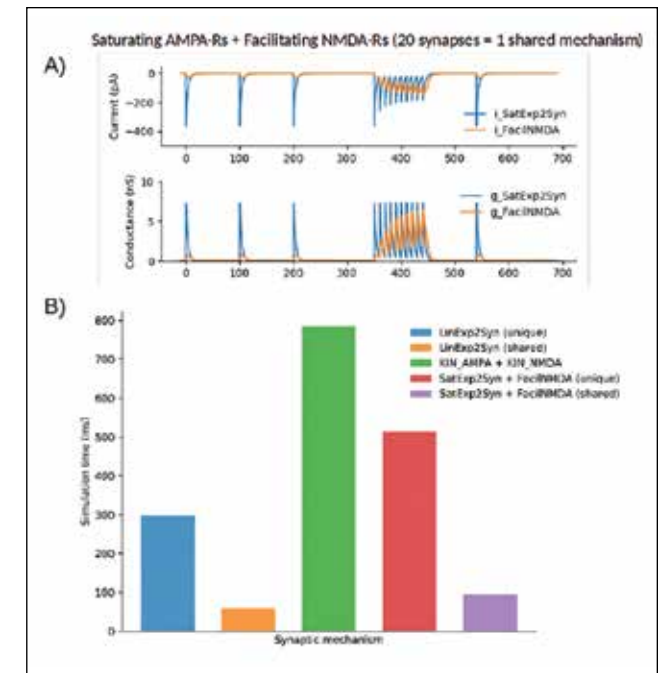


Figure 2: Synaptic mechanisms for full-scale simulations. (A) Biophysically realistic postsynaptic responses of two major types of synapses, a saturating AMPA and facilitating NMDA receptor. (B) Simulation performance: simplified exponential rise and decay LinExp2Syn; physiologically detailed KIN_AMPA and KIN_NMDA; saturating and facilitating mechanisms optimized for connection sharing SatExp2Syn and FacilNMDA.

WHY BLUE WATERS


Our research requires the simulation of behaviorally relevant brain activity on the scale of tens of seconds. In our first year of using Blue Waters, 10-second simulations of our models took 14 to 19 hours to run on 1,024 Blue Waters CPU nodes. Thanks to the PAID program and the technical expertise of Blue Waters’ staff, we have reduced the runtime to three or six hours on 1,024 or 2,048 nodes, respectively. Only Blue Waters can permit longer simulations of the eventual combined hippocampal model. The important discoveries revealed by our simulations underscore the essential need for public petascale computing resources such as Blue Waters.

PUBLICATIONS & DATA SETS

Raikov, I., A. Milstein, I. Soltesz, Determinants of sparse population coding in a computational model of the rat dentate gyrus. *Society for Neuroscience*, Washington, D.C., 2017.

Milstein, A., et al., Place field translocation by bidirectional behavioral time-scale synaptic plasticity. *Computational and Systems Neuroscience (Cosyne)*, Denver, Colo., 2018.

Figure 1: (A) Firing rate map of granule cells, ordered by peak firing rate. (B) Rhythmic population activity of the major cell types in the network. (C) Consistent with experimental observations, the granule cells in the model have a mean 1 place field and a log-normal distribution of number of place fields.



SIMULATION OF VIRAL INFECTION PROPAGATION THROUGH AIR TRAVEL

Allocation: NSF PRAC/350 Knh

PI: Ashok Srinivasan¹

Co-PIs: Sirish Namilae², Anuj Mubayi³, Matthew Scotch³, Robert Pahle³

Collaborator: Sudheer Chunduri⁴

¹Florida State University

²Embry–Riddle Aeronautical University

³Arizona State University

⁴Argonne National Laboratory

EXECUTIVE SUMMARY

There is direct evidence for the spread of common infectious diseases during commercial air travel, including influenza, SARS (Severe Acute Respiratory Syndrome), tuberculosis, and measles. This has motivated calls for restrictions on air travel, for example during the 2014 Ebola outbreak. However, such restrictions 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. Thus, science-based policy analysis can yield useful insight to decision-makers.

The effectiveness of a policy depends on the 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 cases of national emergencies. Our results identify new boarding procedures that can result in a substantial reduction in the risk of the spread of Ebola and SARS.




Figure 1: Schematic of the SPED fine-scale model combining pedestrian movement and infection dynamics.

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 model the movement of pedestrians during air travel as particles based on a force-field approach proposed by [1]. Both the pedestrian density and speed of the 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 model for infection transmission that accounts for demographic stochasticity and variations in the susceptibility of the population. This approach (Fig. 1) provides insight into the consequences of policy choices that change passenger behavior at individual levels. We input this information to a global phylogeography model to assess the impact of these policies at a global scale.

Inherent uncertainties in human behavior and insufficient data during the initial stages of an epidemic make prediction difficult. Our approach is to parameterize the sources of uncertainty and evaluate vulnerability under different possible scenarios. We use the Blue Waters supercomputer to efficiently deal with the computational load that arises from a large parameter-space, and a low discrepancy parameter sweep to explore the space of uncertainties.

Phylogeography uses genetic mutation information and the geographic locations of viruses to model the spread of epidemics across large geographic scales. We used Blue Waters to analyze 264 full-genome Ebola sequences from Guinea, Liberia, Sierra Leone, Italy, the United Kingdom, and the United States. We used the BEAST [5] software installed on Blue Waters to implement the phylogeography model.




Figure 2: Comparison of the convergence of histograms for human contacts with different types of parameter sweep. (Left) Lattice with 1.4 million points. (Center) Low-discrepancy with 32,000 points. (Right) Accurate results.

RESULTS & IMPACT

In prior work, we used the above approach with Ebola. We studied the impact of different procedures for boarding, disembarkation, and seat assignment on infection spread. 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. We have also 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 infections 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.

We have extended our approach to other directly transmitted diseases including SARS and influenza. This requires changes in approach that include aerosol and fomite transmission mechanisms, while pedestrian movement accounts for proximity among infectious and susceptible individuals. We have successfully extended the application to other high-density areas such as airport security-check areas and a generic airport gate. We also found effective strategies to mitigate spread by using temporary walls for queue management instead of ropes.

We developed a low-discrepancy parameter sweep that reduces by one to three orders of magnitude the number of parameter combinations that ought to be tried over the conventional lattice-based sweep (Fig. 2). We used number-theoretic properties of a low-discrepancy sequence to balance the load on the Blue Waters machine [4].

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 step. Consequently, rapid turn-around 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 a short

time. Toward that end, we obtained support from the Blue Waters team to optimize parallel I/O in our code to reduce simulation time by a factor of two.

PUBLICATIONS & 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* (2016).

Namilae, S., et al., Self-Propelled Pedestrian Dynamics Model: Application to Passenger Movement and Infection Propagation in Airplanes. *Physica A: Statistical Mechanics and its Applications*, 465 (2017), pp. 248–260.

Namilae, S., et al., Self-Propelled Pedestrian Dynamics Model for Studying Infectious Disease Propagation during Air-Travel. *Journal of Transport & Health*, 3:2 (2016) p. S40.

Namilae, S., et al., Multiscale Model for Infection Dynamics During Air Travel. *Physical review E*, 95:5 (2017), p. 052320.

Derjany, P., et al., Effect of Pedestrian movement on infection transmission during air travel: A modeling study. *Transportation Research Forum Proceedings* (2017).

Derjany, P., et al., Computational Model for Pedestrian Movement and Infectious Diseases Spread During Air Travel. *AIAA Modeling and Simulation Technologies Conference* (2018), p. 0419.

Namilae, S., A. Mubayi, and A. Srinivasan, Model based policy analysis for infection spread during air transportation. *International Conference of Transportation and Public Health* (2018).

Chunduri, S., et al., Parallel Low Discrepancy Parameter Sweep for Public Health Policy. *IEEE/ACM International Symposium on Cluster, Cloud, and Grid Computing* (2018).

HPC-BASED HIGH-RESOLUTION QUANTITATIVE MRI

Allocation: Illinois/100 Knh

PI: Brad Sutton¹

Collaborators: Curtis L. Johnson², Alex Cerjanic¹, Aaron T. Anderson¹

¹University of Illinois at Urbana-Champaign

²University of Delaware

EXECUTIVE SUMMARY

Newly emerging high-resolution, quantitative magnetic resonance imaging techniques are powerful tools for clinicians and researchers to apply to study problems in neurology, neurosurgery, and neuroscience. For example, magnetic resonance electrography enables mapping the structure, function, and health of the human brain by measuring the mechanical properties of neural tissue. However, such techniques are hampered by excessive image reconstruction and analysis times that limit their adoption into standard practice. Software developed for the unique Blue Waters architecture that leverages the combination of CPU and CPU-GPU resources allows for reconstructions and analysis at a scale commensurate with data collection at the Beckman Institute and Carle Neuroscience Institute. In a cohort of mesial temporal lobe epilepsy patients, magnetic resonance elastography (MRE) has shown it can identify affected regions as stiffer than similar regions in healthy controls. Since surgical resection is the most viable treatment, MRE has shown potential as a specific biomarker of disease.

RESEARCH CHALLENGE

Quantitative magnetic resonance imaging (MRI) techniques provide an added dimension to standard neuroimaging routines: they allow researchers and clinicians to measure the physical characteristics of the human brain. Over the last few years, our team has become a leader in MRE development and application,

enabling the study of the mechanical properties of the brain. Indeed, the research team has made significant advances in developing high-resolution methods for MRE based on the advanced image acquisition and reconstruction approaches [1–3] and robust mechanical inversion schemes [4–5]. Through these developments, we have been able to map the brain at unprecedented resolution and probe specific neuroanatomical structures, including white matter tracts [6,7] and subcortical gray matter [8].

Additionally, the pathogenesis of neurodegenerative diseases such as dementia and Alzheimer’s disease (AD) are not well understood, but there are indications of microvascular changes being an important biomarker. To understand the changes in the microvascular architecture that accompany AD, investigators have been limited to sophisticated histological and immunohistological methods to laboriously document changes in postmortem (after death) samples [8–11]. These histological techniques, while capable of visualizing the degradation of the microvasculature, one microscope field at a time, fundamentally lack the scale to be incorporated in prospective studies tracking individuals at risk of developing AD or other forms of dementia, or to incorporate information about functional measures of blood flow. In light of these limitations, the need for a quantitative set of noninvasive measures for brain microvascular architecture that can scale to population-level studies is clear.




Figure 1: (Top) MPRAGE, FLAIR, and MRE-Stiffness images from a patient with mesial temporal lobe sclerosis. (Bottom) The ratio of stiffness of epileptogenic to nonaffected hippocampus (HC) is significantly greater compared to controls. The ROC curve indicates that the ratio of HC stiffness dissociates patients from controls with high accuracy.

METHODS & CODES

We used two codes in this work: PowerGrid for iterative MRI image reconstruction [12] and nonlinear inversion for iterative MRE property reconstruction [4,5].

PowerGrid is an open-source platform developed by our group for performing iterative MRI image reconstruction [12], and is a flexible, GPU-based MRI reconstruction platform that includes support to model and correct for several common data artifacts. Implemented with initial assistance from NVIDIA and NCSA, PowerGrid has shown good scaling across multiple GPU nodes on Blue Waters via MPI for performing the reconstruction of a single subimage from a single subject.

Nonlinear Inversion (NLI) is an iterative, finite-element method (FEM)-based inversion algorithm for converting MRE shear displacement fields into shear modulus property maps [4,5]. NLI was originally developed by researchers at Dartmouth College and has been adapted for brain MRE in collaboration with the project research team. The current implementation of NLI divides the brain into smaller regions (called “subzones”) for FEM meshing and optimization of properties at a local level, which are then reconstituted as a global solution. Importantly, this subzoning process allows for parallelization across CPU cores.

RESULTS & IMPACT

The codes and computational resources have been applied to several neuroimaging studies, including a mesial temporal sclerosis project (Fig. 1). Our data suggest that MRE is a useful tool for identifying patients with mesial temporal lobe epilepsy due to an asymmetry in hippocampal scarring. This project has also identified an unanticipated and intriguing outcome: Previously, based on available imaging tools, researchers and physicians have believed that the nonsclerotic hippocampus was structurally normal. However, our work has revealed significant softening of the nonsclerotic hippocampus compared to healthy controls, which may provide an important biomarker for early disease detection. Because of poor pharmaceutical outcomes but excellent surgical outcomes in this patient group, early detection is critical for successful treatment before significant and irreversible damage occurs. Thanks to the contributions of Blue Waters, we now have sufficient pilot data to write a grant to translate this technology into the clinic for early detection.

With PowerGrid and Blue Waters we have scaled reconstructions for a single subimage of a single subject to over 200 GPUs (Fig. 2) with significant speedup via MPI parallelization and the high-speed interconnect present in the Cray XE/XK system. MPI has not been widely employed in MRI imaging, and to our knowledge, has not been combined with GPU acceleration for parallelizing single-image reconstructions. Currently, MRI imaging is limited to what can be reconstructed in a reasonable time with a workstation attached to the scanner. This provides poor efficiency in using the limited time that a patient is in the scanner to obtain high-spatial-resolution images with quantitative information relevant




Figure 2: Initial scaling results from PowerGrid suggest acceptable speedup and efficiency when scaling to 200 GPUs (XK nodes), including communication and Shifter containerization overhead. This scaling empowers our workflow to distribute a much larger reconstruction problem across multiple GPUs.

to disease progression. We hope to develop new applications and techniques that can be made tractable by exploiting cutting-edge resources such as Blue Waters.

WHY BLUE WATERS

The Blue Waters system provides a unique resource for MRE because it allows for rapid transfer of imaging data from the scanners and the low-latency memory transfers necessary for highly parallel computations. For PowerGrid, Blue Waters provides an environment with a large number of GPUs, as well as support for containerization technology, Shifter, that supports our development workflows.

PUBLICATIONS & DATA SETS

McIlvain, G., E. Telzer, and C. Johnson. Mechanical Properties of the *In Vivo* Adolescent Human Brain. *Developmental Cognitive Neuroscience*, in review (2018).

Anderson, A., et al., Multi-Excitation MRE in Aging Human Brain. *Proceedings of the 26th Annual Meeting of the International Society for Magnetic Resonance in Medicine* (Paris, France, 2018).

Anderson, A., 2018. *Magnetic Resonance Elastography and Nonlinear Inversion Problem in the Aging Brain*, doctoral dissertation, University of Illinois at Urbana-Champaign (2018), <http://hdl.handle.net/2142/101335>.

Johnson, C., et al., Double dissociation of structure–function relationships in memory and fluid intelligence observed with magnetic resonance elastography. *NeuroImage*, 171 (2018), DOI:10.1016/j.neuroimage.2018.01.007.

AUTOMATIC KNOWLEDGE BASE CONSTRUCTION AND HYPOTHESIS GENERATION: ANTIBIOTIC RESISTANCE MECHANISMS FOR *ESCHERICHIA COLI*

Allocation: NSF PRAC/350 Khn
PI: Ilias Tagkopoulos¹

¹University of California, Davis

EXECUTIVE SUMMARY

Antibiotic resistance is one of the leading threats to global health, food security, and development according to the World Health Organization [1]. The construction of a cohesive knowledge base for antibiotic resistance that can be a source for machine learning methods will have a broad impact in the field and eventually enable an artificial intelligence (AI) system to automate knowledge discovery in unprecedentedly efficient and unbiased ways. With this goal in mind, we built a knowledge base housing one million facts for *E. coli* antibiotic resistance mechanisms that are specifically structured for efficient machine learning. On top of this knowledge base, we trained the multilayered machine learning method for generating novel hypotheses of antibiotic resistance. The cross-validation results showed that the predictor can achieve an AUC (Area Under the Curve) of 0.868 for the ROC (Receiver Operator Characteristic) and Average Precision of 0.24, surpassing the baselines. The proposed framework is generically applicable, and we plan to make the tool publicly available so that anyone can apply it to their domain of interest.

RESEARCH CHALLENGE

Antibiotic resistance can affect anyone of any age or nationality. It is a natural phenomenon that is accelerated by our way of life and overuse of antibiotics in livestock and medicine. Aside from the global threat of antibiotic resistance, its immediate impact results in longer hospital stays, higher medical costs (estimated at \$20 billion annually in the United States alone), and increased mortality [1].

Constructing a knowledge base for antibiotic resistance that can be applied to machine learning methods for generating hypotheses will accelerate the overall discovery process in unprecedentedly efficient and unbiased ways [2–4]. A large number of knowledge bases have recently been created using graph representation, which stores factual information in the form of relationships among entities. These include YAGO [5], NELL [6], Freebase [7], and Google Knowledge Graph [8] for storing general facts about people, cities, and the like. Notable examples in the biological domain include KnowLife [9] and BioGraph [10].

However, key challenges exist to enabling the creation of a machine-learnable knowledge base for *E. coli* antibiotic




Figure 1: Schematic diagram of the proposed framework for knowledge base construction and novel hypothesis generation.




Figure 2: The ROC for each of the ER-MLP, PRA, and Stacked models (left). The Precision Recall curve for Stacked models (right).

resistance. Although knowledge conflicts exist and are reported in the literature [11], these inconsistencies are not curated in existing knowledge bases. Even in the literature, these conflicts are quantified only in part [12]. The inconsistencies are especially prevalent between high-throughput measurements and biological networks, making it nontrivial to draw biologically meaningful conclusions in an automated way [13]. Several truth discovery methods have been proposed over the past decade, and they have been successfully applied in diverse domains.

A primary application domain involves web sourcing, where information conveyed on a web page conflicts with other web pages [7]. Many truth discovery methods were first developed to resolve web-source conflicts. These methods employ one of three possible approaches: iterative methods, optimization methods, and probabilistic methods [14,15]. Recently, the use of a link prediction method has been proposed to decide truth among conflicts [16]. Some notable recent work in the biological sciences includes inconsistency repair in the *E. coli* gene regulatory network using answer set programming [13] and inconsistency resolution in signal transduction knowledge using integer linear programming [17].

Another challenge is that negative findings are not curated as well as positive findings in existing biological knowledge bases despite their importance in training the machine learning models to classify what knowledge is likely true [16]. Finally, existing knowledge bases do not annotate temporal information about antibiotic exposure despite its importance in the emergence of antibiotic resistance [18,19], which obscures the understanding of the time series dynamics of antibiotic resistance mechanisms.

METHODS & CODES

We built a generic framework that first constructs an inconsistency-free knowledge graph for a specific domain and then trains the hypothesis generator on top of this knowledge

graph that is based on the multilayered machine learning method. This method has been published elsewhere [20], and we optimized the published code to best utilize its principles for our needs.

RESULTS & IMPACT

In this work, we developed an inconsistency-resolved, machine learning-friendly, time series knowledge graph for *E. coli* antibiotic resistance. The knowledge graph incorporates a total of one million triples from 10 sources where the five distinctive durations of antibiotic exposure exist, ranging from 30 mins to seven days. Furthermore, this knowledge base houses positive and negative findings and thus facilitates training and evaluation of hypothesis generators that are built with machine-learning methods. On top of this knowledge base, we trained the multilayered machine-learning method that employs a path-ranking algorithm (PRA) and entity-relation multi-layered perceptrons (ER-MLP) for generating novel hypotheses of antibiotic resistance. The cross-validation results showed that the predictor can achieve an AUC of 0.868 for ROC and Average Precision of 0.24, surpassing the baselines. The proposed framework is generically applicable, and we believe this tool can accelerate knowledge discovery in an unbiased way by automatically generating a novel hypothesis from a knowledge base in a specific domain. This framework will be publicly available for use in other domains of interest.

WHY BLUE WATERS

Access to Blue Waters was critical to the success of this project, and project staff helped us to best utilize the Blue Waters high-performance computing resource for our needs by helping us to efficiently run large jobs in parallel and by responding to our requests in a timely manner.

MOLECULAR MECHANISM OF LIPID AND ION TRANSPORT IN PHOSPHOLIPID SCRAMBLING

Allocation hours: Illinois/600 Knh

PI: Emad Tajkhorshid¹

Collaborators: Tao Jiang¹, H. Criss Hartzell², Kuai Yu²

¹University of Illinois at Urbana-Champaign

²Emory University School of Medicine

EXECUTIVE SUMMARY

From bacteria to mammals, different phospholipid species are segregated between the inner and outer leaflets of the cellular membrane by adenosine triphosphate (ATP)-dependent lipid transporters. Disruption of this segregation by ATP-independent phospholipid scrambling is a key step in cellular signaling, e.g., inducing programmed cell death and blood coagulation. The mechanism by which scramblase (the protein responsible for phospholipid translocation) catalyzes rapid exchange of lipids between the two leaflets of a bilayer has been long sought. Using extensive molecular dynamics (MD) simulations on Blue Waters, we showed that a hydrophilic track formed on the surface of a scramblase serves as the pathway for both lipid and ion translocations, and that Ca^{2+} ion binding controls the open/closed transition of the track. This microscopic view of the lipid transport process sheds light on how lipophilic molecules can permeate specialized proteins to travel between the two leaflets of the cellular membrane—a process that is of broad physiological and biomedical relevance.

RESEARCH CHALLENGE

Different phospholipid species are distributed asymmetrically between the two leaflets of the cellular membrane. Dissipation of this asymmetry in response to the elevation of cytoplasmic Ca^{2+} concentration is a ubiquitous signaling mechanism critical for diverse cellular events including blood coagulation, bone mineralization, and cell–cell interaction [1–3]. Phospholipid scrambling is mediated by phospholipid scramblases, which harvest the energy of the phospholipid gradient to drive nonspecific and bidirectional transport of phospholipids between leaflets. Proteins responsible for Ca^{2+} -activated lipid scrambling belong to the TMEM16 superfamily of membrane proteins, with some members being Ca^{2+} -activated Cl^- channels, while others function as Ca^{2+} -activated scramblases and nonselective ion channels. Despite the remarkably diverse functions of TMEM16 proteins, both subfamilies share a common dimeric architecture and mode of Ca^{2+} activation [4–6]. However, the absence of phospholipids and ionic substrates in the solved structures leaves the question of how both lipids and ions are conducted unanswered. It also




Figure 2: Lipids lining the hydrophilic track on the surface of the nhTMEM16 scramblase play a structural role in forming a “proteolipidic” pore, which is likely to be used by ions to cross the membrane. The permeating Na^+ ions during the simulations are shown in time series snapshots (blue spheres).

remains elusive how the same architecture accommodates such diverse functions.

METHODS & CODES

We performed extended MD simulations on the atomic models of the nhTMEM16 scramblase in asymmetric lipid bilayers in the presence of Ca^{2+} ions at the activation binding sites. To uncover the nature of Ca^{2+} dependence, we also simulated a Ca^{2+} -free conformation of the protein in the same condition. To examine the ion permeation properties of the scramblase protein, we extended the fully equilibrated Ca^{2+} -activated simulation under multiple levels of applied transmembrane voltage, from which we calculated the ionic conductivity across the membrane. All MD simulations were carried out on Blue Waters using the NAMD (NANoscale Molecular Dynamics) simulation package [7].

RESULTS & IMPACT

Our simulations reveal a significant deformation of the membrane structure induced by the nhTMEM16 scramblase protein due to its surface hydrophobicity (Fig. 1). The bending and thinning of the lipid bilayer primes lipid translocation by greatly reducing the energy barrier for hydrophilic head groups to move across the membrane. As the simulation was extended, a membrane-spanning lipid translocation track appeared on the protein surface (through hydration and occupancy of lipid head groups), effectively connecting the inner and outer leaflets of the bilayer (Fig. 1). In an aggregate 3 microseconds of Ca^{2+} -activated simulation, we observed one spontaneous full scrambling event through this track under equilibrium conditions and four full scrambling events in the presence of voltage. The observed track provides a hydrophilic environment for head groups to translocate between leaflets while the hydrophobic acyl chains are exposed

to the hydrophobic phase of the bilayer. Simulations indicate that Ca^{2+} binding stabilizes the open conformation of the track by altering the structure of the lining transmembrane helices.

Our simulations also provide mechanistic insights into the ion channel properties of TMEM16 proteins, revealing that the membrane-spanning lipid track forms an ion-conducting “proteolipidic” pore between the protein and lipid head groups (Fig. 2). This flexible pore structure explains a number of unusual features of TMEM16 ionic currents, especially their highly variable ionic selectivity and ability to permeate large ions. In addition, key amino acids predicted by our simulations to enhance scrambling activity have been used successfully to experimentally engineer scramblase activity in a homologous Ca^{2+} -activated ion channel, thus providing insight into the evolutionary relationship of the TMEM16 family members.


WHY BLUE WATERS

The high-performance architecture of Blue Waters makes it an excellent computing resource for our scientific research. The GPU-accelerated simulation program NAMD has been extensively tested and optimized for Blue Waters. The large number of GPUs available in the XK nodes significantly increased our overall scientific productivity. Finally, the technical support provided by the expert scientists of the Blue Waters team has greatly facilitated the accomplishment of our research goals.

PUBLICATIONS & DATA SETS

Jiang, T., K. Yu, H.C. Hartzell, and E. Tajkhorshid, Lipids and ions traverse the membrane by the same physical pathway in the nhTMEM16 scramblase. *eLife*, 6:e28671 (2017), DOI:10.7554/eLife.28671.

Figure 1: (Top) Membrane deformation induced by nhTMEM16, shown by the average phospholipid phosphate density over the simulation (orange surface) compared to the initial distribution (tan spheres). (Bottom) Representative simulation snapshots showing the empty track at $t = 0$, which becomes fully occupied by lipids in both subunits during the simulation.



DI

GA

TN

STRUCTURE AND FUNCTION OF BACTERIAL CHEMOSENSORY ARRAYS

Allocation: NSF PRAC/8,100 Knh

PI: Emad Tajkhorshid¹

Co-PIs: James C. Phillips¹, John E. Stone¹, Rafael C. Bernardi¹

Collaborators: Keith Cassidy¹, Peijun Zhang², John S. Parkinson³, Ariane Briegel⁴

¹University of Illinois at Urbana-Champaign

²Oxford University

³University of Utah

⁴Leiden University

EXECUTIVE SUMMARY

Chemotaxis—the ability of a cell or organism to direct its movement in response to environmental chemicals—is a ubiquitous biological behavior. In chemotactic bacteria, extended transmembrane protein clusters called chemosensory arrays process complex sensory signals to affect motile behavior and thereby aid in survival. Remarkably, the molecular architecture of the chemosensory array is universally conserved across species, providing a novel target for broadly applicable antimicrobial therapies. Despite significant effort, however, the array's large size and dynamic nature have thwarted high-resolution structural characterization by conventional techniques, preventing a detailed description of array function. Utilizing large-scale molecular simulations on Blue Waters, we have constructed the first atomically resolved model of a complete chemosensory array, enabling the study of critical aspects of sensory signal transduction as never before.

RESEARCH CHALLENGE

Chemotaxis plays a critical role in colonization and infection by many human and plant pathogens [1]. With antimicrobial-resistant infections projected to kill more people worldwide per year than cancer by 2050 [2], a comprehensive understanding of the chemotaxis response could have a significant impact on societal wellbeing.

The streamlined chemotaxis machinery of the bacterium *Escherichia coli* is the best understood biological signal transduction system, and serves as a powerful tool for investigating cellular behavior and motility [3]. Essential to their chemotaxis response, *E. coli* cells possess a highly ordered sensory apparatus known as the chemosensory array (Fig. 1), which consists of hundreds of basic core-signaling units comprised of three components: transmembrane chemoreceptors, the histidine kinase CheA, and the adaptor protein CheW [4,5]. Despite important progress in the characterization of chemosensory arrays using a battery of genetic, biochemical, and structural techniques, the lack of a high-resolution description of the array's intact, multicomponent structure has prevented a detailed understanding of the molecular mechanisms underlying array function. Hence, the immediate goal

of this project is to provide a high-fidelity, atomistic model of the complete chemosensory array. The resulting model will enable computer simulations of the molecular mechanisms underlying cooperative sensory signal transduction, and thereby aid in the development of novel antibiotics to control multidrug-resistant strains of bacteria.

METHODS & CODES

To characterize the structural and dynamical properties of the chemosensory array, we carried out extensive all-atom molecular dynamics (MD) simulations using NAMD [6], a highly scalable MD code optimized to make peak use of Blue Waters' petascale features. The all-atom model depicted in Fig. 1 was constructed using a variety of so-called hybrid modeling techniques [7]. We analyzed simulation trajectories using several dimensionality-reducing techniques, especially k-medoids clustering [8].

RESULTS & IMPACT

To understand the underlying molecular mechanism of chemosensory array activation, regulation, and cooperativity, it is essential to understand the precise interactions among chemoreceptors, CheA, and CheW in the context of their native array organization and as they undergo dynamical changes in response to chemical signals. Recently, our collaborator Peijun Zhang developed a novel *in vitro* reconstituted-monolayer system, producing array specimens ideal for high-resolution characterization by cryo-electron microscopy (cryoEM). This analysis resulted in a density map of the cytoplasmic portion of the *E. coli* core-signaling unit (CSU) at 9 Å resolution. Using Molecular Dynamics Flexible Fitting (MDFF) simulations [9], we integrated homology models based on existing crystallographic structures with our cryoEM data to produce an atomic model of the intact *E. coli* cytoplasmic CSU. Using Generalized Simulated Annealing [10] to further explore the conformational spaces of especially flexible array components, we were able to identify a number of novel CheA conformations contained within the averaged cryoEM data and, thereby, to significantly advance the mechanistic description of kinase activation. These findings have produced a number of residue-based predictions that have




Figure 1: All-atom model of a complete bacterial chemosensory array (20 million atoms). Environmental ligands are sensed by binding to chemoreceptors (red), which then transmit signals across the cell's inner membrane to cooperatively regulate the activity of the CheA kinase (blue) with the help of the adaptor protein CheW (green).

subsequently been validated via functional assays in live cells by our collaborator Sandy Parkinson (manuscript forthcoming).

In addition, based on extensive mutagenesis data from the Parkinson Lab, as well as crystallographic and disulfide crosslinking data, we have developed a model of the complete *E. coli* serine chemoreceptor, including the previously uncharacterized transmembrane bundle and HAMP domains. Using this model, along with *in vivo* cryoEM data from our collaborator Ariane Briegel, which provided a view of the membrane-proximal regions of the chemoreceptors in multiple signaling states, we were able to extend our cytoplasmic model to include intact chemoreceptors and membranes as well as to elucidate critical signaling-related interactions among chemoreceptors. Taken together, the boundary-pushing simulations described above represent the first major effort to bring the immense power of modern supercomputing to bear on the problems of bacterial chemotaxis. These simulations have already helped to fill in the missing atomistic detail necessary for elucidating several long-standing problems in the field and provide in general a much-needed molecular platform on which to generate wholly new hypotheses.

WHY BLUE WATERS

To properly investigate the long-range and highly cooperative mechanisms of sensory signal transduction within the chemosensory array, it is necessary to simulate large, multicomponent molecular systems comprised of millions of atoms on the order of microseconds. Such simulations require sustained access to thousands of tightly coupled processors and are able to benefit tremendously from GPU acceleration. Thus, the unique petascale capabilities of the Blue Waters machine have provided a premier computing infrastructure without which this work could not have been completed.

PUBLICATIONS & DATA SETS

Cassidy, C.K., et al., CryoEM and computer simulations reveal a novel kinase conformational switch in bacterial chemotaxis signaling. *Elife*, 4:e08419 (2015).

LARGE-SCALE COARSE-GRAINED SIMULATIONS OF HIV-1: NEW THERAPEUTIC TARGETS

Allocation: NSF PRAC/60 Knh

PI: Gregory Voth¹

¹University of Chicago

EXECUTIVE SUMMARY

A key step in the proliferation of human immunodeficiency virus type-1 (HIV-1) is the so-called maturation of viral particles that have been released from infected cells. Proteins within these particles are “activated” through proteolytic cleavage and subsequently assemble into conical cores that house the viral genome. Hence, one viable therapeutic strategy is to develop drugs that prevent proteolytic activity; e.g., protease inhibitors or maturation inhibitors. Recently, a promising class of antiretroviral drugs (GS-CA1 from Gilead Sciences) has been proposed that targets the capsid protein directly. Enabled by Blue Waters, we performed large-scale coarse-grained (i.e., with reduced representation) molecular simulations of capsid proteins to elucidate the putative mechanism of these capsid inhibition drugs.

RESEARCH CHALLENGE

The infectious form of HIV-1 is prepared during a stage called “maturation,” during which viral proteins are processed through proteolytic cleavage and freed capsids assemble into protein

cores that host the viral genome. Understanding the molecular mechanisms of this process is important since disruption of capsid assembly may be a viable target for therapeutics. Drugs that specifically target capsid proteins (e.g., in comparison to other viral constituents, such as protease) are known as capsid inhibitors. One such drug—GS-CA1—has recently been reported as a potential long-acting, highly potent antiretroviral. However, the mechanisms of action of these drugs remain unclear. A fundamental understanding of these mechanisms will facilitate the future design of this relatively unexplored class of antiretroviral drugs.

METHODS & CODES

To investigate the self-assembly of the mature capsid core, we leveraged previously developed coarse-grained (CG) molecular models of the relevant viral components, which are based on experimental data [1]. These CG models are computationally efficient representations of molecules, thereby enabling computer simulations to be performed at time- and length-scales that are

otherwise inaccessible. We performed molecular dynamics simulations using our CGMD software package that was developed during a previous PRAC proposal, which was designed to take full advantage of Blue Waters’ hardware and network capabilities [2]. Our simulation results were used to understand potential mechanisms of capsid inhibition drugs during HIV-1 maturation, a key step in the viral lifecycle.

RESULTS & IMPACT

Building on the biophysical insights we have developed during our previous CG simulation studies [1,3], we investigated the putative mechanism of capsid inhibitors on mature capsid assembly. We found that even at small concentrations of drugs, which over-stabilize protein assembly intermediates, a plethora of nonideal capsid assembly pathways emerge (Fig. 1). As a result, the population of canonical and infectious capsid cores is dramatically reduced. Here, the mechanism of action is notably different from classic antiretrovirals; whereas most drugs aim to disrupt or weaken certain interactions, the current class of drugs aims to strengthen interactions. We anticipate that these insights can be easily extended to other viral systems and justify further research into capsid inhibition. Our simulation results are also exemplary of an emerging paradigm for biomedical research in which theory and experiment combine to accelerate drug design.

WHY BLUE WATERS

The primary benefit of CG models is that their reduced representations enable access to large numbers of molecules (e.g., proteins) that interact over long lengths of time. In biology, especially in the context of macromolecular complexes, these length- and time-scales are necessarily large and are otherwise inaccessible using conventional molecular simulations. However, communication bottlenecks during runtime prevent efficient use of standard high-performance computing resources. It was therefore crucial to have access to large amounts of computational power, especially combined with superior network performance. The Blue Waters compute platform thus presented a natural choice for our work; the combination of leadership-class compute capabilities with cutting-edge network hardware allowed us to successfully investigate a system of significant biomedical interest. Our large-scale CG simulations were performed using techniques developed for previous work on the Blue Waters platform [2], while existing relationships with Blue Waters technical project staff greatly assisted in their deployment.

PUBLICATIONS & DATA SETS

Pak, A., et al., Immature HIV-1 lattice assembly dynamics are regulated by scaffolding from nucleic acid and the plasma membrane. *Proc. Natl. Acad. Sci. U.S.A.*, 114:47 (2017), pp. E10056–E10065.

Grime, J., et al., Coarse-grained simulation reveals key features of HIV-1 capsid self-assembly. *Nat. Commun.*, 7 (2016), p. 11568.

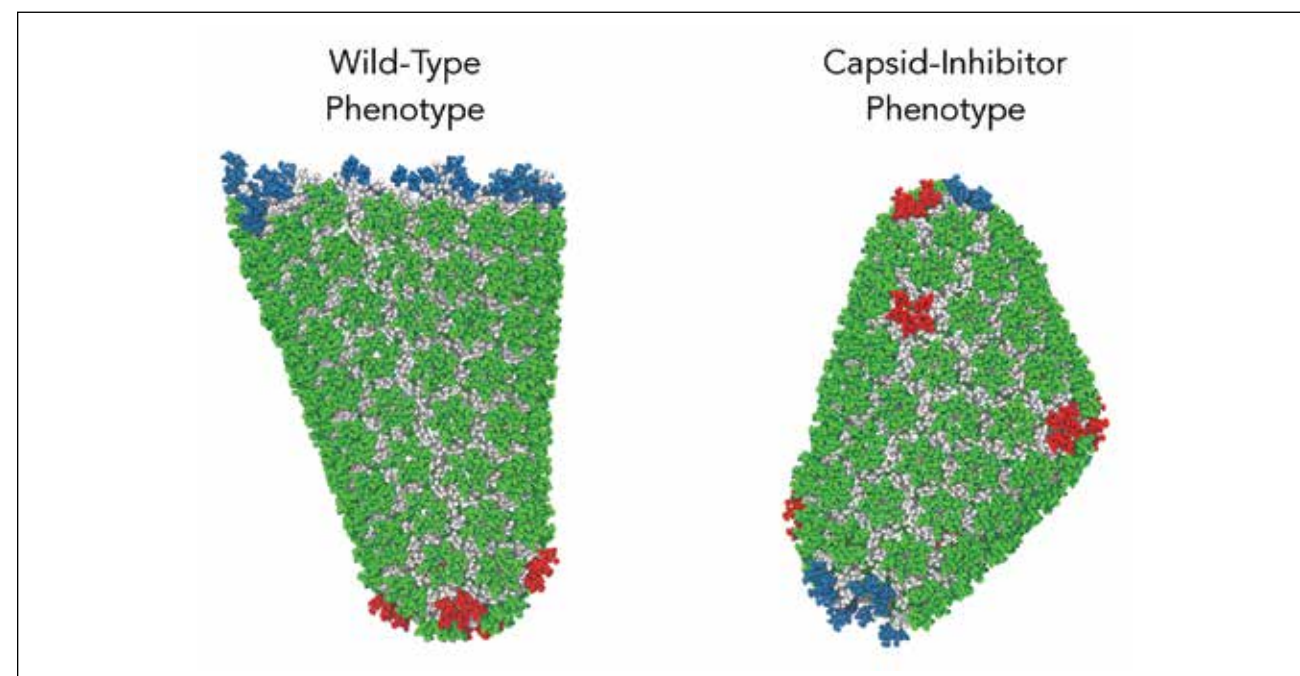


Figure 1: Comparison of example end-point structures from wild-type (left) and capsid-inhibited (right) assembly pathways; the N-terminal domains of edge-case dimers, hexamers, and pentamers are shown in blue, green, and red, respectively. In the latter case, accelerated pentameric assembly increases structural polymorphism with notable curvature variability and incomplete enclosures.

IMPROVING ACCURACY AND SCALABILITY FOR CORE BIOINFORMATICS ANALYSES

Allocation: Illinois/125 Knh

PI: Tandy Warnow¹

Collaborators: Erin Molloy¹, Michael Nute¹, Ehsan Saleh¹, Kodi Collins²

¹University of Illinois at Urbana-Champaign

²University of California, Los Angeles

EXECUTIVE SUMMARY

This project developed new methods for large-scale evolutionary tree construction and multiple sequence alignment that can be used to address fundamental science problems such as “How did life evolve on Earth?” and “What function does this protein have?” The most important outcome is the discovery we made regarding protein sequence alignment methods: Our research suggests that BALi-Phy, the leading statistical method for multiple sequence alignment, has the best accuracy of all methods tested on simulated data sets, but is less accurate than standard multiple sequence alignment methods when evaluated on protein benchmark data sets. While the cause for this difference in performance between biological and simulated data is not yet known, each of the most likely explanations (i.e., either model misspecification or errors in the protein benchmark data sets) presents troubling ramifications for other problems in biology, including molecular systematics and protein structure and function prediction.

RESEARCH CHALLENGE

Much biological research—including the estimation of evolutionary histories, the prediction of protein structure and function, and the detection of positive selection—requires that a set of molecular sequences first be “aligned” with each other. Furthermore, multiple sequence alignment of large data sets is necessary for many biological studies. Most obviously, the construction of the tree of life will require millions of sequences, spanning large evolutionary distances. Less obviously, protein structure prediction also benefits from large data sets: the most accurate protocols for predicting the structure of an unknown protein from its sequence of amino acids begins by collecting a very large number of related sequences and then computing a multiple sequence alignment on that set.

Unfortunately, large-scale multiple sequence alignment is enormously difficult to perform with high accuracy. The only methods that have been able to run on ultra-large data sets (with up to one million sequences) are PASTA [1] and UPP [2], which use a divide-and-conquer approach to scale other alignment methods to large data sets. Both PASTA and UPP have excellent accuracy on simulated data sets, but less is known about their accuracy on biological data sets, and especially on protein data sets, where alignment estimation may be enabled through the use of inferred or known structural elements. BALi-Phy [3], one of the most promising approaches, infers an alignment under a statistical model of sequence evolution and is expected to have the best accuracy of all methods. Yet, BALi-Phy is too computationally intensive to use directly on large data sets. Alternatively, BALi-Phy can be incorporated into PASTA and UPP so that it scales to large data sets [4]. These “boosted” versions of BALi-Phy have outstanding accuracy on simulated nucleotide data sets [4]; however, when this project began, it was not known if BALi-Phy or the “boosted” version of BALi-Phy would provide improved accuracy on biological data sets (nucleotides or amino acids) in comparison to standard multiple sequence alignment methods. Efficient and scalable multiple sequence alignments that have improved accuracy on ultra-large biological data sets, and especially for protein sequences, would therefore provide major benefits for many downstream analyses.




Figure 1: Accuracy on simulated data sets (each with 27 sequences) for different multiple sequence alignment methods, showing averages (over 20 replicates) for modeller score (precision) and SP score (recall). Note that BALi-Phy has the highest accuracy of all methods under all conditions. (Figure taken from [5].)




Figure 2: Average accuracy on four biological benchmark collections, each with at most 25 amino acid sequences (1,192 sequence data sets in all) for different multiple sequence alignment methods, showing modeller score (precision) and SP score (recall). Note that BALi-Phy tends to have the best average recall of all methods for all benchmarks. (Figure taken from [5].)

METHODS & CODES

We performed two major studies regarding protein sequence alignment. In the first study [5], we compared BALi-Phy to leading protein sequence alignment methods on data sets from four established benchmark collections of protein sequences. Since BALi-Phy is computationally intensive, we limited the study to small data sets. The second study evaluated the impact of integrating the best-performing methods from this first study into PASTA.

RESULTS & IMPACT

BALi-Phy has outstanding accuracy on the simulated data sets (Fig. 1), clearly dominating all the other methods with respect to both recall (*sum of pairs* or SP score) and precision (modeller score). Yet, on the biological data sets (Fig. 2), BALi-Phy has much poorer precision and recall. In fact, BALi-Phy typically has only average recall and often is among the poorest of the top alignment methods. The best-performing multiple sequence alignment methods in this study have been integrated into PASTA (thus attaining improved scalability and reduced running time, while maintaining accuracy), and the new version of PASTA is available at [6] in open-source form.

The distinction in accuracy on biological data sets and simulated data sets is troubling and requires further investigation. BALi-

Phy’s excellent accuracy on simulated data is expected since the simulated data sets are generated under statistical models of sequence evolution that are close, even if not identical, to the statistical models under which BALi-Phy performs its inference. However, since BALi-Phy is so much less accurate on biological data sets, this suggests that the protein data sets have evolved under processes that are quite different from the ones that are well modelled by BALi-Phy. While it has always been expected that there would be some level of model misspecification (as no model is perfect), for there to be a substantial difference in relative accuracy between simulated and biological data sets suggests that the level of model misspecification must be quite large. This would be a troubling conclusion, since many bioinformatics analyses are performed under statistical models similar to the one assumed in BALi-Phy. However, there are other potential explanations, one of which is that the reference alignments in the biological benchmarks may themselves not be highly accurate (i.e., they may be inferred through a combination of information about structural features in the proteins and then interpolation among the structurally derived parts of the alignment using software tools). If this is a reason for the discordance, then BALi-Phy may still be useful, but the benchmarks must be questioned. Future research is needed to explore these possible explanations, as well as the others that we posit, as discussed in [5].

WHY BLUE WATERS

This study used 230 CPU years for the BALi-Phy analyses alone and would not have been feasible on other computational systems available to the project team.

PUBLICATIONS & DATA SETS

Nute, M., E. Saleh, and T. Warnow, Benchmarking statistical alignment. *bioRxiv* (2018), DOI:10.1101/304659.

Nute, M., et al., The performance of coalescent-based species tree estimation methods under models of missing data. *BMC Genomics*, 19: Supplement 5 (2018), DOI:10.1186/s12864-018-4619-8.

Christensen, S., et al., OCTAL: Optimal completion of gene trees in polynomial time. *Algorithms for Molecular Biology*, 13:6 (2018), DOI:10.1186/s13015-018-0124-5.

Mirarab, S., et al., PASTA github site (<https://github.com/smirarab/pasta>), accessed May 13, 2018.

Collins, K., and T. Warnow, PASTA for Proteins Data (BALiBASE). University of Illinois at Urbana-Champaign (2018), DOI:10.13012/B2IDB-4074787_V1.

Christensen, S., et al., Datasets from the study “OCTAL: Optimal Completion of Gene Trees in Polynomial Time.” University of Illinois at Urbana-Champaign (2018), DOI:10.13012/B2IDB-1616387_V1.

Nute, M., et al., Data from: The Performance of Coalescent-Based Species Tree Estimation Methods under Models of Missing Data. University of Illinois at Urbana-Champaign (2017), DOI:10.13012/B2IDB-7735354_V1.

SOCIAL SCIENCE, ECONOMICS, & HUMANITIES

ECONOMICS

POLITICAL SCIENCE

264 *Enabling Redistricting Reform: A Computational Study of Zoning Optimization*

266 *Policy Responses to Climate Change*

ENABLING REDISTRICTING REFORM: A COMPUTATIONAL STUDY OF ZONING OPTIMIZATION

Allocation: Illinois/400 Knh
PI: Wendy Cho¹
Co-PI: Yan Liu¹
Collaborator: Bruce Cain²

¹University of Illinois at Urbana-Champaign
²Stanford University

EXECUTIVE SUMMARY

We have developed a scalable computational tool for redistricting that synthesizes and organizes massive amounts of computation and data to evaluate voter 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 analysis.

RESEARCH CHALLENGE

In the United States, political redistricting occurs every 10 years following the decennial census. It is intended to provide fair representation for all communities and interest groups. This process is hampered when those drawing the maps unfairly advantage a particular partisan or racial group over others. Despite broad disdain for the practice of gerrymandering, the Supreme Court has found it difficult to identify a workable standard by which gerrymandering might be regulated. Moreover, in more than two-thirds of the states, the majority and self-interested

political party in the lower state house has the responsibility of devising the electoral map, creating a process that is inherently skewed toward biased outcomes.

In the context of litigation, there are insufficient tools to synthesize redistricting data to provide a sufficient basis from which to analyze and decide the legal issues. The problem is, 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 to be capable of being decided by legal principles, the court has yet to identify manageable standards under which one could measure and identify a partisan gerrymander. The failure of the legal system in this political realm has significant ramifications for our democratic system of governance.

At the initial map drawing stage, computation has improved significantly in the last two decades, but this advancement has served only to further isolate and insulate the process since only politicians and those in their employ have had access and the knowledge to exploit these tools. In essence, technological advances have only served to increase the manipulability of the redistricting process by those with nefarious intent. Our work intends to reverse this trend.




Figure 1: By applying the path relinking heuristics on spatial configurations, the spatial crossover operation is able to generate multiple new solutions that propagate alleles from.

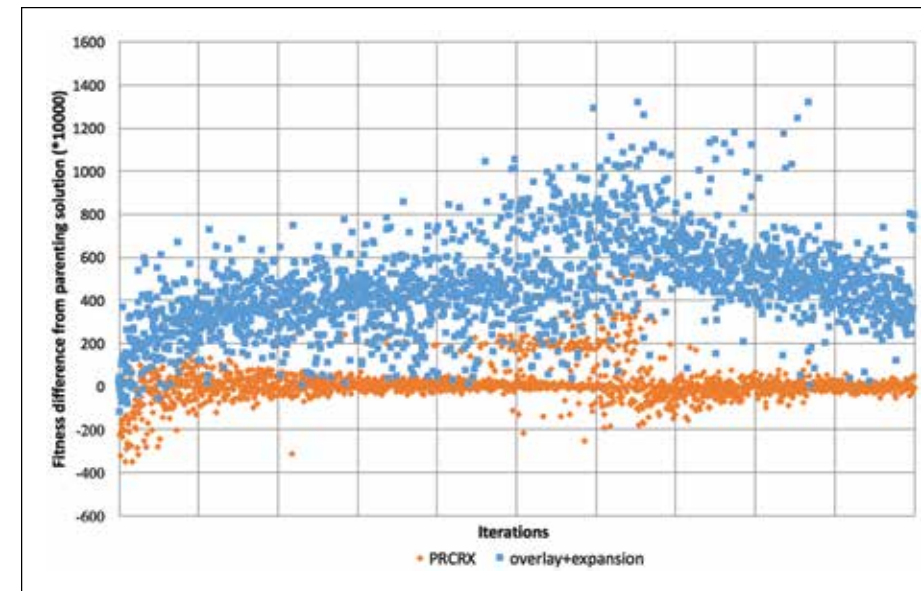


Figure 2: Performance comparison: Spatial Path Relinking crossover operator and the basic overlay + expansion crossover. Both take the same input. Fitness difference between the input and output solution is plotted. Negative difference value means fitness improvement.

Using the Supreme Court’s articulated legal reasoning and mandates, we have developed a computational redistricting tool that formulates redistricting as a combinatorial optimization problem, with objectives and constraints defined to meet legal requirements. Drawing electoral maps amount to arranging a finite number of indivisible geographic units into a smaller number of districts where every unit must belong to exactly one district and no districts are empty. The redistricting problem is similar to the set-partitioning problem that is known to be NP-hard and, thus, the time required to solve the problem increases very quickly as the size of the problem grows. Our scalable evolutionary computational approach utilizes massively parallel high-performance computing for redistricting optimization and analysis at fine levels of granularity.

We intend for our tool to be used both at the initial map-drawing stage as well as at the litigation stage. At the map-drawing stage, the tool is far more advanced than the extant tools that politicians employ. It also incorporates the idiosyncratic legal criteria required of electoral maps so that those without this specialized knowledge are still able to utilize it. At the litigation stage, it enables the statistical analyses that the Supreme Court has found illusive. Both potential uses are intended to improve outcomes from the redistricting process.

METHODS & CODES

Our algorithm, Parallel Evolutionary Algorithm for Redistricting (PEAR), is implemented in ANSI C. It can be compiled on Linux and OS X 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.

We devised the evolutionary algorithm operators specifically for the redistricting application. The effectiveness and efficiency are enhanced by the explicitly spatially aware mutation and crossover operators. These operators are much more efficient in navigating spatially constrained decision space where the solution space exhibits extremely high ruggedness that classic EA operators could not explore comprehensively. The crossover operator employs a path relinking heuristic on adjacency planar graphs and generates new solutions by randomized graph and geometry calculation of suitable neighborhoods. 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 addressing fine levels of spatial granularity. We leveraged and enhanced a scalable Parallel Genetic Algorithm library to develop PEAR for the computationally intensive redistricting problem. PEAR provides a powerful and computationally scalable redistricting tool that has never before existed by incorporating a set of spatial configuration operators and spatial EA operators to handle spatial characteristics and the associated computational challenges, and by harnessing massive computing power. It incorporates a novel spatial crossover operator that incorporates the optimization literature on path relinking methods for spatial constraints.

Our work has been referenced in 11 amicus briefs and discussed in oral arguments before the Supreme Court in the landmark

Continued on page 267

POLICY RESPONSES TO CLIMATE CHANGE

Allocation: GLCPC/250 Knh

PI: Yongyang Cai¹

Co-PIs: Kenneth Judd², William Brock³, Thomas Hertel⁴

Collaborators: Simon Scheidegger⁵, T.J. Canann⁶, Carlos Rangel⁷

¹The Ohio State University

²Hoover Institution

³University of Wisconsin

⁴Purdue University

⁵University of Zurich

⁶University of Minnesota

⁷Penn State University

EXECUTIVE SUMMARY

Cai, Brock, Xepapadeas, and Judd [1] have built a model of Dynamic Integration of Regional Economy and Spatial Climate under Uncertainty (DIRESCU), incorporating a number of important climate science elements that are missing in most integrated assessment models. These include spatial heat and moisture transport from low latitudes to high latitudes, sea level rise, permafrost thaw, and tipping points. Using this more realistic model of the world economy and climate, we study policy responses to climate change under cooperation and various degrees of competition among regions. We find that other assessment models that are missing elements of climate science lead to significant bias in important policy variables such as the social cost of carbon and adaptation.

RESEARCH CHALLENGE

Leading integrated assessment models assume that climate damages are related to the mean surface temperature of the planet. But climate science shows that when the climate cools or warms, high-latitude regions tend to exaggerate the changes seen at lower latitudes due to spatial heat and moisture transport. This effect is called polar amplification (PA). Thus, the surface temperature anomaly is differentiated across spatial zones of the globe. The low- (high-) latitude regions would be hotter (colder) if poleward heat transport were absent; hence, damages in the low-latitude regions would be higher since they are already under heat stress and transporting some of that heat poleward helps relieve this heat stress.

PA will accelerate the loss of Arctic sea ice, a potential meltdown of the Greenland and West Antarctica ice sheets, which could cause serious global sea level rise. Moreover, PA will lead to faster thawing of the permafrost, which is expected to bring about widespread changes in ecosystems and damage to infrastructure, along with the release of greenhouse gases that exist in permafrost carbon stocks. Furthermore, PA will also affect the likelihood of tipping points, such as the “nearest” three potential tipping points located in the high latitudes of the Northern Hemisphere

(Arctic summer sea ice loss, Greenland ice sheet melt, and boreal forest loss).

METHODS & CODES

We developed the DIRESCU model to include spatial heat and moisture transport from low latitudes to high latitudes, sea level rise, permafrost thaw, and tipping points. To model spatial heat and moisture transport, we disaggregate the globe into two regions: Region 1 is the region north of latitude 30°N to 90°N (called the North), while Region 2 is the region from latitude 90°S (the South Pole) to 30°N (called the Tropic-South). The disaggregation also makes clear their significant economic difference, since most countries in the Tropic-South are poor and most countries in the North are rich.

To address the tipping points and solve the dynamic stochastic programming problem, we adapt the computational method in DSICE [2], developed by Cai and Judd in the past four years using GLCPC allocations on Blue Waters. The computational method is parallel backward value function iteration using 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. Our code shows high parallel efficiency, with an almost linear speed-up from 30 nodes to 5,000 nodes.

RESULTS & IMPACT

In 2017, Cai, Judd, and Steinbuks published a paper in *Quantitative Economics* [3] that develops a nonlinear certainty equivalent approximation (NLCEQ) method to solve efficiently and in parallel huge-dimensional dynamic stochastic problems without exogenous trends, by using Blue Waters resources. We also extended NLCEQ to solve high-dimensional dynamic stochastic problems with exogenous trends and applied it to analyze the effect of climate and technological uncertainty in crop yields on the optimal path of global land use in [4].

Also in 2017, Yeltekin, Cai, and Judd published a paper in *Operations Research* [5] that developed a parallel algorithm that can solve supergames with states, which models strategic interactions among multiple players, by using Blue Waters resources.

In 2018, Cai, Brock, Xepapadeas, and Judd released a National Bureau of Economic Research working paper [1] that is under review for publication in a prestigious economic journal. The paper builds the DIRESCU model, studies optimal climate policies under cooperation and various degrees of competition among regions, and finds that excluding some of the elements of climate science leads to significant bias in important policy variables such as the social cost of carbon and adaptation.

WHY BLUE WATERS

Our parallel computational package requires low-latency communications because the algorithm uses the master-worker structure and needs frequent communications between the master and workers. Our problems are large. For example, the DIRESCU model has 10 continuous state variables and one binary state variable, as well as eight continuous decision variables, and a more than 500-year horizon. It corresponds to solving a Hamilton–Jacobi–Bellman equation with 10 or 11 state variables. Using our efficient parallel algorithm, we solved it with one specification case in 3.4 wall-clock hours with 102 computer nodes on Blue Waters. Moreover, we have solved the model with many specification cases for analysis. In addition, the largest problem we solved for DSICE used 3,459 computer nodes and took 7.5 wall-clock hours on Blue Waters. Blue Waters allows us to solve these large problems efficiently as has already been shown in our previous work.

PUBLICATIONS & DATA SETS

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.

Continued from page 265

partisan redistricting case, *Gill v. Whitford*. Our computational approach to redistricting reform has garnered media attention from popular outlets (Vox.com, *Chicago Inno*, *Reason*, the *Washington Post*); computing outlets (Cray Inc., TOP500, *Communications of the ACM*); and outlets aimed at the science and mathematics communities (*Quanta Magazine*, Science Node, *WIRED*, *Nature*). Most of the discussion has focused on the impact of our work on gerrymandering litigation. However, our work is also applicable to earlier stages of the redistricting process, which is now evolving and becoming clearer to practitioners and politicians as the project progresses.

WHY BLUE WATERS

The PEAR library is designed for extreme-scale redistricting applications. From the beginning, it was intended to scale to all of the processor cores on Blue Waters through nonblocking MPI communication calls. The computational approach implemented in 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 & DATA SETS

Cho, W.K.T., and Y.Y. Liu, Sampling from Complicated and Unknown Distributions: Monte Carlo and Markov Chain Monte Carlo Methods for Redistricting. *Physica A*, (2018), DOI:10.1016/j.physa.2018.03.096.

Cain, B., et al., A Reasonable Bias Method for Redistricting: A New Tool for an Old Problem. *William & Mary Law Review*, 59: 5 (2018).

Cho, W.K.T., and Y.Y. Liu, Massively Parallel Evolutionary Computation for Empowering Electoral Reform: Quantifying Gerrymandering via Multi-objective Optimization and Statistical Analysis. *SC17: The International Conference for High Performance Computing, Networking, Storage and Analysis* (Denver, Colo., November 12–17, 2017).

Cho, W.K.T., Measuring Partisan Fairness: How well does the Efficiency Gap Guard against Sophisticated as well as Simple-Minded Modes of Partisan Discrimination? *University of Pennsylvania Law Review Online*, 166 (2017).

Liu, Y.Y., and W.K.T. Cho, A High-Performance Evolutionary Computation Framework for Scalable Spatial Optimization. *ICCS: International Conference on Computational Science: Science at the Intersection of Data, Modelling and Computation* (ICCS, Wuxi, China, June 11–13, 2018).

GRADUATE FELLOWS

- 270** *Quantifying Forest Drought Resistance*
- 272** *Multiscale Biophysical Interactions: Inertia, Growth, and Sargassum Seed Populations*
- 274** *Simulating the Circumgalactic Medium: The Need for Robust Cosmic Ray Transport*
- 276** *How the Solar System Survived a Violent Epoch of Dynamical Instability*
- 278** *Sediment Transport in Estuaries: Estimating Bed Shear Stress Distribution from Numerically Modeled Tides in an Energetic Estuary*
- 280** *Machine Learning Harnesses Molecular Dynamics to Develop Therapeutic Strategies for Alzheimer's Disease and Chronic Pain*
- 282** *Scaling Relationships across Modeling Resolutions in Mountain Headwaters: Understanding Climate Change Impacts on Rocky Mountain Hydrology in a Numerical Modeling Context*
- 284** *Electronic Structure of NiFe Oxyhydroxide Active Sites: Understanding an Effective Oxygen-Evolving Catalyst at a Quantum-Mechanical Level*
- 286** *Refining the Conformational Ensembles of Flexible Proteins Using Simulation-Guided Spectroscopy*
- 288** *Designing Materials in Flow: Multiscale Simulations of Complex Fluid Rheology*
- 290** *High Accuracy Radiative Transfer in Cloudy Atmospheres*
- 292** *High-Fidelity Blade-Resolved Wind Farm Simulations*
- 294** *Discovering Hundreds of New Exoplanets with K2*
- 296** *Toward Quantitative Prediction of Defect Tolerance in Semiconductors: Leveraging HPC and Physical Insights to Design Radically Cheaper Energy Materials*
- 298** *Wires Within Wires: A Multiscale Model for Computational Investigation of Bioelectronic Protein Design*
- 300** *Computer-Aided Formulation Design for Improved Nanoformulation of Anticancer Therapeutics: Using Simulation to Improve the Effectiveness of Cancer Drugs*
- 302** *Factor Separation Analysis of Urban Impacts on a Simulated Supercell*
- 304** *General Relativistic Neutrino Radiation Transport: Understanding the Driving Force Behind Cosmic Explosions*
- 306** *Beyond Navier–Stokes and Molecular Dynamics: Understanding Biomacromolecular Dynamics through the Development of Multiscale Hybrid and Hydrodynamic Simulation*
- 308** *The Impacts of Hydrometeor Centrifuging on Tornado Dynamics: Improving the Realism of Tornado Simulations*
- 310** *Dynamics of Colloidal Glass-Formers*
- 312** *Magnetic Reconnection in Laser-Driven Plasmas: From Astrophysics to the Laboratory in Silico*
- 314** *Energetic Dynamics of Ocean Basin Model with Rotation, Surface Wind, and Buoyancy Forcing*
- 316** *Genomic Perspectives on the Amphibian Tree of Life*

QUANTIFYING FOREST DROUGHT RESISTANCE

Elizabeth Agee, University of Michigan
2016-2017 Graduate Fellow

EXECUTIVE SUMMARY

Changing patterns of precipitation and rising temperatures have highlighted the vulnerability of tropical rainforests to water limitation. Current land surface models are unable to fully capture the plasticity of the forest to water limitation caused by El Niño-induced droughts. This work focuses on examining the role of root water uptake to individual and community drought resilience. Using a modified version of the PFOTRAN model, we simulated root water uptake for more than 3,000 individuals located within the eastern Amazon River basin. Because the below-ground nature of root systems prevents *in situ* observation, an ensemble of simulations has been developed that test various plant traits including various structures and hydraulic functions. Analysis is ongoing, but preliminary data have highlighted the ability of root systems to dynamically shift uptake when shallow layers of the soil water supply have been exhausted.

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 that were triggered by strong shifts in sea surface temperature caused by the El Niño–Southern Oscillation. 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. 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. We modeled root architectures that represent the structural and spatial distribution of roots using the open source RootBox model [3]. Each tree system was assigned hydraulic parameterization (e.g., root hydraulic conductivity and water potential thresholds) based on statistically generated water usage strategies. Water usage 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 was coupled with the massively parallel flow and transport model, PFLOTRAN [4], using hybridization techniques [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, we employed ensemble modeling approaches 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, located in the eastern Amazon River Basin.

RESULTS & IMPACT

We have crafted several different model scenarios that explore the contributions of various root system traits. Because root systems are obscured by the soil from direct observation, we relied

on ensemble modeling techniques that allowed us to test various combinations of traits. For example, we tested the response of water uptake to various scenarios of rooting depth and lateral spread. Rooting depth is the maximum depth the root system grows, and lateral spread is its horizontal extent from the base of the tree. Both traits affect the volume of soil trees can extract resources from, but they also incur construction and maintenance costs. Benefits gained from increasing lateral extent can be limited by competition from neighboring trees. Fig. 1 highlights the canopy structure (a) and examples of different rooting depths and lateral spreads tested (c and d). These traits are modeled individually and concurrently for both wet and dry seasons (b) to highlight shifts in water uptake during water limitation.

Simulations encompass more than 3,000 individuals of varying size and water demands. Preliminary results highlight the contribution of root traits to both individual and community integrity. Analysis is ongoing, but this work represents one of the largest modeling studies of three-dimensional root water uptake ever attempted. Results from this work can enhance next-generation earth system models by identifying key traits for water uptake processes.

WHY BLUE WATERS

Blue Waters is critical to the ongoing success of this project. Simulations of this complexity and scale require the computational power of this system to make meaningful analyses. Not only are the simulation domains complex, multiple simulations are needed to account for system uncertainty. 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 fifth-year PhD student studying environmental engineering. She is working under Dr. Valeriy Ivanov at the University of Michigan and plans to graduate in the summer of 2019.

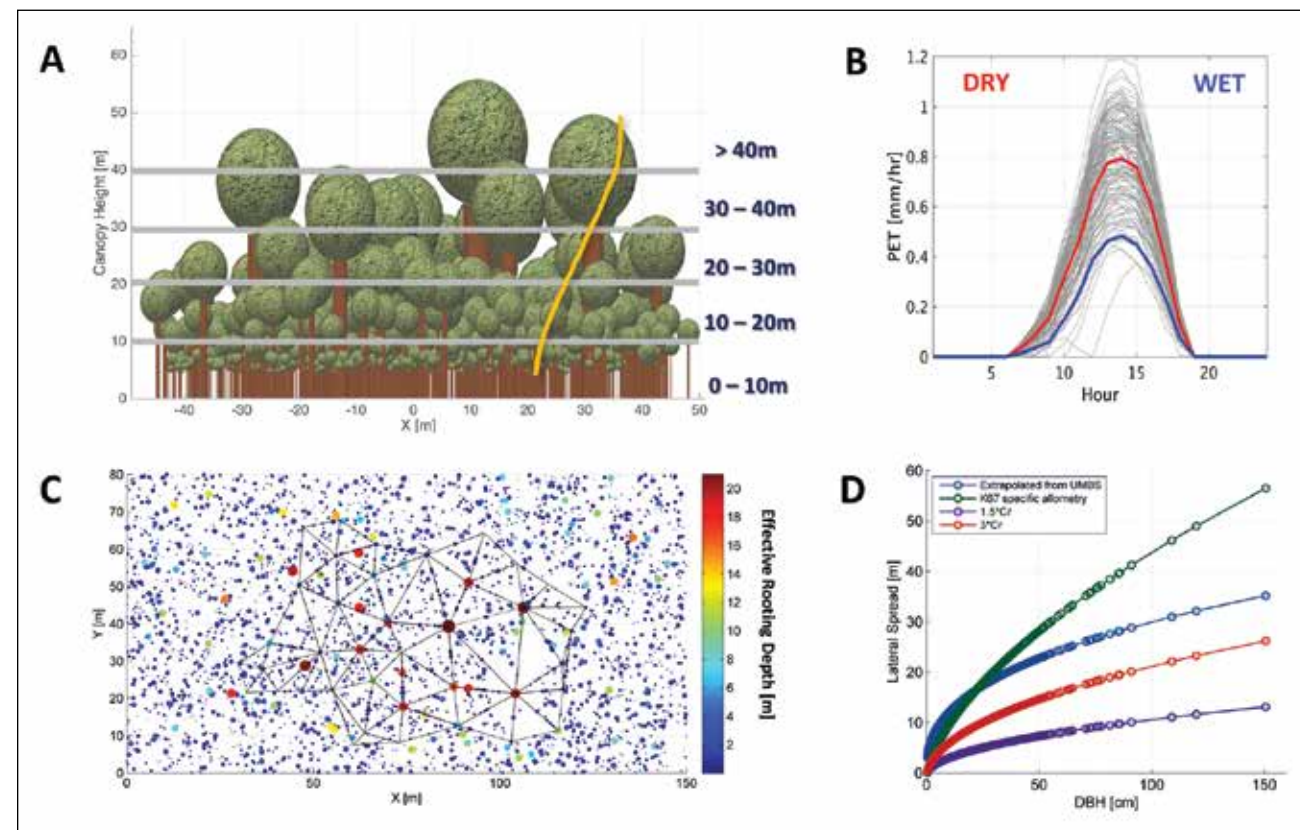


Figure 1: Canopy structure of monitored trees within the Tapajós National Forest. (b) Mean daily potential evapotranspiration for both wet (blue) and dry (red) seasons. (c) Dots represent individual trees in the domain with color indicating rooting depth. (d) Scenarios of lateral spread, which vary with tree diameter at breast height (DBH).

MULTISCALE BIOPHYSICAL INTERACTIONS: INERTIA, GROWTH, AND SARGASSUM SEED POPULATIONS

Maureen Brooks, University of Maryland Center for Environmental Science Horn Point Laboratory
2015–2016 Graduate Fellow

EXECUTIVE SUMMARY

Sargassum is a macroalgae, or seaweed, that spends its entire life cycle floating on the surface of the Atlantic Ocean. This makes it an ideal model organism for understanding mesoscale biological–physical interactions. It is also of interest because it supports a unique open ocean ecosystem, but also has negative consequences for coastal communities when it washes ashore. This study examines how inertial forces influence the pathways *Sargassum* follows across the Atlantic and its access to nutrients vital for growth. This research uses a coupled model system that bridges scales from nutrient chemistry up through basinwide ocean circulation. Blue Waters resources allowed for model implementation at high spatial and temporal resolution. Results suggest that inertial forces enhance *Sargassum* connectivity between the tropics and the Gulf of Mexico, alter rates of entrainment in eddies, and lead to patchier growth.

RESEARCH CHALLENGE

Pelagic *Sargassum* is unique in the ocean. It is the only macroalgae that does not have a benthic-attached stage in its life cycle. It can be found throughout the tropical and subtropical Atlantic Ocean and serves as a keystone species [1]. In recent years, it has also washed up on beaches in increasing quantities, where it poses economic challenges, causes human health concerns, and can be harmful to wildlife [2]. My previous Blue Waters-supported work has suggested the existence of nursery regions in the Gulf of Mexico and western tropical Atlantic that exert strong influence on the basinwide distribution of *Sargassum*. However, the mechanisms that allow *Sargassum* to be retained in these regions and seed the rest of the basin were not understood.

Recent studies of Lagrangian particles have shown that trajectories can shift when inertial interactions due to particle size and density are considered [3]. Buoyant *Sargassum* rafts are also more likely to become entrained in eddies where they can

potentially experience increased nutrient availability due to a shift in the nutricline brought on by eddy pumping. This work compares *Sargassum* growth and trajectories in simulations with and without these inertial interactions to determine their impact on the distribution of *Sargassum* in the Atlantic.

METHODS & CODES

A system of four coupled models was developed to simulate *Sargassum* growth and transport. A 1/12°-resolution Hybrid Coordinate Ocean Model (HYCOM) [4] implementation was run over an Atlantic domain from 15°S to 64°N and 100°W to 15°E. The vertical structure includes 28 hybrid-coordinate layers that can transition between sigma- and z-coordinates, with increased resolution near-surface to better capture dynamics relevant to *Sargassum* physiology. A biogeochemical model adapted from the work of Fennel [5] captures planktonic ecosystem dynamics and biologically mediated nutrient cycling.

A Lagrangian particle model tracks the trajectories of individual *Sargassum* colonies, and a *Sargassum* physiology model calculates growth and mortality within each particle. For this study, I have modified the HYCOM particle-tracking code to allow for buoyant *Sargassum* and to sample and interpolate the physical and biological parameters from the circulation and biogeochemical models that are necessary for calculating *Sargassum* growth. Inertial effects were implemented using a simplified Maxey–Riley equation applicable for a quasigeostrophic flow [3]. Lagrangian particle motion accounts for finite particle size and the difference in density between the *Sargassum* particles and ambient seawater.

The *Sargassum* physiology model uses light, temperature, and nutrient availability to determine growth rate. Model *Sargassum* reproduces via vegetative propagation, where a new *Sargassum* particle can be initialized in place when conditions are favorable. The coupled model estimates of *Sargassum* biomass were validated using monthly satellite climatologies derived from 10 years of observations [6].

RESULTS & IMPACT

Model experiments comparing *Sargassum* growth and distribution in inertial versus noninertial particles highlight several

new biological–physical interactions. Including inertial forces due to the size and density of individual rafts significantly alters the pathways by which *Sargassum* moves between regions of the Atlantic. Connectivity analysis showed year-round reductions of almost 50% in transport from the tropics to the Sargasso Sea. Instead, *Sargassum* that started in the tropics either remained there or was transported into the Caribbean and Gulf of Mexico. This is consistent with previous work that showed the strong influence of the western tropics and Gulf of Mexico on the overall *Sargassum* distribution.

Analysis of the *Sargassum* physiology also yielded new insights. When modeled as noninertial water parcels, *Sargassum* experiences moderate growth conditions across most of its range (Fig. 1a). In contrast, accounting more realistically for the physical properties of *Sargassum* rafts yields patchier growth rates as the rafts get entrained in eddies. It also leads to local conditions of higher growth in the western Gulf of Mexico, the Caribbean, and the tropics.

Together, these results point to the tropics as the major potential source for the *Sargassum* responsible for inundation events in the Caribbean and give an additional mechanism leading to the high biomass that makes them so difficult to manage. Understanding this pathway and the oceanographic conditions that drive it is key to being able to predict and eventually mitigate these costly events.

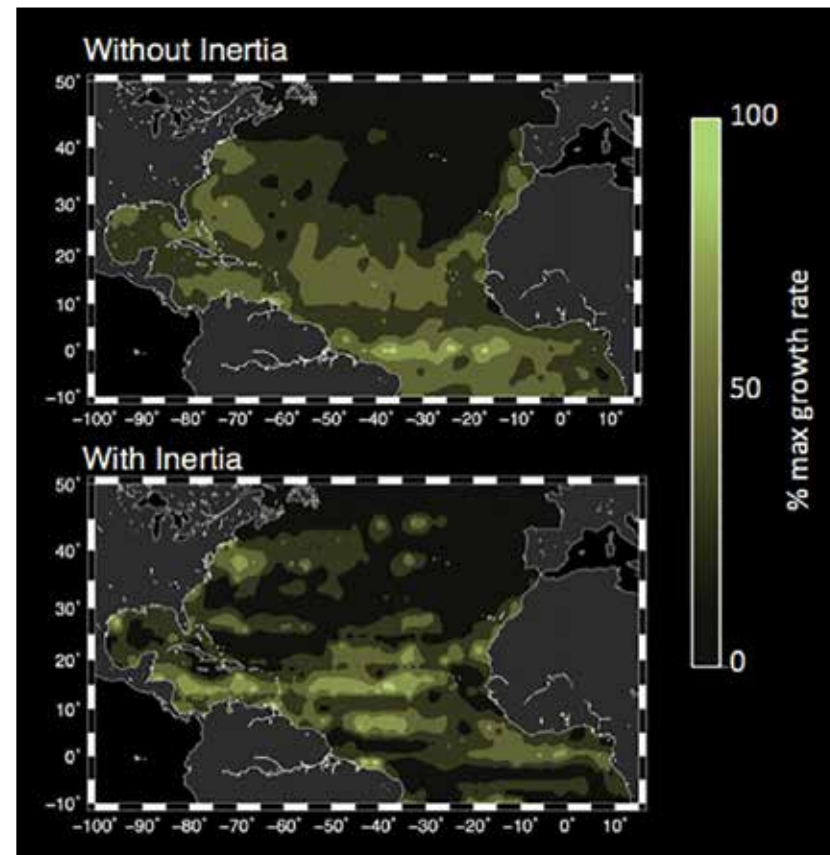
WHY BLUE WATERS

The scope of this project is reliant on Blue Waters. This study combines high-resolution ocean circulation with ocean biogeochemistry, Lagrangian particle tracking, and individual-based organism physiology across orders of magnitude of spatial scales. Blue Waters' performance helped accommodate the high computational cost of running this complex system of models. In addition, the responsiveness and professionalism of the NCSA staff has been of great value in implementing this code on Blue Waters.

PUBLICATIONS & DATA SETS

Brooks, M., et al., Factors controlling the seasonal distribution of pelagic *Sargassum*. *Mar. Ecol. Prog. Ser.*, in revision (2018).

Figure 1: *Sargassum* growth, shown as % maximum growth rate, from model *Sargassum* particles. (a) Control; (b) Inertial experiment. Note patchiness and zonal spreading consistent with entrainment in eddies.



Maureen T. Brooks is in the fifth year of a Marine–Estuarine–Environmental Sciences PhD 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.

SIMULATING THE CIRCUMGALACTIC MEDIUM: THE NEED FOR ROBUST COSMIC RAY TRANSPORT

Iryna Butsky, University of Washington
2016–2017 Graduate Fellow

EXECUTIVE SUMMARY

The majority of galactic baryons reside outside of the galactic disk in the diffuse gas known as the Circumgalactic Medium (CGM). While state-of-the-art simulations excel at reproducing galactic disk properties, many struggle to drive strong galactic winds or to match the observed multiphase structure of the CGM with thermal supernova feedback. To remedy this, recent studies have included nonthermal cosmic ray (CR) stellar feedback prescriptions to drive strong outflows and to better match observed low-ion column densities in the CGM. However promising, these results depend strongly on the choice of CR transport mechanism and its constant parameter values, thus weakening the predictive power of such simulations. This work uses a suite of simulated isolated disk galaxies to demonstrate that the invoked approximation of CR transport affects the predicted temperature and ionization structure of the CGM and motivates the need for a detailed parameter study.

RESEARCH CHALLENGE

The majority of baryonic matter exists as the diffuse gas that surrounds galaxies, commonly referred to as the CGM. Recent observations, such as those from the Cosmic Origins Spectrograph on board the Hubble Space Telescope, have revolutionized our

understanding of the CGM. The new picture of the CGM is that of a dynamic, multiphase gas shaped by the interplay of metal-rich galactic outflows and pristine inflows from the intergalactic medium. The inferred gas properties deduced from absorption features in the spectra paint a picture of a CGM with gas clouds spanning wide ranges of temperatures and pressures, spatially coexisting [1].

While state-of-the-art simulations excel at modeling galactic disk structures, they struggle to reproduce the observed ionization structure of the CGM. Particularly puzzling is the presence of two gas states out of pressure equilibrium: low-density, cool (10⁴K) and warm gas that originates in the disk and hot (10⁶K) gas native to the halo. A possible explanation for this is that the CGM is constantly being replenished by strong galactic winds and is being collisionally ionized by cosmic rays therein.

Cosmic rays are charged particles that have been accelerated to relativistic velocities in extreme shocks such as supernovae. As they travel around magnetic field lines, the induced current exerts an effective pressure on the thermal gas. Simulations of galaxy evolution that include cosmic rays in supernova feedback have been more successful at driving strong outflows and reproducing the observed CGM ionization structure than those with purely thermal supernova feedback [2]. However promising, these simulations lack predictive power because the parameters of simulating cosmic ray physics are poorly constrained. In fact, there is no consensus on the preferred approximation of cosmic ray transport: streaming or diffusion.

METHODS & CODES

Bulk cosmic ray motion is simulated as a relativistic fluid that is coupled to the thermal gas. In this approach, cosmic ray motion relative to the gas is restricted to either diffusion or streaming. Physically, this choice depends on the source of turbulence in the magnetic field lines that is scattering the cosmic rays. During my time as a Blue Waters Graduate Fellow, I implemented cosmic ray diffusion and streaming in the astrophysical simulation code Enzo [3]. We used Enzo to simulate a suite of isolated Milky Way-type galaxies in which supernova feedback injects 10% of its energy as cosmic ray energy. The galaxy models differ in their cosmic ray transport mechanism.

RESULTS & IMPACT

The CGM in the isolated galaxy models is enriched solely by the outflows that expel gas from the disk. Models with cosmic ray feedback drive strong galactic winds that populate the CGM




Figure 2: The distribution of the ratio of cosmic ray pressure to thermal pressure as a function of spherical radius from the galaxy center. The pixels are colored by the density of the thermal gas.

with metals (Fig. 1). However, the cosmic ray transport models predict different ionization structures for the warm and cool gas components. Only the galaxy model with cosmic ray streaming develops the patchy, multiphase CGM structure inferred from observations.

The primary differences between the galaxy models stem from the distribution of cosmic ray pressure, which is responsible for driving outflows and for altering the temperature and ionization structure of the CGM (Fig. 2). In the diffusion model, cosmic ray pressure is most influential near the disk, at high gas densities. In the streaming model, the cosmic ray pressure support is most effective at large radii and low gas densities. This fundamental difference between the two CR transport prescriptions cannot be reconciled by trivially altering constant parameter values of CR transport.

Cosmic rays are observed to be a dynamically important component of galaxies [4] and likely play a significant role in shaping the observed ionization structure in the CGM. However, in order for galaxy simulations with cosmic ray physics to hold predictive power, they must first develop robust models of cosmic ray transport.

WHY BLUE WATERS

Although the CGM spans hundreds of kiloparsecs (1kpc = 3.09x10¹⁶ km), its structure is shaped by gas interactions on subparsec scales. Even with adaptive mesh refinement, capturing this dynamic range of physical and temporal scales is a formidable task and galaxy simulations have yet to resolve the CGM in such detail. Furthermore, realistic simulations of cosmic ray transport must include magnetic fields, which are computationally expensive compared to purely hydrodynamic simulations. Therefore, these simulations require the use of massively parallel, high-performance supercomputers such as Blue Waters. In addition to Blue Waters' computational resources, I have benefited from the help of the User Support team and my point of contact, Roland Haas.

PUBLICATIONS & DATA SETS

Butsky, I., and T. Quinn, The Role of Cosmic Ray Transport in Shaping the Simulated Circumgalactic Medium. Submitted to *The Astrophysical Journal* (2018), arXiv:1803.06345




Figure 1: The metallicity (metal-enriched gas that traces galactic outflows) and the line-of-sight densities of H I and O VI (tracing cool and warm gas respectively) for galaxy models with differing CR transport prescriptions, taken after 13 Gyr of evolution.

Iryna Butsky is a third-year PhD student in astronomy. She is working under the supervision of Tom Quinn at the University of Washington, and she plans to graduate in 2021.

HOW THE SOLAR SYSTEM SURVIVED A VIOLENT EPOCH OF DYNAMICAL INSTABILITY

Matthew Clement, University of Oklahoma
2017–2018 Graduate Fellow

EXECUTIVE SUMMARY

The solar system’s outer planets (Jupiter, Saturn, Uranus, and Neptune) formed in just a few million years, while gas was still present in the Sun’s primordial protoplanetary disk [1]. Although the evolution of these outer planets is well studied [2,3], the leading models seem to be incompatible with the solar system’s terrestrial system (Mercury, Venus, Earth, and Mars) [4]. By performing thousands of N-body simulations of an orbital instability in the outer solar system occurring in conjunction with terrestrial planet formation, we presented the first complete evolutionary model for the solar system that explains both its inner and outer regimes. Additionally, we conducted the largest-ever suite of planet-formation simulations using a realistic model that accounts for the effects of fragmentation as bodies collide [5]. Finally, we used a GPU-accelerated code [6] to accurately model dynamics down to realistic mass resolutions during the solar system’s earliest epoch, and in the young asteroid belt.




Figure 1: Semi-Major Axis vs. orbital eccentricity plot depicting the evolution of a successful system. The size of each point corresponds to the mass of the particle; because Jupiter and Saturn are hundreds of times more massive than the terrestrial planets, we use separate mass scales for the inner and outer planets.

RESEARCH CHALLENGE

Accurately modeling the late stages of planet accretion is subject to numerical limitations and simplifications. In particular, to keep the calculation tractable, most authors [7,8] employ integration schemes that neglect collisional fragmentation. The initial planet-forming disk, which in reality contained millions of solid objects with a range of masses, must be approximated with just over a thousand bodies (the majority of which are assumed not to interact gravitationally with one another). Nevertheless, such studies have proved successful at replicating the general orbits of the inner planets. However, explaining Mars’ small mass (just 10% that of the Earth) and rapid formation (about ten times faster than the Earth, as inferred from isotopic dating [9]) requires substantial modification to the standard theory of planet formation [8,10]. Furthermore, the asteroid belt’s low total mass (only a few percent that of the moon) and unique dynamical structure are still largely unexplained [7,8,10]. Earth and Mars are both in the Sun’s “potentially habitable” zone, yet Mars is small, barren, and unable to support a robust atmosphere. Understanding the dynamical mechanisms that prevented Mars from growing into an Earth-like planet will give us insight into how special our own world really is.

METHODS & CODES

For our fragmentation simulations, we use a modified version of the Mercury6 hybrid integrator, written in Fortran [11,13]. Our simulations of terrestrial accretion begin with the simplest initial conditions and are consistent with observations of protostellar disks [1,7,8]. To systematically test the effects of a giant planet instability, we perform several batches of integrations and trigger the instability during different epochs of terrestrial growth. To investigate the effect on the asteroid belt, we use a GPU code written in CUDA C (GENGA, [6]) to reperform successful integrations with a larger number of objects in the belt region. Because gas-disk interactions are complex, the exact mass and planetesimal size distribution profiles that emerge from the primordial gas (and go on to form the inner planets) are not well known. Therefore, our “simple” initial conditions might not be representative of physical reality. We investigate this problem as well, using GPU acceleration. We employ a forcing function to mimic the effects of gas drag and utilize a multiannulus approach to track the accretion of millions of small objects in the infant terrestrial disk.




Figure 2: Semi-Major Axis vs. orbital inclination plot comparing the actual asteroid belt (top) with our simulated asteroids (bottom). The dashed lines denote important orbital resonances. Grey points correspond to asteroids on Mars-crossing orbits, which will be removed naturally during subsequent evolution.

RESULTS & IMPACT

Our work offers a simple and elegant explanation for Mars’ small size and rapid growth. Our instability simulations consistently outperform our control runs when measured against a variety of success criteria. These include the orbits, masses, and geological formation timescales of the planets; the size and orbital structure of the asteroid belt; and the water content of Earth. When the instability occurs approximately 1–10 million years after gas disk dispersal, “Mars” is just one of several Mars-sized objects with similar orbits (Fig. 1). The frequent perturbations from the increasingly eccentric orbits of Jupiter and Saturn quickly cause these objects (and similar bodies in the asteroid belt) to either be ejected from the system or scatter inward towards the proto-Earth (sometimes to deliver water). In successful simulations, Mars undergoes no further accretion events after the instability, while Earth and Venus continue to grow (thus matching their relative geological formation times [9]). Additionally, we find that accounting for collisional fragmentation results in fully grown systems of terrestrial planets that are better matches to the actual

solar system in terms of their orbital excitation (eccentricities and inclinations) and planet spacing (particularly that of Earth and Venus). Furthermore, the instability proves successful at reproducing the broad orbital distribution of large objects in the asteroid belt (Fig. 2). Thus, an early dynamical instability among the giant planets can simultaneously explain the structure of both the inner and outer solar system.

WHY BLUE WATERS

Blue Waters’ unique capabilities and highly supportive staff were crucial to the success of this project. Our study relied heavily on the use of GPU accelerators on Blue Waters’ XK nodes. Having the ability to efficiently run large suites of GPU-accelerated jobs led us to seek out a Blue Waters allocation. Furthermore, our initial work on Blue Waters has spurred several follow-on projects and greatly accelerated my progress in graduate school.

PUBLICATIONS & DATA SETS

Clement, M., et al., Mars’ growth stunted by an early giant planet instability. *Icarus*, 311 (2018), DOI:10.1016/j.icarus.2018.04.008.

Matthew Clement is a third-year PhD candidate in astrophysics at the University of Oklahoma. He is working under the supervision of Nathan Kaib and expects to receive his doctorate in August 2019.

SEDIMENT TRANSPORT IN ESTUARIES: ESTIMATING BED SHEAR STRESS DISTRIBUTION FROM NUMERICALLY MODELED TIDES IN AN ENERGETIC ESTUARY

Salme Cook, University of New Hampshire
2017–2018 Graduate Fellow

EXECUTIVE SUMMARY

Ocean tides constitute a significant forcing mechanism and are responsible for the transport of salt, sediment, nutrients, and pollutants in most coastal environments. Estuaries are the areas where rivers flow into the ocean and an increase in human population density in their vicinity has a significant anthropogenic impact on them. This impact, in turn, alters the productivity of estuarine environments and results in increased nutrient loading and amplified suspended sediment that reduces water quality. To accurately predict sediment transport, a detailed understanding of the bed shear stress that drives sediment erosion, suspension, and deposition is essential. In this work, we implemented and verified a high-resolution three-dimensional coupled hydrodynamic–wave–sediment transport numerical model for a tidally dominated estuary located in the U.S. Gulf of Maine. When paired with observational data sets, the model successfully predicted the shear stress distribution from the tidal channels across the mudflat.

RESEARCH CHALLENGE

More than half of the world’s population lives within 50 miles of a coast, and two-thirds of the world’s largest cities are located near estuaries [1]. Estuaries are fertile ecosystems that provide both invaluable ecosystem services and considerable economic benefits to society. The rise in land development and associated increases in impervious surface cover have led to a decline in water quality and estuarine health by depositing higher loads of




Figure 1: The bathymetry and topography of the Piscataqua River–Great Bay estuary for the 30-meter grid.

sediment, nutrients, and pollutants. In order to guide legislative policies and conservation efforts, scientists and environmental managers are relying increasingly on computer models that represent the hydrodynamics, sediment transport, and associated biogeochemical fluxes within these systems. Numerical modeling is a cost-effective way to predict the impact of sediments on estuary-wide nutrient loads, a potentially significant but so far largely neglected source of nutrients.

In this project, we applied a coupled ocean-sediment transport model to a New England estuary to resolve the importance of tidally induced sediment transport of fine-grained sediment, specifically in estuarine mudflats. Both the suspension and resuspension of fine-grained material are driven by the critical shear stress that is strong enough to overcome the cohesive forces at the bed. Tidal currents are a relatively well-understood forcing mechanism. However, in combination with wind-driven circulation and the wetting and drying of mudflats during a tidal cycle, tidal currents create spatial and temporal variabilities that are hard to sample in the field.

METHODS & CODES

In this work, we used the Regional Ocean Modeling System (ROMS) [2] within the Coupled–Ocean–Atmosphere–Wave–Sediment Transport coupled model framework [3,4]. What makes the estuary we simulated unique is the high-resolution bathymetry and topography used to create the model grid (Fig. 1) as well as extensive observational data sets (1972–present) available for verification of modeled flows. Our first goal was to verify the 30-meter model grid for tides and meteorological forcing using different bottom boundary conditions that directly affect the estimates of shear stress and, thereby, sediment transport. We tested four different bottom boundary conditions for purely tidal and for tidal plus meteorological forcing. We ran each model for 30 days, saving model state data at five-minute intervals. Using time series and statistical analyses, we determined the best-fit bottom boundary condition for the next model runs. Differences between simulations using models with and without meteorological oscillations were found to be negligible, suggesting that interactions among the tides and other low-frequency (weather-forced) mean flows are weak and can be ignored when considering tidal dynamics. This result directly informs the next set of models that will include wind-driven circulation that might have an episodic but nonetheless significant effect on the shear stress distribution within tidally modulated estuaries.

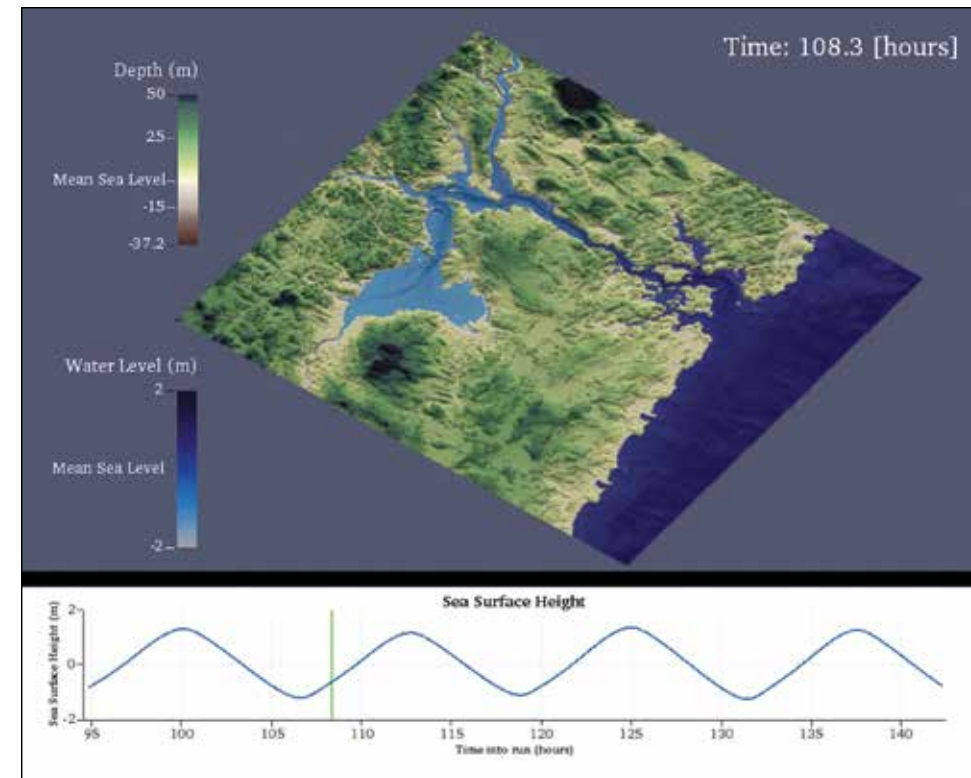


Figure 2: The Piscataqua River–Great Bay estuary at mid-tide—108.3 hours into the model run. Darker blue colors indicate higher water levels than lighter blue colors. The differences in the tidal level from the Atlantic Ocean to the estuary drive large tidal currents that are important for sediment transport.

RESULTS & IMPACT

Comparison of simulation results to experimental observations suggests that the model reproduces the nonlinear tidal behavior and accurately describes the sea surface elevations and velocity throughout the estuary. This result supports the estimates of shear stress distribution, sediment transport, and nutrient fluxes caused by the tidal forcing. Our results suggest that nutrient fluxes form sediment during a typical tidal cycle that are significant and should not be neglected when estimating nutrient loads in estuaries with strong currents and tidal mudflats. Scientists, land managers, and legislators should be aware of this finding when funding projects and determining best practices for estuarine management.

Modeled currents have also aided in the planning of several field studies, particularly the timing and location of deploying instrumentation, which in the past have been difficult to plan efficiently. Future work will include a comparison of a coarse (30-meter) and fine (10-meter) grid to determine how grid resolution affects estimated shear stress distributions. Since grid resolution is directly proportional to computational needs, this is important for estimating computational requirements for such projects. This could potentially lead to better estimates of shear stress, sediment transport, and nutrient loading, particularly

for estuaries with small-scale bathymetric features like the tidal channels in this estuary.

With the assistance of the NCSA Visualization Group, the modeled tidal currents and shear stress have been visualized in new and exciting ways (see Figs. 1 and 2). These figures and movies have captured the interest of groups around campus, local scientists, and government agencies.

WHY BLUE WATERS

The major limitations to this research project in the past have been the accessibility of computational resources to resolve these processes at the necessary temporal and spatial scales, and availability of observational data to verify model results. The Blue Waters system provided the required computational power to test models using a higher-resolution 10-meter grid, which was previously infeasible. Further, the project support staff were an invaluable asset in getting this project up and running on Blue Waters.

PUBLICATIONS & DATA SETS

Cook, S., T.C. Lippmann, and J.D. Irish, Modeling nonlinear tidal evolution in an energetic estuary. *Ocean Modeling*, under review (2018).

Salme Cook is a sixth-year PhD student in oceanography at the University of New Hampshire. She is working under the supervision of Tom Lippmann and expects to graduate in May 2019.

MACHINE LEARNING HARNESSSES MOLECULAR DYNAMICS TO DEVELOP THERAPEUTIC STRATEGIES FOR ALZHEIMER'S DISEASE AND CHRONIC PAIN

Evan N. Feinberg, Stanford University
2017–2018 Graduate Fellow

EXECUTIVE SUMMARY

The arc of drug discovery entails a multiparameter optimization problem spanning vast length scales. The key parameters range from solubility (angstroms) to protein–ligand binding (nanometers) to *in vivo* toxicity (meters). Through feature learning—instead of feature engineering—deep neural networks promise to outperform both traditional physics-based and knowledge-based machine learning models for predicting molecular properties pertinent to drug discovery. To this end, we developed the PotentialNet family of graph convolutions. These models are designed for and achieve state-of-the-art performance for protein–ligand binding affinity. We further validated these deep neural networks by setting new standards of performance in several ligand-based tasks. Finally, we introduced a cross-validation strategy based on structural homology clustering that can more accurately measure model generalizability, which crucially distinguishes the aims of machine learning for drug discovery from standard machine learning tasks.

RESEARCH CHALLENGE

Most FDA-approved drugs are small organic molecules that elicit a therapeutic response by binding to a target biological macromolecule. Once bound, small molecule ligands either inhibit the binding of other ligands or allosterically adjust the target's conformational ensemble. Binding is thus crucial to any behavior of a therapeutic ligand. To maximize a molecule's therapeutic effect, its affinity—or binding free energy—for the desired targets must be maximized while simultaneously minimizing its affinity for other macromolecules. Historically, scientists have used both cheminformatics- and structure-based approaches to model ligands and their targets, and most machine learning approaches use domain expertise-driven features.

More recently, deep neural networks (DNNs) have been translated to the molecular sciences. Training most conventional DNN architectures requires vast amounts of data. For example, ImageNet currently contains over 14 million labeled images. In contrast, the largest publicly available data sets for the properties of

drug-like molecules include PDDBind 2017, with a little over 4,000 samples of protein–ligand co-crystal structures and associated binding affinity values; Tox21 with nearly 10,000 small molecules and associated toxicity endpoints; QM8 with around 22,000 small molecules and associated electronic properties; and ESOL with a little over 1,000 small molecules and associated solubility values. This scarcity of high-quality scientific data necessitates innovative neural architectures for molecular machine learning.

METHODS & CODES

In this project, we generalized a graph convolution to include both intramolecular interactions and noncovalent interactions between different molecules. In particular, we described a staged gated graph neural network, which distinguishes the derivation of differentiable bonded atom types from the propagation of information between different molecules. We implemented the models in PyTorch, a cutting-edge deep-learning framework. We trained and evaluated our models on publicly available data sets, including Tox21 for toxicity, ESOL for solubility, and PDDBind for protein–ligand affinity.

RESULTS & IMPACT

Spatial Graph Convolutions exhibit state-of-the-art performance in affinity prediction. Whether based on linear regression, random forests, or other classes of DNNs, RF-Score, X-Score, and TopologyNet are machine learning models that explicitly draw upon traditional physics-based features. Meanwhile, the Spatial Graph Convolutions presented here use a more principled deep-learning approach. Input features are only basic information about atoms, bonds, and distances. This framework does not use traditional hand-crafted features

such as hydrophobic effects, π -stacking, or hydrogen bonding. Instead, higher-level interaction “features” are learned through intermediate graph convolutional neural network layers. In light of the continued importance and success of ligand-based methods in drug discovery, we benchmarked PotentialNet on several ligand-based tasks: electronic property (multitask), solubility (single task), and toxicity prediction (multitask). We observed statistically significant performance increases for all three prediction tasks. A potentially step change improvement was observed for the QM8 challenge, which also reinforced the value of the concept of stages that privilege bonded from nonbonded interaction.

WHY BLUE WATERS

The Blue Waters supercomputer, in particular the many GPU nodes, as well as the outstanding staff, were quite important in enabling us to run massively parallel hyperparameter searches to train the optimal deep-learning models for drug discovery tasks.

PUBLICATIONS & DATA SETS

Farimani, A.B., E.N. Feinberg, and V.S. Pande, Binding Pathway of Opiates to μ Opioid Receptors Revealed by Unsupervised Machine Learning. *arXiv preprint arXiv:1804.08206* (2018).

Feinberg, E.N., et al., Spatial Graph Convolutions for Drug Discovery. *arXiv preprint arXiv:1803.04465* (2018)

Feinberg, E.N., et al., Machine Learning Harnesses Molecular Dynamics to Discover New μ Opioid Chemotypes. *arXiv preprint arXiv:1803.04479* (2018). Feinberg, E.N., V.S. Pande, A.B. Farimani, and C.X. Hernandez, Kinetic Machine Learning Unravels Ligand-Directed Conformational Change of μ Opioid Receptor. *Biophysical Journal*, 114:3 (2018), p.56a.




Figure 1: Visual depiction of multistaged spatial gated graph neural network. Stage 1 entails graph convolutions over only bonds, which derives new node (atom) feature maps. Stage 2 entails both bond-based and spatial distance-based propagation of information. In the final stage, a graph gather operation is conducted over the ligand atoms.

As a fifth-year PhD candidate in biophysics at Stanford University, Evan N. Feinberg worked under the direction of Vijay S. Pande and Kerwyn C. Huang. He received his degree in September 2018.

SCALING RELATIONSHIPS ACROSS MODELING RESOLUTIONS IN MOUNTAIN HEADWATERS: UNDERSTANDING CLIMATE CHANGE IMPACTS ON ROCKY MOUNTAIN HYDROLOGY IN A NUMERICAL MODELING CONTEXT

Lauren Foster, Colorado School of Mines
2017–2018 Graduate Fellow

EXECUTIVE SUMMARY

One in 10 Americans sources water from the Colorado River [1], and 85% of that streamflow is generated in Rocky Mountain headwater catchments [2] that have been shown to be especially sensitive to a changing climate. Global climate models and regional hydrologic models are known to perform poorly in these regions [4–6], casting doubt on water supply forecasts for the next century. In this work, we developed relationships across scales to better model climate change impacts on headwater hydrology. First, we described a new method to parameterize high- and hyper-resolution models where traditional calibrations are infeasible. Second, we demonstrated that high-resolution simulations show more sensitivity to climate changes, indicating that the coarse-resolution models used now may overpredict future water supply. Finally, we compared extensive remote-sensing, point, and field observations to modeling predictions of snowpack, allowing for cross-validation across spatial and temporal scales.

RESEARCH CHALLENGE

Mountain regions are complex, causing nonlinear interactions between water and energy fluxes that are difficult or impossible to model with simple algorithms. Understanding how these interactions scale across modeling resolutions is critical to improving predictions of water supplies under climate change. This problem is important to municipalities, farmers, and water managers across the Southwest United States who depend on water supplies generated in Rocky Mountain headwaters. Local and federal governments need to develop new policies to cope with snowpack reductions in the Rockies. Finally, as hydrologists push toward a new grand challenge of hyper-resolution modeling, new parameterization techniques across scales will be needed.

METHODS & CODES

We modeled a representative headwater catchment in the physically based code Parflow-CLM [7,8] at 1-km and 100-m




Figure 1: Streamflow plotted in the East River and three sub-basins for the initial 100-m (high-resolution) parameterization and the 100-m parameterization after applying the new scaling technique indicate a dramatic improvement in scaling model inputs to match winter baseflow, spring peak flow, and summer monsoon flow.




Figure 2: At higher resolution, larger reductions to streamflow are predicted in response to increasing temperatures. This pattern likely is driven by increases in evapotranspiration at 100 m as temperatures increase. Additionally, temperature increases (black lines) are shown to have a larger effect than precipitation changes (shaded bars), especially for evapotranspiration.

resolution. These models were used to do extensive sensitivity experiments across parameter sets and future climate projections to understand how results and predicated behavior changed across different scales.

RESULTS & IMPACT

Results from the first study (Foster, et al., *Hydrologic Processes*, in review, 2018) have produced a new method that is consistent with the underlying physically based mathematics to scale a critical hydrologic parameter—hydrologic conductivity—across modeling resolutions. These results allow calibration of a coarse-resolution model that can then be applied at scales where calibration is impossible. Given that hyper-resolution models are considered a current “grand challenge” in the field of hydrology [9], these results are critical to both hydrology and physical modeling fields.

The second study (Foster, et al., *Geophysical Research Letters*, in preparation, 2018) demonstrated that climate uncertainty is on a similar scale to modeling uncertainty, highlighting that existing predictive models may overestimate water supplies for the coming century and that modeling uncertainty needs to be included to better predict future water availability.

WHY BLUE WATERS

The Blue Waters project provided the graduate support to make this work possible. Further, it provides one of the best high-performance computing systems for applications, such as this one, that require parallelization. In addition, the consistent advice and guidance of my Blue Waters point of contact expanded my knowledge of high-performance computing systems and applications.

PUBLICATIONS & DATA SETS

- Foster, L.M. and R.M. Maxwell, Using sensitivity analysis and model resolution to scale effective hydraulic conductivity and Manning's n parameters in a mountain headwater catchment. *Hydrologic Processes*, in review (2018).
- Foster, L.M., K.H. Williams, and R.M. Maxwell, When Does Uncertainty Matter While Modeling Climate Change in Mountain Headwaters? Contrasting model resolution and complexity under a changing climate in an alpine catchment. *Geophysical Research Letters*, in preparation (2018).

Lauren Foster is a PhD candidate in the Hydrologic Sciences and Engineering Department of the Colorado School of Mines, working under the supervision of Reed Maxwell. She expected to graduate in August 2018.

ELECTRONIC STRUCTURE OF NiFe OXYHYDROXIDE ACTIVE SITES: UNDERSTANDING AN EFFECTIVE OXYGEN-EVOLVING CATALYST AT A QUANTUM-MECHANICAL LEVEL

Zachary Goldsmith, University of Illinois at Urbana-Champaign
2017–2018 Graduate Fellow

EXECUTIVE SUMMARY

This project aims to characterize the likely active sites in the oxygen-evolving electrocatalyst NiFe oxyhydroxide at a density-functional theory level. Such a characterization includes understanding the redox chemistry of the catalyst thin film, its metal oxidation states, as well as the nature of the frontier bands in its electronic structure. The work involves edge-terminated Ni and NiFe oxyhydroxide films. Experimental research has shown that terminations of the film are sites at which catalytic intermediates are turned over, going through high Fe oxidation states in the process. These studies describe how the electronic structure responds to the broken periodicity of the film. In addition, Fe has been doped into the film at both the interior and exterior layers. Novel results include the energetic preference for Fe at exterior sites, reliable metal oxidation states, and the characteristic sites expected to be electrocatalytically active.

RESEARCH CHALLENGE

Facilitating the redox chemistry between H₂O and O₂ is essential to the development of renewable energy technologies—particularly that of solar fuel cells. Doing so will require understanding and optimizing efficient and earth-abundant electrocatalysts. NiFe oxyhydroxide is among the most active such oxygen-evolving electrocatalysts. While previous studies have well characterized this material *in operando* [1,2], recent experiments have proven the particularly high activity of Fe dopants at the edges of the material [3] and the highly valent nature of such Fe centers [4]. This project aims to determine the electronic and chemical nature of these particularly active NiFe oxyhydroxide edges using frontier quantum-mechanical simulation techniques. This work, if successful, will allow for experimentalists and computationalists alike to rationally design even better catalysts for the oxygen evolution reaction.

METHODS & CODES

We performed density-functional theory (DFT) calculations with a hybrid functional for Ni-only and NiFe oxyhydroxide nanowire supercells using Quantum-ESPRESSO [5]. The material's periodicity was broken in one in-plane direction such that the systems contain different interior and exterior metal sites. We probed different degrees of protonation and hydroxylation of Ni and Fe edge sites as a proxy for the different reactive intermediates that appear throughout the progression of the oxygen evolution reaction. We determined metal oxidation states using maximally localized Wannier functions; the catalytic activity of given sites was determined qualitatively by the projected density of states around the Fermi energy.

RESULTS & IMPACT

The methodology described above has been used previously to determine the redox activity, oxidation states, and electronic structure of Ni and NiFe oxyhydroxide bulk materials under open-circuit and operating conditions [1]. With this computational framework in place, we can gain a better understanding of the nature of the purportedly active edge sites, particularly in the


doped material. For one, there is a ~ 2 eV energetic preference for Fe to dope an exterior site compared to an interior site. This helps explain experimental evidence that the material's overall catalytic activity is optimized only shortly after introducing Fe dopants to the system [3]. In addition, the projected density of states only shows likely catalytically active Fe oxide motifs when the Fe dopant resides at the material edge. The chemical nature of the material's conduction band minimum likely will determine which sites are most catalytically active. Ongoing calculations and analysis will discern the conditions necessary for achieving high-valent Ni- and Fe-centered reaction intermediates.

WHY BLUE WATERS

This work is enabled by Blue Waters given the considerable computational expense yet high level of parallelizability of hybrid functional DFT calculations. These DFT calculations with hybrid functionals are necessary to compute a qualitatively accurate electronic structure of Ni and NiFe oxyhydroxide. In particular, the modeling of heterogeneous materials interfaces is an important task that can only be done with frontier scientific computing resources such as those of Blue Waters.

Zachary Goldsmith is a fourth-year PhD student who expects to receive his degree in 2019 from the University of Illinois at Urbana-Champaign. He is working under the supervision of Sharon Hammes-Schiffer of Yale University.

Figure 1: Illustration of a single layer of NiFe oxyhydroxide with 25% Fe doping solvated in water. Ni is shown in green, Fe in blue, O in red, H on the NiFe slab in orange, and H in solution in white.



REFINING THE CONFORMATIONAL ENSEMBLES OF FLEXIBLE PROTEINS USING SIMULATION-GUIDED SPECTROSCOPY

Jennifer Hays, University of Virginia
2017-2018 Graduate Fellow

EXECUTIVE SUMMARY

Flexible molecular recognition is a common paradigm in immunity and infection; many pathogens have proteins that are structurally flexible and/or highly tolerant of mutations but still effectively bind to human cells. Determining the structural basis of this recognition experimentally is challenging because of the multitude of structures involved. We have developed a computational methodology that provides a way to select the best experiments to measure these structures and subsequently combine them in an integrated model. The conformations of many important flexible proteins are still not well understood; this methodology will allow researchers to characterize systems that are currently too difficult to understand with existing refinement techniques. A systematic approach to refining these flexible receptor–ligand complexes would help elucidate the fundamental physical principles of receptor–ligand binding and promote better drug design for infectious disease.




Figure 1: Iterative refinement of flexible conformational ensemble using simulation-guided spectroscopy. We measured experimentally derived distance distributions and then integrated them into molecular dynamics simulations using restrained-ensemble or bias-resampling ensemble approaches. The resulting hybrid ensemble was then analyzed using mRMR to prospectively identify a set of optimal experiments to perform.

RESEARCH CHALLENGE

Multistructured proteins play critical roles in infectious disease but are often difficult to characterize; experimental techniques often capture either a subset of the structures at high resolution or a more complete set of structures at low resolution. Double electron–electron resonance (DEER) spectroscopy is a powerful tool for measuring multiple structures, but these experiments are low-throughput, which means that it is critical to select only the very best, most informative experiments. We have developed a model-free simulation-based approach for selecting a set of optimal DEER experiments and integrating the resulting data to estimate the full set of structures at high resolution (Fig. 1). This combined experimental and computational approach will allow us to rapidly study and even redesign flexible proteins and thus accelerate the study and treatment of infectious disease.

METHODS & CODES

DEER allows measurement of the distances between nearby pairs of chemically labeled amino acids in a protein. An ideal set of these pairs would have two properties: each selected pair should resolve as many other distances in the system as possible, and each selected pair should resolve distances that are distinct from those determined by each of the other measured pairs. These criteria are optimally satisfied by selecting pairs using the information–theoretic criteria of maximum relevance and minimum redundancy (mRMR) [1].

We incorporated the analyzed DEER data from sets of optimal experiments into our scheme as distance distributions. These distributions are used to drive restrained-ensemble Molecular Dynamics (MD) simulations in which the simulation distance distribution is biased toward the experimental distribution. We performed simulations using a modified version of the restrained-ensemble method described in [2]; code to perform these simulations is available at <https://github.com/kassonlab/restrained-ensemble>.

For DEER-derived distributions with well-separated probability modes, such as would happen for a system with distinct open and closed states, current state-of-the-art methods of incorporation fail to fully integrate the data. As a result, we have also developed a new methodology for integrating the spectroscopic data into MD simulations (Fig. 2).

RESULTS & IMPACT

We tested our iterative refinement methodology on the Neisserial virulence-associated protein Opa60, selecting a set of optimal experiments using the mRMR criteria. We performed those experiments and then incorporated the resulting experimental data into an MD simulation. In comparison to experiments selected according to current structure-guided methods, simulation-guided spectroscopic measurements are significantly more informative. By integrating the data in restrained-ensemble MD, we were able to obtain not only a refined ensemble of apo Opa60 but also structural insight into how Opa60 engages cellular receptors. After two rounds of mRMR-guided refinement, we were able to identify specific residue–residue interaction patterns that were not determined in a structure-guided refinement approach. Further experiments revealed the subset of conformations responsible for binding Opa’s target receptor.

We are currently expanding this iterative methodology to include more robust ways of incorporating distance distributions into MD simulation. In preliminary studies of the open/closed conformational equilibria of syntaxin-1a, a protein involved in SNARE complex formation in synaptic exocytosis, we have successfully and robustly incorporated DEER experimental data that resisted previous methods of incorporation. Specifically, these bias-resampling simulations facilitate the study of transitions between the open/closed conformations, whereas alternative methods prohibit sampling of transitions and, in the case of syntaxin-1a, sampling of the open state.

Many bacterial pathogens improve their ability to evade the human immune response by having proteins that are structurally flexible and/or highly tolerant of mutations, thus preventing recognition by immune cells. However, these proteins are also able to engage receptors on nonimmune cells and trigger infection. By selecting sets of optimal experiments and incorporating these informative results into an estimate of the conformational ensemble, experimentalists can now study biological systems that were once prohibitively complex or expensive. This systematic approach to refining flexible receptor–ligand complexes helps elucidate the fundamental physical principles of receptor–ligand binding and promotes better drug design for infectious disease.

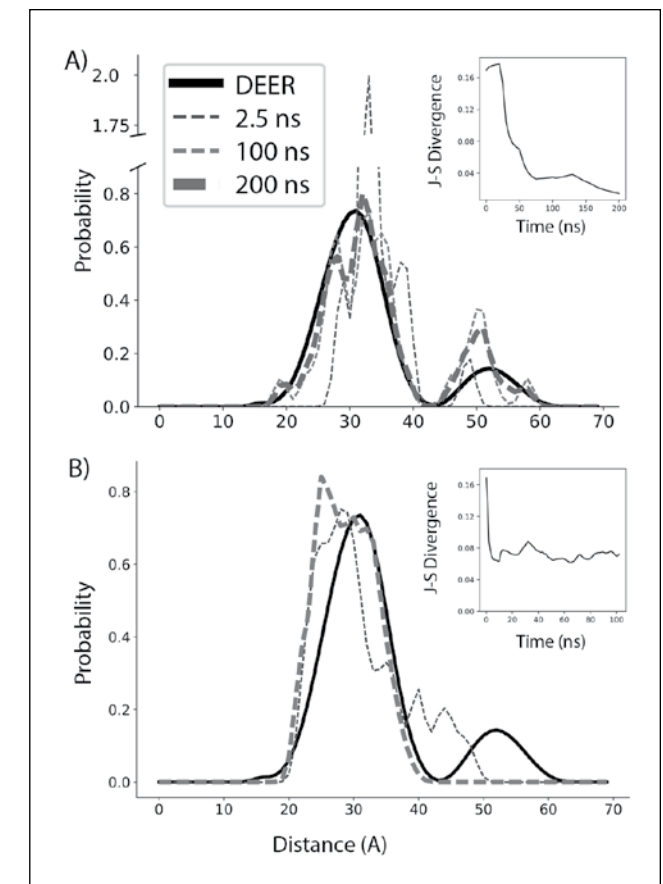


Figure 2: Novel methods of incorporating DEER data facilitate sampling of both the open and closed states of syntaxin-1a. Bias resampling of DEER distributions in ensemble simulations samples both peaks of the distribution (A), while current state-of-the-art methods for incorporating DEER data do not permit sampling of the far peak (B). The difference between the MD and DEER distributions is quantified as Jensen–Shannon divergence in the insets of (A) and (B).

WHY BLUE WATERS

Access to Blue Waters has greatly accelerated the time-to-completion of this project. An enormous amount of molecular dynamics sampling is required to capture the full set of structures of a flexible system; that sampling would not have been possible without an petascale system like Blue Waters. We have run massively parallel MD simulations, scaling to multiple nodes for a single ensemble member and to many ensemble members.

Jennifer Hays is in the fourth year of a PhD program in biomedical engineering at the University of Virginia (UVA). She is working under the direction of Peter Kasson of UVA and Uppsala Universitet and expects to receive her degree in July 2019.

DESIGNING MATERIALS IN FLOW: MULTISCALE SIMULATIONS OF COMPLEX FLUID RHEOLOGY

Michael 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. MPCD is a mesoscale particle-based simulation method for hydrodynamics and is particularly useful for modeling complex fluids and soft matter. Our software will enable studies of complex fluids at length and time scales that would otherwise be inaccessible. We additionally modeled the migration of droplets flowing in microchannels, finding that polymer macromolecules can drive droplets toward the channel center. Such flow-induced focusing of soft materials has important implications for processes ranging from cell sorting to oil recovery.

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.




Figure 1: Snapshot of a droplet in a polymer solution under flow simulated using the DPD method. The different DPD particles are: droplet (orange), polymers (green), solvent (blue), and channel walls (grey).

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 dominates 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

We applied two particle-based mesoscale simulation methods to simulate complex fluids in flow: multiparticle collision dynamics (MPCD) [1,2] and dissipative particle dynamics (DPD) [3,4]. Both MPCD and DPD significantly accelerate simulations of complex fluids compared to explicit-solvent molecular dynamics models. They employ coarse-grained representations of the solvent that still faithfully reproduce relevant physics such as long-ranged hydrodynamic interactions. However, MPCD and DPD models often still require large numbers of particles with simple interactions, naturally inviting a parallel computational approach. We have leveraged the massive parallelism of graphics processing units (GPUs) with the HOOMD-blue [5–7] simulation package to model complex fluids under flow using MPCD and DPD.

RESULTS & IMPACT

We developed the first massively parallel implementation of MPCD for GPUs, which we released as open-source in HOOMD-

blue (version 2.3.0). Flexible initialization and simulation setup are supported through a scriptable Python user interface, while the underlying algorithms have been aggressively optimized to exploit the parallelism of the GPU. Our MPCD software scales efficiently up to 1,024 nodes on Blue Waters. GPU acceleration on the XK nodes exhibited a roughly 3× speedup compared to the XE nodes, which is close to the maximum theoretical performance ratio. The complex fluid and soft matter research communities will significantly benefit from the newly developed MPCD software, which will permit studying processes at physically relevant length and time scales that would be otherwise inaccessible.

We also simulated the migration of droplets in microchannels using DPD (Fig. 1), which is a more mature method than MPCD for modeling multiphase fluids. We found that the addition of polymers to the continuous (solvent) phase induced a migration of the droplets toward the middle of the channel. This flow-induced focusing effect may be exploited during oil recovery to enhance the mobility of oil droplets or used to manipulate soft biological materials such as cells in lab-on-a-chip devices.

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 MPCD software at scale.

PUBLICATIONS & DATA SETS

Howard, M.P., A.Z. Panagiotopoulos, and A. Nikoubashman, Efficient mesoscale hydrodynamics: Multiparticle collision dynamics with massively parallel GPU acceleration. *Comput. Phys. Commun.*, 230 (2018), DOI:10.1016/j.cpc.2018.04.009.

Michael Howard received a PhD in chemical engineering in May 2018 from Princeton University, where he worked under the supervision of Athanassios Z. Panagiotopoulos.

HIGH ACCURACY RADIATIVE TRANSFER IN CLOUDY ATMOSPHERES

Alexandra Jones, University of Illinois at Urbana-Champaign
2014–2015 Graduate Fellow

EXECUTIVE SUMMARY

One of the most important roles clouds play in the atmosphere is in redistributing the radiative energy from the sun and that which is emitted from the Earth and atmosphere. However, radiative transfer in the atmospheric sciences is generally modeled crudely because of the perceived computational expense. Therefore, we developed a highly accurate open-source model that uses Monte Carlo methods to capture the 3D transfer of radiation between the atmosphere and clouds over a broad range of the electromagnetic spectrum.

RESEARCH CHALLENGE

Because of the important role that clouds play in the atmosphere in redistributing the radiative energy from the sun, Earth, and atmosphere as well as the ubiquity of cloud coverage it is imperative that we correctly model the interactions between clouds and radiation in order to accurately predict and observe weather and climate. However, modeling of radiative transfer tends to be crude 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.

METHODS & CODES

A model that treats broadband integration and 3D radiative transfer in a highly accurate and unbiased way is needed to serve as a standard of comparison for 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 sample at high resolution the broad range of the electromagnetic spectrum. It is called MCBRaT-3D and is available for public use and development. Unlike the direct approach to solving the radiative transfer equation, the Monte Carlo approach has the potential to be perfectly parallel, since the random samples are independent of one another.

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 improved to include integration over the electromagnetic spectrum to produce the broadband 3D model discussed above (MCBRaT-3D). 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, which are the largest and highest-resolution publicly available databases of their kind. The tools and workflow to create and subset them will also be made available.

As a usage example, these databases can be mined to update the decades-old broadband parameterizations of cloud radiative properties that are still in wide use today. Each of these products 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 in the literature with long heritage have been conducted to provide the first set of benchmark simulation results that can be used to evaluate other models. As an example, the figure shows cross-sectional snapshots of the heating rate through a cloud field due to the sun (top), due to emission of radiation by the system (bottom), and the net heating from the two (middle). Results such as these can be used as accuracy benchmarks for simpler, less computationally expensive models to characterize their bias.

WHY BLUE WATERS

Access to debugging and profiling tools such as DDT and CrayPat allowed me to streamline the development process. Having access to a point of contact on the Science and Engineering

Application Support staff helped me think through issues and find tailored solutions for my problems that would have otherwise had me stuck for weeks. The quick responsiveness of the Blue Waters' staff 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.

PUBLICATIONS & DATA SETS

Jones, A.L., and L. Di Girolamo, Design and Verification of a New Monochromatic Thermal Emission Component for the I3RC Community Monte Carlo Model. *Journal of the Atmospheric Sciences*, 75:3 (2018), pp.885–906.

Jones, A.L., Alexandraljones/imc-emission: Code base plus select benchmark results. *Zenodo*, (2017), DOI:10.5281/zenodo.574872.

Jones, A.L., AlexandraLJones/MCBRaT3D: Initial public release (Version V1.0.0-alpha). *Zenodo*, (2018), DOI:10.5281/zenodo.1313839.

Alexandra L. Jones received her PhD in atmospheric science in May 2016 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.

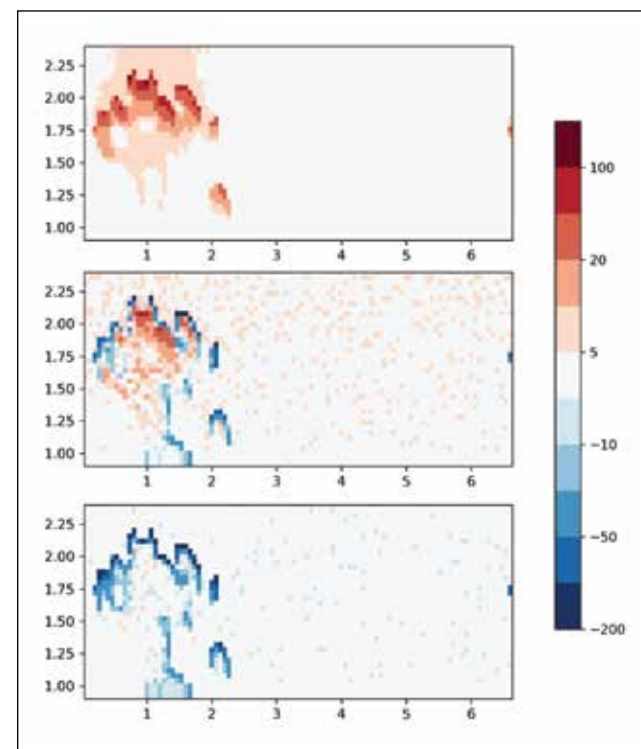


Figure 1: Vertical cross section of the radiative heating rate through a field of shallow cumulus clouds as calculated by MCBRaT-3D. The top image is the heating rate due to solar radiation; the bottom image is the heating rate due to radiation emitted by the clouds, atmosphere, and surface. The middle image is the net heating rate from the two sources.

DI

HIGH-FIDELITY BLADE-RESOLVED WIND FARM SIMULATIONS

Andrew Kirby, University of Wyoming
2016–2017 Graduate Fellow

TN

EXECUTIVE SUMMARY

Wind energy is an emergent renewable energy source throughout the world. Costs have dropped dramatically over the past two decades, making it a desirable alternative to fossil fuels. Improvement in wind energy application simulation technologies may have a profound economic impact through improved wind plant efficiency. Complete wind farm simulations will elucidate the flow physics that govern overall wind plant performance, including complex blade aerodynamics, turbine–turbine wake interference, and complex terrain effects. High-fidelity simulations of complete wind farm installations using blade-resolved models for wind turbines in complex terrain and atmospheric conditions will set unprecedented milestones in wind farm modeling capabilities. 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.

FS

RESEARCH CHALLENGE

Predicting wind farm performance represents a complex problem that spans spatial and temporal scales over 10 orders of magnitude from the continental scales that govern wind patterns to the thin boundary layers over the wind turbine blades. Improved prediction of wind farm productivity requires good resolution of flow structures and reliable modeling of turbulent eddies in this entire length-scale range. The temporal scales are also quite disparate: wind turbine rotation periods are of the order of seconds, while atmospheric inflow modes have time periods ranging from a few minutes to hours. Structural mechanics further compound this problem by introducing yet another set of temporal scales such as those due to significant elastic vibrations of the rotor blade and masts, which can be one to two orders of magnitude smaller than the wind turbine rotor period.

The aim of this work is to demonstrate the feasibility of simulating entire wind farm installations, including atmospheric

turbulence effects in complex terrain with high-fidelity methods with the goal of impacting wind farm siting decisions and wind turbine design practices. A major objective is to advance the state of the art of computational methods for achieving realistic and affordable petaflop simulations of time-dependent flows over complex configurations with moving bodies and overlapping mesh systems. Unsteady flows in nature are dictated by vortex structures and turbulent eddies that span relatively large scales of atmospheric dynamics to relatively small scales near moving wall boundaries (ocean surface, wind turbines, aircraft, etc.).

METHODS & CODES

At the University of Wyoming, we have developed a high-fidelity multiscale modeling methodology that can accurately predict the performance of wind turbines by reliably modeling the entire range of these spatial and temporal scales. The software framework uses the Large-Eddy Simulation (LES) approach for prediction of the turbulent flow fields in the off-body region. We used a multimesh framework to provide accurate and efficient prediction capabilities for vortex-dominated wind turbine flow fields.

We realized the multimesh paradigm by using multiple flow solvers, with each code optimized for the corresponding mesh type. In regions near wind turbine blades and towers, an unstructured mesh was used to handle the complex geometry and thin boundary layer regions. These mesh systems rotate with the blades and are overset on a background dynamically adaptive Cartesian mesh that is responsible for simulating the wakes propagating downstream of the turbines. A fully parallelized and load balanced adaptive mesh refinement capability was incorporated in the Cartesian mesh system to enable propagation of highly resolved wake features over long distances. Flow field values were interpolated in the regions of overlap between these two mesh systems at each timestep, and these overlapping patterns were computed dynamically in parallel using efficient search algorithms.

Nearly all software used on our framework was developed in-house at the University of Wyoming with the exception of the p4est




Figure 2: Instantaneous axial momentum at multiple downstream positions of the National Renewable Energy Laboratory WindPACT-1.5MW wind turbine.

adaptive mesh refinement framework. The two flow solvers we developed are NSU3D, an unstructured 3D finite-volume solver, and dg4est, a high-order discontinuous Galerkin finite-element solver. The overset solver we used was 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

Understanding the aerodynamics of the wind turbine is an essential aspect for energy production optimization, not only for the individual turbine but also for the complete wind farm. Exploration of wind turbine yawing for wind farm optimization introduces complex aerodynamics and possible structural effects. These complex aerodynamics, such as flow separation, cannot be captured accurately using lower-fidelity methods, such as actuator disk and actuator line methods. High-fidelity blade-resolved simulations are required for accurate prediction.

As a result of this research, state-of-the-art simulation analysis is now feasible through this computational framework. New pioneering analysis is now accessible, enabling the fundamental understanding of these complex wind turbine wake physics and their interactions. In particular, blade-resolved simulations allow for the study of the impact of complex aerodynamics on wake characteristics through a first-principles viewpoint.

At present, coupled experimental and computational studies are underway at the University of Wyoming (see Figs. 1 and 2)

with the overarching goal of providing the first validation wake data and first validated wind energy computational framework. By generating validated wake simulated data, lower-fidelity models can be more accurately derived with the possibility of improving full wind farm simulations at low cost thereby improving wind farm layout and wind farm efficiency. This can have a far-reaching impact on the renewable energy sector around the globe.

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 optimizing throughput.

PUBLICATIONS & DATA SETS:

Kirby, A., et al., Wind Turbine Wake Dynamics Analysis Using a High-Fidelity Simulation Framework with Blade-Resolved Turbine Models. *Proceedings of the 2018 Wind Energy Symposium* (AIAA SciTech Forum, Kissimmee, Florida), DOI:10.2514/6.2018-0256.

Kirby, A., et al., Visualization and Data Analytics Challenges of Large-Scale High-Fidelity Numerical Simulations of Wind Energy Applications. *Proceedings of the 2018 AIAA Aerospace Sciences Meeting* (AIAA SciTech Forum, Kissimmee, Florida), DOI:10.2514/6.2018-1171.



Figure 1: Wind turbine wake slice visualizing instantaneous normalized absolute tangential flow velocity demonstrating the wake propagation downstream of a small-scale rotor design with a two-meter diameter.

Andrew Kirby completed a PhD in mechanical engineering in May 2018. He worked under the supervision of Dimitri Mavriplis at the University of Wyoming.

DISCOVERING HUNDREDS OF NEW EXOPLANETS WITH K2

Ethan Kruse, University of Washington
2017–2018 Graduate Fellow

EXECUTIVE SUMMARY

The Kepler mission proved that space telescopes can discover thousands of exoplanets, including ones small and cool like the Earth. As the number and diversity of known exoplanets grow, we learn more about how our own planet formed and evolved, as well as how common other Earthlike planets may be in our universe. Both the K2 and upcoming TESS missions provide millions more star systems to search for planets, but there are not any mission-funded efforts to do so.

We developed a flexible, comprehensive planet search pipeline to discover exoplanets around the K2 and TESS targets. It has discovered hundreds of new planet candidates in K2 and TESS promises to deliver even more. These new planets orbit stars vastly different from the ones the Kepler mission studied. Combining the exoplanets from this research with the known ones from Kepler, we can learn how planets form and how common they are around stars of all types.

RESEARCH CHALLENGE

Kepler’s haul of over two thousand confirmed planets, most smaller than Neptune, changed the field of exoplanet research. We learned that small planets are ubiquitous and that planets similar to the Earth may be common in our galaxy [1]. Recently, evidence has emerged that reveals a gap in small-planet radii: Planets are either smaller than about 1.5 times the size of the Earth or more than twice as large, with relatively few planets of intermediate sizes [2]. This radius gap hints that planets with rocky cores may form up to about 1.5 times the size of the Earth, but some have an added puffy hydrogen and helium atmosphere inflating their radii to more than twice the Earth’s. However, all of this work has been done with planets around stars like our sun; we still know little about planets around other types of stars, including our galaxy’s smallest and most numerous red dwarfs.

The Kepler mission has been succeeded by K2, which studies more diverse stars. Yet because of the telescope’s reduced data quality in its extended mission, searching for planets is not as easy,

and there is no official effort to do so. The goal of this project is to develop a pipeline that accounts for the increased noise of K2 data and to search for planets. By finding planets around smaller stars, we will learn if the radius gap observed around larger stars holds true and get a better handle on what causes it in the first place. We can also constrain the occurrence rate of small planets around these smaller stars to see how common Earth-sized planets are in the much closer “habitable zones” surrounding the smaller stars.

METHODS & CODES

The raw K2 data is very noisy due to the telescope’s loss of fine pointing. First, we developed a technique that separates the instrumental noise from the astrophysical noise, in a process called pixel-level decorrelation. We run this processing pipeline (called EVEREST) on every star to create light curves with noise at about a factor of four better than the raw light curves: a precision that allows for planets to be found again [3].

These EVEREST light curves are then brought over to Blue Waters to search for planets. I have developed a general-purpose transit search pipeline based on an algorithm called QATS [4]. My transit pipeline can find not just the usual periodic planets but also planets that are overlooked in other planetary searches. Namely, those that transit only once or twice as well as those with transit timing variations due to perturbations by other planets in the system. Altogether, we have developed the most sensitive and comprehensive planet search pipeline for K2 data.

RESULTS & IMPACT

We searched the first two years of K2 data for new exoplanet candidates and discovered over 700. As with the original Kepler mission, most of these planets are smaller than Neptune, but unlike Kepler the majority orbit stars smaller than the sun and are closer, brighter, and easier to study in depth. Furthermore, because the K2 stars are scattered across the sky in different environments, we can learn about planet populations throughout the galaxy.

Identifying these candidates is just the first step. Adding a large, diverse pool of planet candidates to our sample will enable a suite of follow-up studies. Detailed observations of individual candidates will teach us about planet compositions and atmospheres. Comparing statistical samples of planets in different environments will help test planet formation and evolution theories.

WHY BLUE WATERS

Searching for planets is an enormously computationally intensive task. Going in, we don’t know a potential planet’s period, transit depth and duration, or ephemeris. Thus, searches are effectively over a four-dimensional grid of all possible parameters, and searching for a single star takes around an hour. With over 200,000 stars to search in the first two years of K2 data alone, large computing power is necessary. Access to Blue Waters speeds up development and processing, ensuring that our planet candidates get out in a timely manner for quick follow-up by the community.

As a sixth-year PhD student in astronomy at the University of Washington, Kruse works under the supervision of Eric Agol.

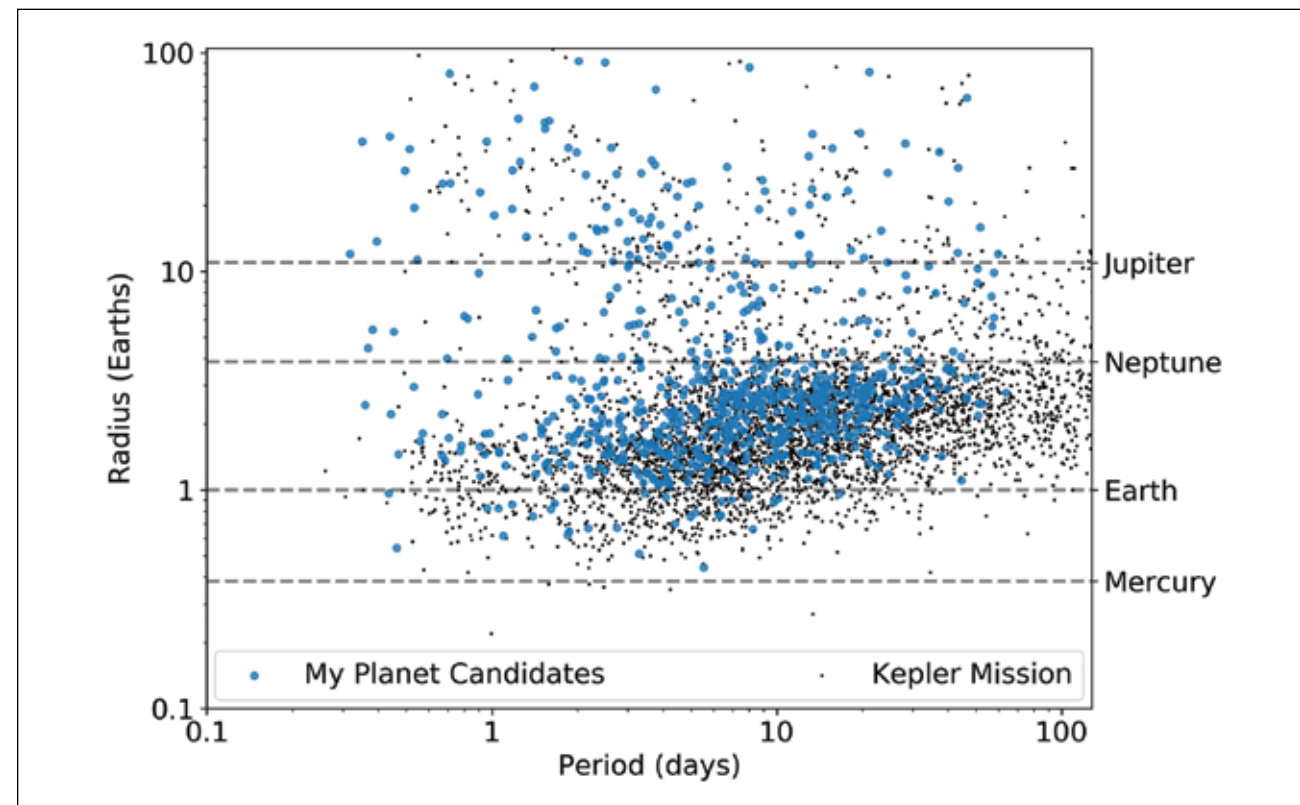


Figure 1: Our K2 exoplanet candidates with those from the original Kepler mission, showing we are finding a similar population of planets. A K2 campaign is limited to 80 days of observation compared to Kepler’s four years, preventing us from finding the longer periods and smaller planets.

TOWARD QUANTITATIVE PREDICTION OF DEFECT TOLERANCE IN SEMICONDUCTORS: LEVERAGING HPC AND PHYSICAL INSIGHTS TO DESIGN RADICALLY CHEAPER ENERGY MATERIALS

Rachel Kurchin, Massachusetts Institute of Technology
2017–2018 Graduate Fellow

EXECUTIVE SUMMARY

Climate change is a critical challenge confronting this generation, and if we are to address it, part of the solution must include a decarbonized electricity sector. Photovoltaic (PV) solar cells will undoubtedly play a vital role in this transition. My research aims to discover new materials that have a high potential to drastically lower the cost of PV solar cells. Performance of these new materials is less sensitive to imperfections in their atomic structure, permitting the use of simpler methods for their manufacturing, such as solution processing. We approach this problem through quantum mechanical calculations of the atomic-scale physics of the defects themselves to build a detailed understanding of what makes defects detrimental in materials such as silicon—the dominant PV material today—and relatively benign in others. This will eventually enable the design of defect-tolerant materials for the next generation of PV technologies.

RESEARCH CHALLENGE

My work aims to understand the Ångstrom-scale physical mechanisms underlying defect-tolerant behavior in some semiconductors. Materials such as silicon—the active material in approximately 90% of the PV market today—require an extremely high degree of purity and crystalline perfection in order to perform well, while some newly discovered materials can achieve high efficiencies despite the presence of many defects. However, all such materials prompt concerns surrounding elemental scarcity, toxicity, and stability. Understanding the physics behind the defect-tolerant behavior would enable us to engineer similar materials that do not suffer from these drawbacks and could potentially revolutionize the PV industry and more effectively combat climate change.




Figure 2: This figure, in turn, shows strategies for realizing shallow anion vacancies, which are related to crystal structure rather than chemistry. Part (a) shows a symmetry-enforced ionic structure, while (b) illustrates low anion coordination.

METHODS & CODES

We seek to understand the physics of crystalline defects through quantum mechanical calculations performed within the density functional theory formalism as implemented in the VASP code. We computed formation energies and charge transition levels of various defects of interest in order to understand trends in this behavior across a wide array of compounds.

RESULTS & IMPACT

The key impact of this work would be a drastic improvement in the ability of computer simulations to inform and guide materials design and discovery—the tasks that, at the moment, are mostly done by trial and error or intuition. Our work is part of a coming revolution in HPC-aided precision materials design that is capable of targeting specific properties and functionalities of interest to science as well as society.

WHY BLUE WATERS

The funding from the Blue Waters graduate fellowship, the community of fellows and of NCSA staff, more broadly, was key to my intellectual independence as well as my becoming a member of the HPC community.

PUBLICATIONS & DATA SETS

Kurchin, R.C., et al., Structural and Chemical Features Giving Rise to Defect Tolerance of Binary Semiconductors. *Chemistry of Materials*, 30:16 (2018), pp. 5583–5592.




Figure 1: This figure shows molecular-orbital-type band structures of ns^2 compounds featuring deep (a) and shallow (b) cation vacancies. A key finding of our work was the importance of aligning the orbital energies of the anion- p and cation- s states to enable larger band dispersion and hence shallow cation vacancy states.

Rachel Kurchin is a fourth-year PhD student in materials science and engineering at the Massachusetts Institute of Technology. She works with Tonio Buonassisi and expects to graduate in December 2019.

WIRES WITHIN WIRES: A MULTISCALE MODEL FOR COMPUTATIONAL INVESTIGATION OF BIOELECTRONIC PROTEIN DESIGN

Rachael Mansbach, University of Illinois at Urbana-Champaign
2017–2018 Graduate Fellow

EXECUTIVE SUMMARY

Certain proteins containing aromatic groups can assemble under appropriate conditions to form fluorescent mini-“wires” that can be used for electronic applications such as photovoltaic cells, light-emitting diodes, or pH sensors. Such organic electronics are desirable due to their ease of manufacture and nontoxicity. In order to produce rational principles for the design of such proteins, it is of great importance to understand their assembly on a molecular scale. Using Blue Waters, we employed coarse-grained molecular dynamics to study how changing the chemical properties of a series of proteins containing aromatic centers changes the properties of their amorphous aggregation. We illuminated generic properties of aggregation and identified five potential chemistries for further study. In the future, this work could help lead to the production of new biocompatible electronic devices.

RESEARCH CHALLENGE

The DXXX series is a group of small rod-like proteins with central aromatic cores possessing the ability to aggregate under acidic conditions into fluorescing semiconductive nanostructures [1]. However, the fluorescent properties of the nanostructures are limited by the extent to which the aromatic cores overlap. This overlap, in turn, is controlled by multiple factors, including the chemistry of the amino acid side chains and the kinetics under which aggregation occurs. The study of such protein aggregation at a molecular level is hindered by the comparatively large length scales and long time scales on which such aggregation occurs. Understanding the key determinants governing assembly is crucial in providing rational precepts for molecular design and engineering.

METHODS & CODES

In order to understand the effects of side chain chemistry on the aggregation of the DXXX series and to identify specific candidates with desirable optical properties, we created a simple model of the DXXX series, in which a single monomer was represented as a set of rigidly constrained beads. By changing the interactions of the different beads, we modeled changing the chemistries of the aromatic cores and the side chains. Using this inexpensive model to reach previously inaccessible length and time scales, we conducted Langevin dynamics in the HOOMD 2.1.7 simulation suite [2,3]. We performed five independent simulations of systems of 10,000 peptide monomers for 660 microseconds each using sixty different sets of model interaction parameters and analyzed the resulting properties of aggregation.

RESULTS & IMPACT

From our analysis of the simulations of large-scale aggregation, we identified the most salient interaction characteristic controlling the formation of aggregates likely to possess desirable optical properties. This characteristic is the overall “stickiness” of the amino acid side chains, represented in the model by the interactivity of the beads that represent the side chains. When the magnitude of the interactivity of the side chain beads becomes smaller than that of the beads representing the aromatic cores, desirable aggregation strongly increases because core—core interactions become more favorable than side chain—side chain interactions.

Our work also showed that the size of the side chains controls the small-scale morphology of the aggregates. At small scales (~10 nm), increasing the radius changes the aggregates from flat ribbons to twisted fibers, but at large scales (>30 nm), due to the

Figure 1: Visualization of a rigid body, patchy model of a DXXX peptide. The small, green *A* beads represent cofacial aromatic interactions; the large, red *SC* beads represent side chain interactions; and the large, blue *BB* beads represent noncofacial aromatic interactions.





Figure 2: Snapshots of the assembled morphologies for selected parameter settings (a–h). Side chain beads (*SC*) are transparent and blue, noncofacial aromatic core beads (*BB*) are red, and cofacial core beads (*A*) are green. The small-scale morphology depends on the specific parameters, but all molecules form porous branched networks at large scales.

intrinsic peptide geometry, all systems generically form a porous, branched network.

We defined two metrics by which to measure the quality of the assembled aggregates: (1) the rate at which “optical clusters”—aggregates expected to possess fluorescent properties—grow, and (2) the degree of one-dimensional order of the resulting aggregates that form. Using these measures, we performed a multiparameter optimization over the space of the sixty parameters to identify six model molecules expected to display the best aggregate properties. From these six parameter sets, we employed a mapping to identify the five peptide sequences corresponding to these optimal parameters. These five chemistries are expected to display rapid aggregation into thin wires with desirable optical properties.

Overall, this work characterizes the interactions and assembly of the DXXX series from the microscopic to the mesoscopic level, thus providing new fundamental understanding of the important molecular determinants of assembly behavior. This understanding enables rapid screening over molecular parameter space and the identification of chemistries predicted to favor assembly of large, linear aggregates with desirable optical properties. It also provides new rational design principles by which to engineer self-assembling peptides to fabricate large assemblies for bioelectronic applications. Further, it forms the coarsest level in a hierarchy of models of varying resolutions by which to perform high-throughput virtual screening of molecular space to efficiently discover and engineer these molecules and guide and accelerate experimental synthesis and characterization.

WHY BLUE WATERS

Performing molecular dynamics simulations, even at coarse-grained resolution, over sixty different parameter sets would have been prohibitively expensive without access to the computational resources provided by Blue Waters. Access to multiple XK GPU nodes enabled us to generate the necessary simulation data. In addition, the generous storage space on Blue Waters made it much easier to run in parallel by eliminating concerns tied to the fact that each of the $60 \times 5 = 300$ simulation runs generated between 10 and 20 gigabytes of data. Furthermore, the close support of the project staff, in particular our point of contact, was invaluable in enabling us to get up and running quickly. Finally, access to a larger big data community is particularly important for a highly interdisciplinary application such as ours, providing a pool of expertise we might otherwise have been unable to access.

PUBLICATIONS & DATA SETS

Mansbach, R., and A. Ferguson, A Minimal Patchy Particle Model for a Family of Self-assembling π -conjugated Optoelectronic Peptides. (To be submitted, 2018.)

Sixth-year PhD candidate Rachael Mansbach received her degree in physics in August 2018 from the University of Illinois at Urbana-Champaign, where she worked under the direction of Andrew L. Ferguson.

COMPUTER-AIDED FORMULATION DESIGN FOR IMPROVED NANOFORMULATION OF ANTICANCER THERAPEUTICS: USING SIMULATION TO IMPROVE THE EFFECTIVENESS OF CANCER DRUGS

William Payne, University of Nebraska Medical Center
2017–2018 Graduate Fellow

EXECUTIVE SUMMARY

Typical anticancer drugs produce harmful side effects and can be ineffective. Computational tools that elucidate the fundamental interactions in advanced drug formulations help to understand and predict the composition of next-generation therapeutics. Choosing an ideal drug carrier vehicle is difficult, and there are very few established design processes for the development of new nanoformulations. This has resulted in calls from government, industry, and academic institutions for the development of new design processes. To overcome the challenges preventing successful clinical adoption, and to further integrate nanoformulation into the industrial research and development process, quantitative, standardized methods of evaluating and designing nanomedicine must be developed. We used molecular simulation to show how nanoparticle formulations form on a molecular level in order to inform the design of new and better formulations. By observing the interactions among molecules, we can develop methods to formulate drugs that have better targeting, fewer side effects, and cheaper developmental costs.




Figure 1: Simulation of a polymeric nanoparticle to be used for drug delivery. These polymers assemble around hydrophobic pockets, which trap hydrophobic drug molecules.

RESEARCH CHALLENGE

Development of new nanoformulations usually requires extensive synthesis and experimental evaluation. Ways to guide experiments and to better understand what makes a “good” formulation could dramatically reduce development time and cost. This research is of importance to clinicians in cancer research as well as materials scientists and cancer researchers who are seeking to develop better cancer therapies.

METHODS & CODES

We used the molecular dynamics programs GROMACS and NAMD to model molecular systems across multiple size scales. We began by modeling single polymer strands and then graduated to multistrand systems and systems that included polymer molecules as well as organic dye molecules acting as surrogate drug molecules. By using dye molecules, we could compare our simulation results to experiments, further expanding the information available on the interactions in nanoformulations that are of importance. We used fluorescence spectroscopy as well as static and dynamic light scattering to observe the simulated systems in an experimental setting to confirm and explain the simulation results.




Figure 2: Simulation of different polymer strands interacting with dye molecules. This process demonstrates drug loading into a polymeric nanoformulation, but the use of dye molecules in place of drugs allows for better experimental evaluation.

RESULTS & IMPACT

This work was among the first to investigate self-assembled polymeric nanoformulations using simulations and correlating simulations to experimental observations. The simulations have enabled our lab to focus experimental investigation of new lead formulations as well as to begin working with new polymers in a more effective, methodical design process.

WHY BLUE WATERS

As an experimentalist seeking to develop computational chemistry skills, the support, tools, and knowledge provided by the Blue Waters staff dramatically reduced the learning curve. The Blue Waters supercomputer increased my ability to perform simulations, as queue times were shorter and Blue Waters performed simulations faster than other available resources.

William Payne is a third-year PhD student in pharmaceutical sciences at the University of Nebraska Medical Center. He is working under the supervision of Dr. Aron Mohs and plans to graduate in May 2019.

FACTOR SEPARATION ANALYSIS OF URBAN IMPACTS ON A SIMULATED SUPERCELL

Larissa Reames, University of Oklahoma
2015–2016 Graduate Fellow

EXECUTIVE SUMMARY

The effect of urban areas on weakly forced precipitation events is well documented and the mechanisms by which this happens are generally understood. However, the effects of urban areas on precipitation systems in synoptically active regimes, particularly severe convection, are relatively unstudied. This investigation used the ARW–WRF model to simulate an isolated supercell interacting with a large Great Plains urban area. We used a factor separation approach to determine the relative importance of roughness and thermal characteristics of urban areas on storm modification. Results generally suggest that surface roughness and its interactions with 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.

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 be able 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. In order to truly understand the nature of these interactions, and thus provide city planning recommendations for the future, it is important to ascertain whether the urban heat island or slower winds induced by increased urban surface roughness result in greater storm modifications.

METHODS & CODES

Using the Weather Research and Forecasting (WRF) [1] model, a community mesoscale numerical weather prediction model, we conducted a total of 334 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 well resolve the atmospheric boundary layer, we used 120 vertical grid points, with 20 of those points in the lowest 1.5 km above ground. In all, we integrated more than 29.7 million points over 75,600 timesteps 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 man-made surfaces (i.e., full physics; ORIG) was placed in 108 gridded locations throughout the domain to determine effects of the city-relative path of the storm. We performed two additional simulations for each urban location: one with only increased surface roughness over the city (ROUG), and one with only the different thermal properties of the urban area represented (THER).

RESULTS & IMPACT

Full- and single-physics urban simulations were compared 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 to establish how the urban effects on the prestorm environment are being represented, as well as changes in supercell structure, dynamics, and evolution. Analyses of the data are still underway, but early results suggest that groups of simulations (each with an urban area in a unique location) that are significantly similar have cities that are more geographically co-located (and more similar to the ORIG groupings) when the urban area is only represented as a roughness element than when it is only parameterized by its thermal properties. This result suggests that urban surface roughness may play a greater role in modifying supercell characteristics than the better-known urban heat island effect.

WHY BLUE WATERS

While HCA has been used previously for attribution of variations in synoptic and mesoscale fields to various factors, this is one of the first times it has been used to analyze storm-scale modifications. Given their large scale of motion, synoptic [$O(1,000\text{ km})$] and mesoscale [$O(100\text{ km})$] phenomena are generally quite predictable; thus, few simulations are required to attribute variations in fields of these scales to modifications in boundary conditions and parameterization options. However, due to their inherently unpredictable nature, to attribute deviations in storm-scale [$O(10\text{ km})$] phenomena 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 quantity of simulations needed to produce significant results required the large computational and data storage capacities of Blue Waters.

PUBLICATIONS & DATA SETS

Reames, L.J., and D.J. Stensrud, Influence of a Great Plains urban environment on a simulated supercell. *Monthly Weather Review*, 146 (2018), pp. 1437-1462.

Larissa Reames graduated from the University of Oklahoma in May 2017 with a PhD in meteorology. There, her work was directed by David Stensrud at Pennsylvania State University.

GENERAL RELATIVISTIC NEUTRINO RADIATION TRANSPORT: UNDERSTANDING THE DRIVING FORCE BEHIND COSMIC EXPLOSIONS

Sherwood Richers, California Institute of Technology
2016–2017 Graduate Fellow

EXECUTIVE SUMMARY

Calculating the behavior of neutrinos within core-collapse supernovae and neutron star mergers is key to understanding the mechanism behind the explosions and the formation of most elements in the universe. I apply a highly accurate Monte Carlo radiation transport method to simulate neutrinos in these environments to understand the detailed structure of the radiation field that current approximate methods are unable to probe. Using these results, I am then able to improve the current approximate methods.

RESEARCH CHALLENGE

The collapse of the cores of massive stars and collisions between neutron stars are known to produce some of the most powerful explosions in the universe, but it is unclear how these occur. The extreme density and temperature of the matter in these events also make them the likely sole source of all the elements in the universe other than hydrogen, helium, and lithium [1–3]. Lurking behind the bright flash and explosive ejection of every imaginable element in both types of events is the elusive neutrino [4,5]. Neutrinos, which ordinarily interact so weakly with matter that experimentalists have great difficulty even detecting them, are the dominant couriers of energy and determine which elements

form in the ejecta. Since we cannot see the interiors of these events where all the explosion-launching work is being done, we have to rely on numerical models to advance our understanding of the nonlinear dynamics. Simulating the behavior and impact of neutrinos remains particularly challenging, since it is a seven-dimensional problem (three spatial dimensions, two independent directions, neutrino energy, and time). I use an advanced Monte Carlo method to calculate the structure and impact of the six-dimensional (ignoring time-dependence) distribution of neutrinos in these systems. The results of this work help us to interpret models and observations and to pave the way for new computational methods that are both accurate and efficient.

METHODS & CODES

I calculated the steady-state radiation field in core-collapse supernovae and neutron star mergers using Sedonu [6], an open-source, general relativistic Monte Carlo neutrino radiation transport code. Sedonu samples the trajectories of a large number of individual neutrinos that are emitted from, absorb into, and scatter off a three-dimensional background matter profile. When a large number of these trajectories are calculated and averaged together, the outcome is a very detailed and highly accurate picture of what the neutrino radiation is doing at every point in space, at every energy, and in every direction. The physics that goes into calculating what each individual neutrino does is also fairly complete. General relativity is fully considered by evolving the particles along curved trajectories around the central neutron star or black hole. Detailed absorption, emission, and scattering rates are included via the NuLib library [7], and I include fully energy- and angle-dependent inelastic scattering processes. Once all the trajectories have been simulated, I also use the resulting neutrino distribution function to calculate the rate that neutrinos annihilate with anti-neutrinos.

RESULTS & IMPACT

Understanding the mechanisms behind the explosions from core-collapse supernovae and mergers is a longstanding problem that will likely not be solved without advances in the treatment of neutrino radiation transport. These calculations of Monte Carlo radiation transport are far more accurate than more efficient methods and elucidate where the latter need to be modified to enable high-accuracy simulations without the large expense of Monte Carlo radiation transport calculations. The detailed results pave the way for the next generation of models.

In addition, radiation transport is a problem in a wide variety of disciplines and has been a major component of high-performance computing since its inception. Each application of radiation transport has unique challenges, but algorithm development on state-of-the-art hardware in core-collapse supernovae is coupled to that in accretion disks, weather, and nuclear reactors, to name a few. The code Sedonu is open-source to facilitate rapid advances in the field at large.

WHY BLUE WATERS

Monte Carlo radiation transport requires a large number of test particles to be simulated to get smooth, converged results. In my calculations of the radiation field in three-dimensional neutron star mergers, I simulate around 1.3 trillion particles. Doing so requires: (1) a large amount of memory on each node so each node can fit

the entire problem, and (2) many nodes to repeat the problem enough times to get a solution with little noise. Blue Waters is the prime computer of its time for doing these calculations. The project staff have been very valuable in helping to optimize the code to make effective use of vector parallelism.

PUBLICATIONS & DATA SETS

Richers, S., et al., A Detailed Comparison of Multidimensional Boltzmann Neutrino Transport Methods in Core-collapse Supernovae. *The Astrophysical Journal*, 847:2 (2017), DOI:10.3847/1538-4357/aa8bb2.

Richers, S., et al., A Detailed Comparison of Multidimensional Boltzmann Neutrino Transport Methods in Core-collapse Supernovae (dataset). *Zenodo* (2017), DOI:10.5281/zenodo.807765.

Currently a postdoctoral student at North Carolina State University, Sherwood Richers did his Blue Waters Graduate Fellow work at the California Institute of Technology, where he received his PhD in June 2018.




Figure 1: The battle of neutrino transport methods. The left side shows the rate at which neutrinos are depositing (red) or removing (blue) energy from a core-collapse supernova using a discrete ordinates method (top) and my Monte Carlo method (bottom). The difference (right) is small, but significant.

BEYOND NAVIER-STOKES AND MOLECULAR DYNAMICS: UNDERSTANDING BIOMACROMOLECULAR DYNAMICS THROUGH THE DEVELOPMENT OF MULTISCALE HYBRID AND HYDRODYNAMIC SIMULATION

Sean Seyler, Arizona State University
2016–2017 Graduate Fellow

EXECUTIVE SUMMARY

Proteins are nanomachines that perform mechanical and/or chemical work dynamics spanning femtosecond timescales (i.e., covalent bond oscillations) to beyond the millisecond regime (e.g., glucose transport across a lipid membrane). All-atom molecular dynamics (MD) can fully capture solute–solvent interactions but is currently limited to microsecond timescales—orders of magnitude short of many biophysical timescales of interest. One viable means of overcoming this timescale problem is the hybrid atomistic–continuum (HAC) method where, for example, MD is used in a subdomain requiring atomistic detail while a hydrodynamic representation elsewhere captures solvent dynamics.

We developed a numerical fluctuating hydrodynamics (FHD) model to extend the regime of applicability of the Landau–Lifschitz Navier–Stokes equations—popular in contemporary HAC methods. Our model can capture nonlinear transport phenomena—e.g., viscoelasticity, thermoacoustic effects, and anomalous transport—that emerge at the nanoscale in liquid

water and may be necessary to explain the emergence of collective phenomena in so-called active matter systems.

RESEARCH CHALLENGE

At nanoscale, the behavior of soft-matter is dominated by nonequilibrium dynamics where thermal and sometimes active fluctuations inject energy at the smallest scales. So-called active matter systems, which include collectively swimming bacteria and suspensions of energized colloidal particles [1,2] often exhibit unintuitive behavior such as anomalous (non-Brownian) transport [3,4], flocking [1], or even active turbulence [5]. Active turbulence, for instance, arises in systems at near-zero Reynolds number and is characterized by an inverse cascade of energy upward toward larger spatiotemporal scales (where directed flows emerge spontaneously). Subcellular biology harbors many examples of active matter—e.g., zipping of DNA/RNA hairpins, intracellular transport via molecular motors, and enzyme catalysis—and

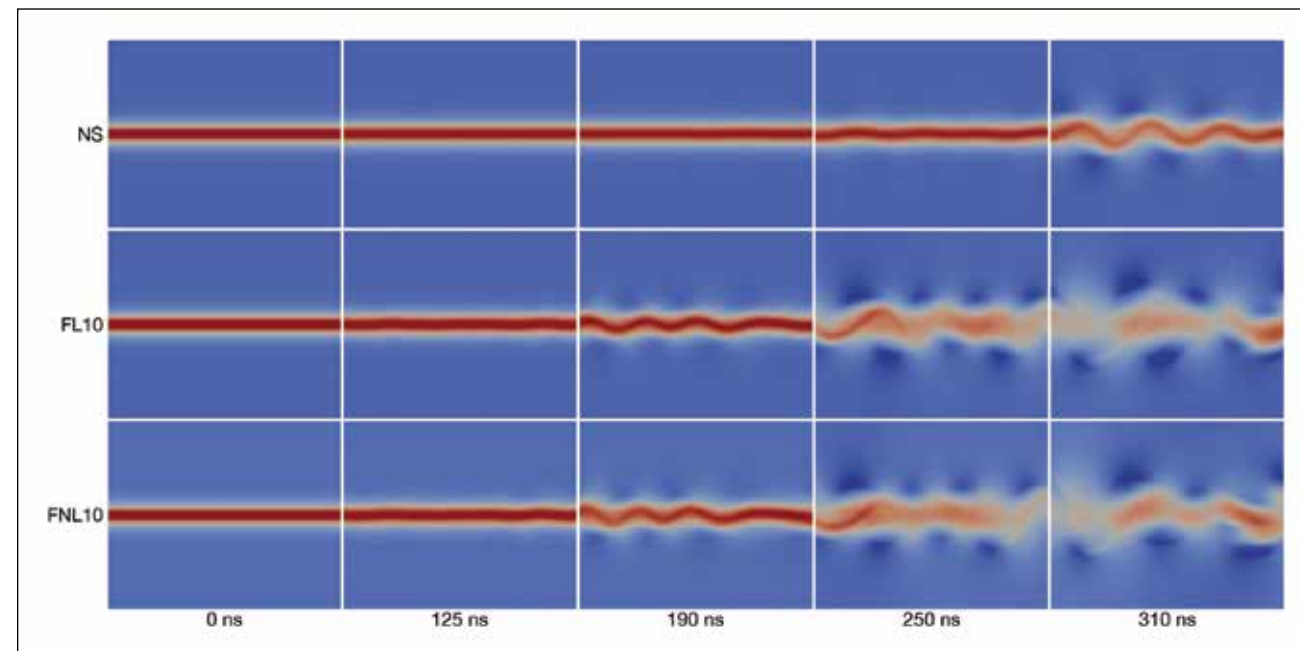


Figure 1: Nanojet comparison for liquid argon (1 bar, 80 K): Navier–Stokes (top), linearized 10-moment FHD (middle), full 10-moment FHD (bottom). Blue (red) corresponds to small (large) velocities (0–100 nm/ns); domain measures 500 $\mu\text{m}/\text{side}$ (80 cells). The 10-moment models differ only slightly, while Navier–Stokes takes longer to produce jet breakup.

exhibits collective organization that can span the entire spatial extent of a cell [4,6,7].

The role of fluctuations and transport is essential to the description of active matter, though surprisingly little is known about precisely how such systems produce emergent behavior [3,6,7]. Our long-term goal is to develop a multiscale hybrid simulation method that can bridge the atomistic–continuum gap by combining an MD and efficient FHD (solvent) model that extends the range of phenomena that may be feasibly studied [8,9]. However, our foremost efforts have been devoted to carefully extending hydrodynamic theory and accompanying numerical models to the nanoscale with the idea that our methods may eventually be incorporated in hybrid simulations.

METHODS & CODES

To construct a HAC method for nanoscale hydrodynamic simulation, it is necessary to augment the Navier–Stokes (NS) equations—a macroscopic hydrodynamic description—with thermal fluctuations [10]. This procedure was first explored by Landau and Lifschitz [11], where stochastic hydrodynamic equations describing mass, momentum, and energy transport were obtained by incorporating fluctuations in the viscous stress tensor and heat flux vector with correlations obeying linear fluctuation–dissipation [8,12,13]. Commonly known as the Landau–Lifschitz Navier–Stokes (LLNS) equations, this model still assumes that hydrodynamic timescales are “long” compared to microscopic collision times, leading to Gaussian white noise fluctuating terms that are spatiotemporally uncorrelated. Though LLNS has proven a powerful approach to modeling many nanoflows, we expect the aforementioned assumptions to break down for dense, inhomogeneous fluids when simulation grid cells approach nanometer dimensions (i.e., several water molecules) and hydrodynamic timescales of interest become comparable to molecular collision times.

We have developed an FHD model that extends the domain of applicability of LLNS. Our implementation is called HERMESHD (Hyperbolic Equations and Relaxation Model for Extended Systems of HydroDynamics) and is inspired by Grad’s 13-moment approximation [14]. The principal difference between LLNS and the fluctuating 13-moment (F13) equations is as follows: Whereas LLNS contains the familiar linear constitutive relations of Navier–Stokes—Newton’s law of viscosity and Fourier’s law (where stresses and heat flux are proportional to, respectively, velocity and temperature gradients)—F13, in effect, promotes these constitutive laws to time-dependent equations describing stress (tensorial, symmetric/traceless—five equations) and heat (vectorial—three equations) transport.

RESULTS & IMPACT

HERMESHD has a Python-wrapped library interface to facilitate rapid prototyping, Pythonic data manipulation, and interfacing with external codes; the code has been made available as open-source code under the MIT license on GitHub. HERMESHD is based on an efficient discontinuous Galerkin spatial discretization and leverages the hyperbolic structure of the F13 equations in a split-level, implicit–explicit scheme: Explicit time advance is carried out with second- or third-order strong-stability preserving (SSP) Runge–Kutta methods, while a locally implicit relaxation method steps over timestep constraints imposed by stiff source terms [15].

Unlike LLNS, the coupling of stress and heat flux equations to the momentum and energy equations in F13 generates memory effects, while thermal fluctuations, modeled as white Gaussian noise, enter through the extended (rather than momentum and energy) equations. As a result of this coupling, F13 naturally gives rise to viscoelasticity, finite-speed thermoacoustic waves, and colored Gaussian noise on small spatiotemporal scales, which may have important consequences for molecular motors and enzyme dynamics; indeed, recent experiments have revealed a connection between heat released during enzyme catalysis and enhancement of diffusion [16].

Previously, HERMESHD had been limited to simulating the 10- and 13-moment linearized equations (with fluctuations); we have since successfully implemented the nonlinear 10-moment equations and are nearing completion of the full (nonlinear) 13-moment model originally proposed by Grad [14] and being developed by others [17,18]. We are currently exploring applications of our numerical models to describing molecular motor efficiency as well as collective protein dynamics.

WHY BLUE WATERS

Proper validation of HERMESHD requires extensive testing against both hydrodynamic test problems and gold-standard atomistic MD simulation, which is particularly computationally expensive in three dimensions. Since HERMESHD uses MPI-based domain decomposition, Blue Waters’ nodes facilitate quick execution of such benchmarks. Blue Waters is also ideal for running benchmark MD simulations, which must be sufficiently large microcanonical systems to avoid spurious effects introduced by thermostats and barostats, to mitigate spurious correlations across periodic boundaries, and to obtain adequate statistical sampling. The quality of these data is essential, as empirical relations that enter into FHD simulation must be calculated numerically from MD simulations of bulk fluids [19].

Sean Seyler received a PhD in physics from Arizona State University in December 2017. He worked under the supervision of Oliver Beckstein.

THE IMPACTS OF HYDROMETEOR CENTRIFUGING ON TORNADO DYNAMICS: IMPROVING THE REALISM OF TORNADO SIMULATIONS

Ronald Stenz, University of North Dakota
2016–2017 Graduate Fellow

EXECUTIVE SUMMARY

Continued population growth in regions prone to tornadoes makes improving the understanding of these violent weather phenomena increasingly important. This research project attempts to improve our understanding of tornadoes by making the simulations used to study these destructive and dangerous weather events more physically realistic. For the first time, we have quantified the impacts that centrifuging of precipitation has on the vorticity budgets of these numerically simulated tornadoes. Preliminary findings have been consistent with radar observations of tornadoes, removing an unrealistic buildup of precipitation in the vortex center of simulated vortices and tornadoes, which has been widely seen in current tornado simulations. Ongoing work uses numerous tornado simulations to evaluate the significance of the inclusion of precipitation centrifuging in tornado dynamics, as well as more generally studying 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 airflow, creating an unrealistic buildup of precipitation in the vortex center, which

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 from each tornado, a better understanding of these destructive weather events is needed in order to improve forecasting, preparedness, and mitigation of their impacts. By including the centrifuging of precipitation in the model we use to learn about tornadoes, our simulations are more consistent with what is observed in nature, facilitating the improvement of our understanding of tornadoes. 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

We employed the widely used CM1 (Cloud Model 1) code for our simulations. CM1 was designed for studying small-scale atmospheric phenomena such as thunderstorms [1] and has also been designed to run efficiently on supercomputers such as Blue Waters. To quantify the impacts that the inclusion of centrifuging has on tornado dynamics, we first ran simulations without centrifuging. Just prior to the formation of a tornado, a checkpoint is saved, allowing the model to be run both with and without centrifuging from this point forward to determine how the centrifuging of precipitation influences the tornado dynamics. To define the magnitude of the centrifuging, an algorithm based on [2] uses trajectories released within the simulation to calculate the curvature of the flow and ultimately determine how quickly precipitation will be centrifuged, or moved outward, from the tornadic circulation. To quantify these impacts over a large sample size, we used atmospheric profiles of temperature, moisture, and wind from atmospheric soundings that were in close proximity to observed supercells [3] as the environmental conditions for our simulations of storms and their resulting tornadoes. A subset of these environments, which have been known to produce simulated tornadoes in previous research, was used for this study.

RESULTS & IMPACT

We completed idealized simulations and full-scale storm simulations (with a resulting tornado) with and without

centrifuging. In simulations without centrifuging, an unrealistic maximum of precipitation develops within 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. Similar to radar observations of tornadoes, the removal of precipitation from the vortex center is completed within several minutes in both types of simulations. Continued optimization and improvement to this centrifuging algorithm is in progress, with the goal of sharing these findings and eventually the centrifuging code to allow future research to benefit from the improved realism of the tornado simulations. Potential findings from this study on both the importance of centrifuging and also more general findings about how tornadoes work have the potential to improve future forecasting of tornadoes and also 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 provided by the Blue Waters team greatly facilitated accomplishing our research goals.

Ronald Stenz is a fifth-year PhD student in atmospheric sciences at the University of North Dakota. He is working under the supervision of Dr. Matthew Gilmore and hopes to graduate in 2019.

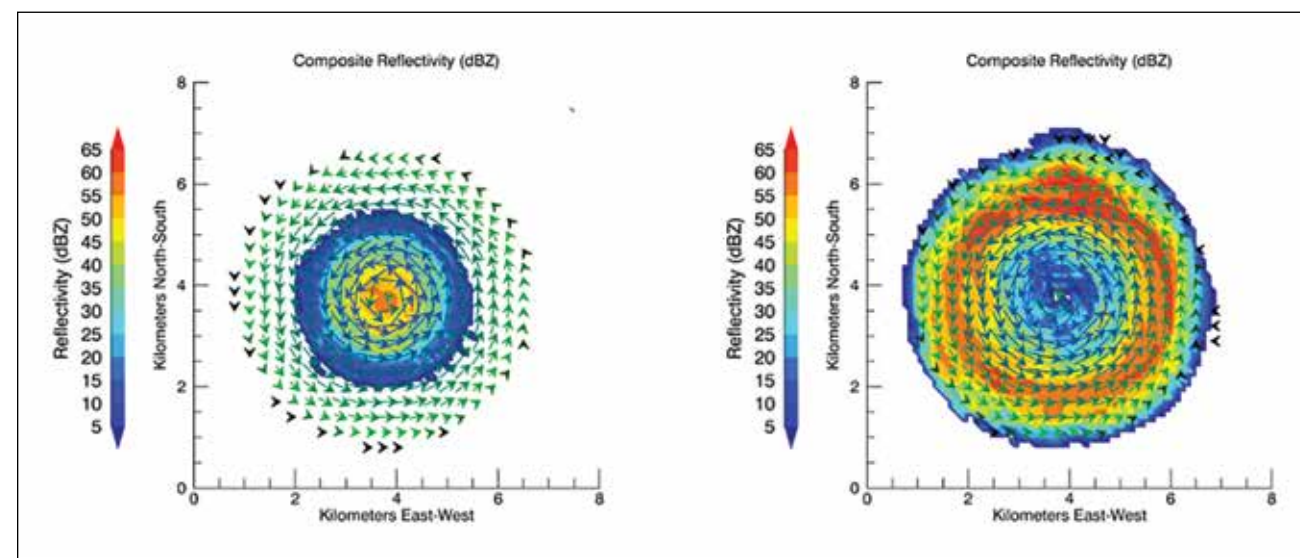


Figure 1: An idealized vortex simulation (left) where an unrealistic buildup of precipitation occurs in the vortex center. The same idealized vortex simulation (right) with centrifuging applied and a physically realistic minimum of precipitation in the vortex center.

DYNAMICS OF COLLOIDAL GLASS-FORMERS

Erin Teich, University of Michigan
2017–2018 Graduate Fellow

EXECUTIVE SUMMARY

Glass formation is a well-known outstanding mystery in the physical sciences. Although it has been an aspect of daily life for millennia, it still lacks a canonical thermodynamic explanation. We utilized our Blue Waters allocation to help shed light on this murky topic by investigating the structure and dynamics of assembly failure in a family of monodisperse colloidal systems composed of particles of related polyhedral shapes. We found that assembly failure arises from a competition between local structural motifs that are preferred in ordered structures composed of particles of similar, but not identical, shapes. Our work demonstrates the power of considering families of related systems when exploring phase behavior, and additionally probes the long-sought nature of the relationship between structure and dynamics in glass-forming systems.

RESEARCH CHALLENGE

Colloidal systems are capable of self-assembling into a wide variety of ordered structures ranging from the simple to the exceedingly complex [1]. Often, however, no such assembly occurs and the system instead remains disordered, displaying dynamical signatures characteristic of glass-formers. An understanding of why systems sometimes avoid crystallization is crucial, both for developing robust methods of colloidal self-assembly and for the development of glass-based technologies including rewritable data storage devices [2] and fiber optic networks. Despite the importance of this understanding, the underlying mechanism of the glass transition remains in contention. This is due in large part to the significant slowing down of any system as it approaches the glass transition [3], requiring investigations of glass formation to resolve system dynamics on time scales that vary by orders of magnitude. We were able to tackle this problem in studying glass-forming dynamics in a model colloidal system by using our

highly parallel simulation package, HOOMD-blue [4], and access to both the Cray XK7 nodes hosted on Blue Waters and its very generous storage capabilities.

METHODS & CODES

We performed hard particle Monte Carlo (HPMC) [5] simulations of model glass and crystal-formers comprised of hard polyhedra contained in the spheric triangle invariant 323 family [6]. These are a set of convex polyhedra formed by truncating the vertices and edges of a tetrahedron by sets of planes at varying radial distances from the polyhedron center. In HPMC simulations, particles have no interactions aside from those of excluded volume, and systems are therefore entirely entropically mediated. It was discovered previously [7] that systems of particles in certain regions of this shape space assemble into a rich variety of colloidal crystals, while systems in other regions of this shape space fail to assemble into any ordered structure at any density studied. We used Blue Waters to corroborate this previous discovery and to verify that these systems do in fact display canonical dynamical signatures that are characteristic of glass-formers.

RESULTS & IMPACT

To probe the structural and dynamical properties of suspected colloidal glass-formers, we simulated systems of 4,096 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 spanning six orders of magnitude, we wrote our system trajectories to disk 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 TB per simulation, each containing about 10 million simulation frames. We were able to store these extraordinarily large trajectories on Blue Waters' generous file system, using our 500 TB allocation, and we analyzed various dynamical signatures associated with canonical glass formation. We found that our colloidal glass-formers exhibited these signatures, including

plateaus in the mean-squared displacement and the real part of the self-intermediate scattering function that indicates caging, and peaks in the non-Gaussian parameter [8] as well as the self-part of the four-point susceptibility [9] that indicates dynamical heterogeneity associated with relaxation events. We were also able to extract relaxation times for our simulated systems and to plot those relaxation times as functions of system density in a variation of the canonical Angell plot [3]. Fig. 1 shows these relaxation times for two example systems and demonstrates that glass-former fragility is dependent on particle shape.

Additional work on other systems demonstrated a structural connection to the dynamical arrest we quantified on Blue Waters: We found that assembly failure in each system we studied arises from a competition among local structural motifs that are preferred in ordered structures composed of particles of similar (but not identical) shapes. Our work, facilitated by the resources of Blue Waters, helps to provide insight into the relationship between structure and dynamics in colloidal glass-formers.

WHY BLUE WATERS

As previously mentioned, to gather trajectory information at a wide range of time scales spanning six orders of magnitude we wrote our system trajectories to disk 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 TB per simulation, each containing about 10 million simulation frames. We were able to store these extraordinarily large trajectories within our 500 TB allocation on Blue Waters' generous file system; this was essential for our research and analysis.

PUBLICATIONS & DATA SETS:

Teich, E.G., G. van Anders, and S.C. Glotzer, Identity crisis in alchemical space drives the entropic colloidal glass transition. Under review (2018).

Erin Teich is a sixth-year PhD student in applied physics. She is working under the supervision of Sharon C. Glotzer at the University of Michigan and planned to graduate in August 2018.




Figure 1: Relaxation time as a function of system density for two model glass-formers comprised of particles of slightly different polyhedral shapes. Note that the fragilities of these systems, or the slopes of these curves, are dependent on particle shape.

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

Sam Totorica, Stanford University
2015–2016 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 intensely studied 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, and 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 clearly show 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, 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 currently helping to guide several experimental programs in the United States. We are now including Coulomb collisions in our simulations, which are showing the importance of collisionality on the structure of the reconnection layer (Fig. 1).

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 (Fig. 2) and have been able to show that in certain regimes SIC can reach a given noise level using one thousand times fewer




Figure 2: (a) Charge density for a simulation of the Weibel instability calculated using SIC and zoomed in on a filament. (b) Momentum distribution at the location indicated by the green circle in (a). (c) Stream number field over the same region and time shown in (a). (d) Transverse momentum along a lineout indicated by the red dotted line in (c).

simulation particles than standard methods. Future work will involve implementing SIC directly in the simulations to reduce noise and unphysical artifacts.

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 & 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), p. 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.

Totorica, S.R., F. Fiuza, and T. Abel, A new method for analyzing and visualizing plasma simulations using a phase-space tessellation. *Physics of Plasmas*, accepted (2018).

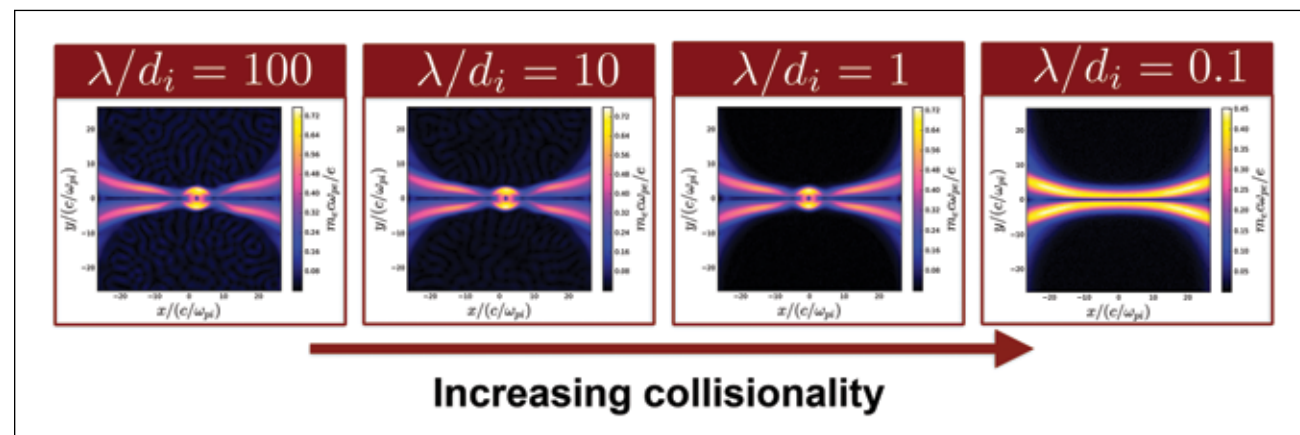


Figure 1: Magnetic field structure from four simulations with the same initial conditions but varied levels of collisionality. As the electron mean free path drops below the ion skin depth, plasmoid formation is suppressed.

Working under the direction of Tom Abel of Stanford University and Frederico Fiuza of the SLAC National Accelerator Laboratory, sixth-year PhD student Sam Totorica expects to receive his degree in 2018.

ENERGETIC DYNAMICS OF OCEAN BASIN MODEL WITH ROTATION, SURFACE WIND, AND BUOYANCY FORCING

Varvara Zemskova, University of North Carolina at Chapel Hill
2014–2015 Graduate Fellow

EXECUTIVE SUMMARY

The sensitivity of the circulation in the Southern Ocean to changes in surface wind stresses is one of the pivotal questions in oceanography. We employed an idealized rectangular rotating model of this ocean basin and forced it at the surface with a buoyancy flux that varied with latitude and with wind stress profiles that differ in shape and magnitude between simulation runs. The resulting flow is primarily wind-driven mean flow and is compensated by small-scale transient eddies. The circulation within the domain is affected by both the magnitude of the wind stress and symmetry/asymmetry between the Easterlies and the Westerlies. As the wind stress increases, the circulation progresses from a buoyancy-driven regime dominated by dense water formation in the south to a wind-driven regime. We found that when the surface winds are too strong or the Westerlies and the Easterlies are of the same magnitude (unlike the modern-day ocean), the circulation differs significantly from the one observed in the ocean under modern-day wind stress and surface buoyancy conditions.

RESEARCH CHALLENGE

The Southern Ocean, the region of the ocean between the Antarctic continent and the continents to the north, plays an important role in the global meridional transport of heat and tracers, oxygenation of the bottom waters via the formation of

dense Antarctic Bottom Water, and sequestration of atmospheric carbon. The energy in the Southern Ocean is drawn both from the strong and consistent zonal winds and differential surface buoyancy forcing. The circulation of the Southern Ocean can be divided into three distinct cells: the dense water formation cell originating near Antarctica (the lower cell), a cell of water originating in the North Atlantic that flows south at mid-depth and upwells in the Southern Ocean (the upper cell), and a surface water cell at mid-latitudes.

The response of the ocean circulation, in the Southern Ocean in particular, to changes in surface wind patterns due to changing climate has been one of the pivotal questions in oceanography. The changes in meridional, and particularly vertical transport, would affect the rates of carbon sequestration or its release back into the atmosphere and the supply of nutrients to support primary production in surface waters and the ocean food web.

METHODS & CODES

We ran five direct numerical simulations of an idealized rotating rectangular ocean basin, resolving the smallest dissipative scales, using the SOMAR code. All simulations were forced with a variable surface density field (dense near the southern end and lighter at the northern end of the domain). One simulation had only differential surface density profiles and no surface wind stress, and the remaining four simulations were forced with surface wind




Figure 1: a) Surface density distribution over the Southern Ocean (ECCO2) and from the simulation boundary conditions. (b) Surface wind stress profiles: WF1 and WF2 are sinusoidal with symmetric Easterlies and Westerlies; WF3 is a polynomial fit of the wind stress over the Southern Ocean; WF4 is twice the magnitude of WF3. The black axes represent the parameters using the dimensional ocean values. The blue axes represent the nondimensionalized values used in the DNS.

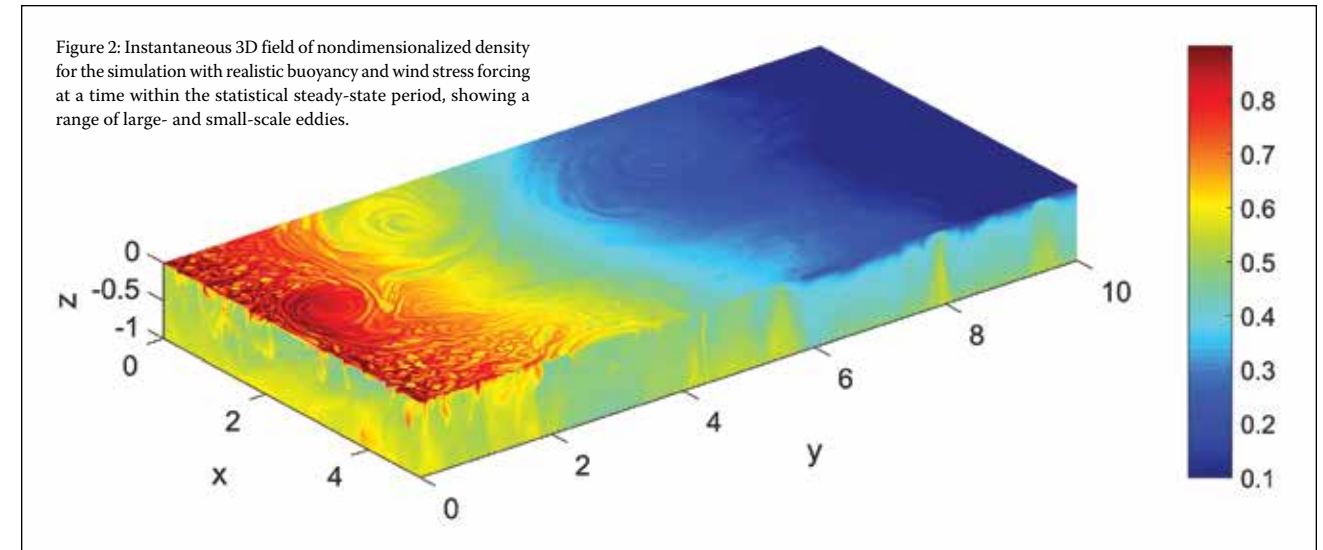


Figure 2: Instantaneous 3D field of nondimensionalized density for the simulation with realistic buoyancy and wind stress forcing at a time within the statistical steady-state period, showing a range of large- and small-scale eddies.

stress profiles of different magnitude and symmetry between the Easterlies and the Westerlies. These simulations were run from arbitrary initial conditions to a statistical steady state, such that both kinetic and available potential energies were not significantly variable in time. Each run was analyzed for the circulation stream functions and terms of the energy budget as well as the exchange rates between the kinetic and available potential energy reservoirs.

RESULTS & IMPACT

Many of the previous studies of the effects of increasing wind stress magnitudes with climate change have only focused on the changes in the Westerlies, disregarding how the interplay between the Easterlies and the Westerlies affects the circulation dynamics in the Southern Ocean. The Easterlies will become particularly important as ice melting progresses near Antarctica, exposing greater surface area affected by these winds.

This work is one of the first direct numerical simulation studies, resolving all energy scales of the system and dividing the flow into mean and turbulent components. Our study shows that as the wind stress, and subsequently the kinetic energy generation of the system, increases, the increase in the mean flow circulation that is wind-driven is compensated for by dissipation via smaller, transient turbulent eddies, and thus the dissipation of available potential energy (irreversible mixing) is unaffected by the wind stress magnitude.

However, the circulation in the basin is sensitive to both the profile shape and the magnitude of the surface wind stress. When there are no surface winds or the wind stress is small, the basin

is dominated by the lower dense water formation cell. When the wind stress is double in magnitude of the modern-day wind stress levels, the basin is primarily wind-driven, dominated by the upper cell of waters traveling from the northern hemisphere at mid-depth and upwelling in the Southern Ocean, while the dense water formation cell disappears.

This lower cell also is not present when the maximum magnitudes of the Easterlies and the Westerlies are equal. The balance between the deep dense water formation cell at the southern end and the wind-driven upwelling cell in the middle, as observed in the Southern Ocean today, is only achieved when the wind stresses are of modern-day magnitude, and the Westerlies are approximately twice as strong as the Easterlies. The disappearance of the dense water formation cell under two other scenarios is particularly important to note because this cell plays a vital role in supplying oxygen to the bottom ocean waters and sequestering atmospheric carbon below the mixed layer.

WHY BLUE WATERS

The direct numerical simulations are highly computationally intensive because of the great spatial resolution required for the small dissipative scales. In addition, these simulations need to be run for 70,000 to 100,000 timesteps, which takes several months, in order to reach a statistical steady state. We are interested in both the temporal-average and the deviations from the mean field, meaning that output files (approximately 2.8 GB per file, one per timestep) have to be stored for further analysis.

Varvara Zemskova is a sixth-year PhD student studying physical oceanography at the University of North Carolina at Chapel Hill. She is working under the supervision of Brian White and planned to graduate in December 2018.

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) have been a longstanding debate in phylogenetics, and 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 support for the three amphibian orders across the genome, we have discovered substantial discordance across genes in which these three possible models are supported. 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. Nonetheless, 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 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 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]. We used the Akaike information criterion [3] to quantify the direction and magnitude of support across genes. We then reconciled gene trees 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 there is substantial variation in the amphibian genome for which of the three possible topologies relating the three amphibian orders is supported across genes. Significant numbers of genes are found to support each of the three competing topologies, suggesting that either there has been an erosion of phylogenetic signal over deep time and/or that population-level processes (large effective population sizes and rapid divergences) have led to incomplete lineage sorting and gene tree-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

Access to the Blue Waters system provided opportunities to leverage large-scale GPU processors to accelerate the tens of thousands of gene tree estimation analyses needed to implement information-theoretic topology testing. Although many of the analyses I intended to execute on Blue Waters were delayed due to unforeseen issues with my data set, the experiences gained from interacting with NCSA staff and through the Blue Waters Symposia have helped to expand my HPC knowledge and will be invaluable assets in other computational aspects of my planned career.

Paul Hime received his PhD in evolutionary biology from the University of Kentucky in August 2017. He worked under the supervision of Dr. David Weisrock.

BLUE WATERS FULL SERVICE ASSESSMENT SUMMARY

A project as large as Blue Waters has many varied aspects to measure and assess: the operational statistics of the machine, the wide-ranging science results, the number of participants in various educational programs, and quality of service measures, to name a few. The pieces of the project may be tracked in different ways, but the goal is the same: how can we ensure that the project is providing the greatest service possible to the science communities.

Since going into full-service operations in 2013, Blue Waters has hosted nearly 4,000 researchers and students and a total of 785 projects from myriads of disciplines and areas, and these numbers continue to grow. The impact and benefits of the Blue Waters Project are truly widespread. Researchers from across the country utilized the system to advance scientific and computational research and data analysis using almost every type of method. Blue Waters serves as a proof of concept that the “convergence”

of high performance computing, big data analysis, and artificial intelligence can efficiently co-exist on the same platform, as targeted with the 2015 National Strategic Computing Initiative goals. Because of its flexibility and system architecture balance, Blue Waters serves a much more diverse research community, using more methods and algorithms and applications, than was envisioned at the start of the project. To date, researchers have produced over 1,000 high impact papers and articles as a result of their time on Blue Waters and many of these highlight results which were only possible on Blue Waters.

Fig. 1 shows the diverse disciplines using the Blue Waters system throughout Project Year Five (PY5), from June 1, 2017–May 31, 2018. During PY5, 1,160 researchers on 387 projects ran over 1.5 million jobs on Blue Waters, consuming over five billion core-hour equivalents. Indeed, science teams in PY5 utilized three percent




Figure 1: Distribution of annual actual usage by discipline June 2017–May 2018.

System Wide Interrupts by Month




Figure 2: The number and categories of System Wide Interrupts per month.

more time than in PY4, 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, perhaps a more familiar term. There are two ways to use the AMD Interlagos processors; 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 in past years primarily because the workload has evolved with more data analytics and our improved ability to increase node utilization by backfilling smaller/shorter jobs (while not delaying any large or very large jobs). This causes the median job size to decrease 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 five petabytes of data in 24-hour periods multiple times during the year. Despite the lower job size median

this year, Blue Waters still exceeds its control metric by having 44.3 percent of all the computational time used by large and very large-scale capability jobs.

As a partial measure of quality of service, the leadership of the Blue Waters Project monitors and reports the time required to respond to and resolve service requests, which includes traditional “trouble tickets,” as well as requests for basic issues, expanded assistance, added services, and suggestions. Over the year, the Project met its two quality of service control metrics by having a human staff member respond within four hours to over 95.6 percent of all service requests and successfully resolving 83.5 percent of all service requests within three business days.

The reliability of Blue Waters during PY5 was outstanding, with a scheduled availability above 99.72 percent. In addition to one unscheduled issue due to a campus wide power event there were only four scheduled system-wide outages, as illustrated in Fig. 2. Three of the four scheduled outages were taken to install critical, time-sensitive security updates. This outstanding availability is due in part to the remarkably low failure rates observed across the system, with the individual node failure rate below 1.5 nodes/day (0.006 percent) for the year, as shown in Fig. 3, and a drive




Figure 3: Average node interrupts per day by node type.

failure rate continuing below 0.43 drives/day (0.0026 percent). The reliability of the Blue Waters system is remarkable given the scale and complexity of the system!

Blue Waters is one of the most heavily monitored systems in existence. Health, status, and performance metrics are collected from all system components and used to diagnose issues, as well as to identify system and application level improvements. At the end of this project year we have collected over 16 trillion datums and are adding over 20 billion datums per day!

Administering Blue Waters effectively involves more than running the physical machine. The Blue Waters Project has successfully managed two significant awards from the NSF: \$226.6M was awarded in 2007 (OCI-0725070) for the development, deployment, and early science periods of the system and \$133.8M was awarded in 2013 (ACI-1238993) for the operations of Blue Waters. Through careful financial management, the deployment award was able to be extended to 2018. This was utilized in full to support the educational and outreach efforts including our internship and fellowship programs and the successful PAID program reported on last year. These efforts are essential to the comprehensive reach and impact of the Project. The full-service operations funding, initially intended for a five-year period, ending in March 2018, is now serving the community for seven years! Blue Waters will run at least into December 2019. The diligent planning and oversight over the awarded funds have demonstrated unquestionable stewardship to the mission of the Project and is evidence that Blue Waters is fulfilling—and exceeding—its mission

of providing a productive and innovative resources for leading-edge research in computational science and engineering.

The Blue Waters Project remains dedicated to continuous improvement, and to serve this goal, external evaluators conduct assessments to validate the effectiveness of the Project’s community and user engagement programs and to disseminate findings and best practices through publication and presentation. This effort is led by Professor Lizanne DeStefano, executive director of the Center for Education Integrating Science Mathematics, and Computing (SEISMC) as well as an associate dean and professor of psychology at Georgia Tech. The formative and summative evaluations address implementation, effectiveness, impact, sustainability, and diversity/inclusion. Assessments are conducted on the Blue Waters Graduate Fellows Program, the Undergraduate Internship Program, the Petascale Institute and workshops, and the annual Blue Waters Symposium. In addition, a survey of all users is utilized for broad input and one-on-one interviews with PRAC Principal Investigators are also conducted to assure the most impact and best value.

The fellowship program enhances research and enables new possibilities by utilizing the unique power of Blue Waters. As noted by one fellow “...the project that I tackled for the Blue Waters Fellowship was really something that I hadn’t proposed in my original PhD program, because it was something I couldn’t accomplish with the scope of resources that I had. So that was absolutely a positive, I just would not have been able to do this work without the Blue Waters’ Fellowship.”

Following former fellows has also demonstrated the impressive impact of the fellowship program. According to one former fellow, “The BW Fellowship was very important to my current and future professional endeavors. The fellowship allowed me the freedom and opportunity to propose and conduct my own research projects. Establishing this confidence and experience helped me obtain my faculty position without a postdoc. Moreover, the connections with NCSA staff and other BW fellows have been useful and will continue to be useful going forward. In particular, the opportunity to collaborate with NCSA and other fellows help our careers advance. Currently, another BW fellow and myself are brainstorming a joint cross discipline NSF proposal coupling our work. We plan to write and submit once they complete their PhD and are either a postdoc or a junior faculty.”

Another fellow summed it up simply by saying “The Blue Waters Fellowship changed my life.”

The Internship Program evaluation also included follow up assessments with past interns. A survey with previous cohorts (2014-2016) of interns shows that the program encourages students to pursue CS major/minors, increases their interest in computational research, and expands their career choices. According to one individual, “This is the most meaningful project I had the opportunity to carry out in my undergraduate program, considering my background from a small liberal arts college. I gained a lot of experiences working with large-scale distributed systems.”

Events from this past year also produced noteworthy findings. The Petascale Institute, which is used to kick off the annual internships, received impressive scores. On a five-point scale (1=Strongly Disagree, 5 = Strongly Agree) ratings included: 4.59 for “I had positive interactions and communication with other participants during the institute”; 4.71 for “I had positive interactions with instructors during this institute”; and 4.59 for “I have a better understanding of high performance computing in general as a result of this experience.”

Participants in the annual Blue Waters Symposium also expressed high satisfaction on the overall experience with a 4.51 rating for achieving the goals for attending the symposium. The attendance at the symposium was at an all-time high this year with nearly 200 attendees.

A user survey was conducted in May 2018. Active users in the previous six months as well as PIs were asked to participate. The survey response rate was nearly 25 percent, a very good rate for email/web surveys. Overall user satisfaction on Blue Waters services rated very highly (4.5) a five-point scale (1=Strongly Disagree, 5 = Strongly Agree). Furthermore, all services surveyed were individually rated over 4. Notably, 95.73 percent of the respondents said that “Blue Waters is very important or essential in conducting their research” and 99.31 percent agreed that “Blue Waters helped them to succeed in their research.” The following response further emphasized the important role of Blue Waters.

“My project has been the seed for several completely novel research directions. What may be unconventional in the impact Blue Waters has had was providing direct contact with like-minded, computing-oriented experts in otherwise disparate fields that has spawned new ideas and research directions. This would not have been possible without the supportive, diverse community of researchers and staff that comprise the Blue Waters Project.”

In summary, Blue Waters has proven to be a valuable asset to the educational research community, which relies on its excellent processing, memory and I/O performance as well as the intellectual services that make this a comprehensive program. Blue Waters is a groundbreaking, multi-faceted project that enables new frontier discoveries in science and engineering, energizes academic research in computer science and engineering, inspires innovations in education, helps build the future high-performance workforce and enhances U.S. competitiveness. The key to its success has been a focus on the system, the services, and the people who make it a valuable resource for science, engineering, and society.

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 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 members are nominated and the committee 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.

SETAC Members

PETASCALE COMPUTING RESOURCE TEAMS

- **Paul Woodward**, University of Minnesota (Chair)
- **Tom Cheatham**, University of Utah
- **David Ceperley**, University of Illinois Urbana-Champaign
- **Tiziana Di Matteo**, Carnegie Mellon University
- **Paul Morin**, University of Minnesota
- **Susan Bates**, NCAR
- **Brian O'Shea**, Michigan State University
- **Manuela Campanelli**, Rochester Institute of Technology
- **Tom Jordan** and **Philip J. Maechling**, University of Southern California
- **Said Elghobashi**, University of CA Irvine
- **Nikolai Pogorelov**, University of Alabama, Huntsville

GREAT LAKES CONSORTIUM FOR PETASCALE COMPUTATION TEAMS

- **H. Birali Runesha**, University of Chicago

INDUSTRY TEAMS

- **Rick Arthur**, General Electric Global Research, Advanced Computing

UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN TEAMS

- **Athol Kembal**, 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.

Measuring Descriptivity of Text Content

J. Stephen Downie

Fluid, Particulate, and Hydraulic System

Said Elghobashi

Protein Sequence Similarity Networks

John Gerlt

Modeling Brain Activity

Michael Insana

Dynamic Brain Network Estimation

Sanmi Koyejo

Dark Energy Survey

Felipe Menanteau

Convection-Gravity Wave Interactions

Stephen Nesbitt

Communication Avoiding Algorithms

Edgar Solomonik

NCSA Industry Staff

Ahmed Taha

Massively Parallel Materials Modeling

Dallas Trinkle

Structural Fluctuations of Protein

Yang Zhang

REFERENCES

SPACE SCIENCE

Balsara, Dinshaw

- [1] Balsara, D.S., Wave Propagation in Molecular Clouds. *Astrophysical Journal*, 465 (1996), pp. 775–794.
- [2] Li, H.-B., and M. Houde, Probing the Turbulence Dissipation Range and Magnetic Field Strengths in Molecular Clouds. *Astrophysical Journal*, 677 (2008), p. 1151.
- [3] 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.
- [4] 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.
- [5] 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.
- [6] 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.
- [7] Xu, S., and A. Lazarian, Damping of Magnetohydrodynamic Turbulence in Partially Ionized Plasma: Implications for Cosmic Ray Propagation. *Astrophysical Journal*, 833 (2016), p. 215.
- [8] 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.
- [9] Burkhardt, B., et al., Alfvénic Turbulence Beyond the Ambipolar Diffusion Scale. *Astrophysical Journal*, 805 (2015), pp. 118–126.

Carrasco Kind, Matias

- [1] Springel, V., The cosmological simulation code GADGET-2. *MNRAS*, 364, p. 1105.

Di Matteo, Tiziana

- [1] Banados, E., et al., An 800 million solar mass black hole in a significantly neutral universe at redshift 7.5. *Nature*, (2017), DOI:10.1038/nature25180.

Haas, Roland

- [1] Ott, C.D., et al., The Progenitor Dependence of Core-collapse Supernovae from Three-dimensional Simulations with Progenitor Models of 12–40 M_{\odot} . *Astrophys.J.*, 855 (2018). DOI:10.3847/2041-8213/aaa967.

Hawley, John

- [1] Balbus, S.A., and J.F. Hawley, A Powerful Local Shear Instability in Weakly Magnetized Disks: I. Linear Analysis. *Astrophysical Journal*, 376 (1991), pp. 214–222.
- [2] Balbus, S.A., and J.F. Hawley, Instability, Turbulence, and Enhanced Transport in Accretion Disks. *Rev. Mod. Phys.*, 70 (1998), pp. 1–53.
- [3] Krolik, J.H., and J.F. Hawley, A Steady State Alignment Front in an Accretion Disk Subject to a Lense–Thirring Torque. *Astrophysical Journal*, 806 (2015), DOI:10.1088/0004-637X/806/1/141
- [4] Hawley, J.F., and J.M. Stone, MOCCT: A Numerical Technique for Astrophysical MHD. *Comp. Phys. Comm.*, 89 (1995), pp. 127–148.
- [5] Evans, C.R., and J.F. Hawley, Simulation of General Relativistic Magnetohydrodynamic Flows: A Constrained Transport Method. *Astrophysical Journal*, 332 (1998), pp. 659–677.
- [6] Sorathia, K., J.H. Krolik, and J.F. Hawley, MHD Simulation of a Disk Subjected to Lense–Thirring Precession. *Astrophysical Journal*, 777 (2013), p. 21.

Huerta Escudero, Eliu (1)

- [1] Usman, S.A., et al., The PyCBC search for gravitational waves from compact binary coalescence. *Classical and Quantum Gravity*, 33:21 (2016), p. 215004.
- [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] George, D., E.A. Huerta, Deep Neural Networks to Enable Real-time Multimessenger Astrophysics. *Physical Review D*, 97 (2018), DOI:10.1103/PhysRevD.97.044039.
- [4] Taracchini, A., et al., Effective-one-body model for black-hole binaries with generic mass ratios and spins. *Physical Review D*, 89 (2014), DOI:10.1103/PhysRevD.89.061502.
- [5] Abbott, B.P., et al., Observation of Gravitational Waves from a Binary Black Hole Merger. *Physics Letters B*, 116 (2016), DOI:10.1103/PhysRevLett.116.061102.

Huerta Escudero, Eliu (2)

- [1] Huerta, E.A., et al., BOSS-LDG: A Novel Computational Framework that Brings Together Blue Waters, Open Science Grid, Shifter and the LIGO Data Grid to Accelerate Gravitational Wave Discovery. *eScience*, (2017), DOI:10.1109/eScience.2017.47.
- [2] Pordes, R., et al., The Open Science Grid. *Journal of Physics: Conference Series*, 78:1 (2007), p. 012057.
- [3] Livny, M., et al., Mechanisms for high throughput computing. *SPEEDUP journal*, 11:1 (1997), pp. 36–40.
- [4] Deelman, E., et al., Pegasus: A framework for mapping complex scientific workflows onto distributed systems. *Scientific Programming*, 13:3 (2005), pp. 219–237.

- [5] Usman, S.A., et al., The PyCBC search for gravitational waves from compact binary coalescence. *Classical and Quantum Gravity*, 33:21 (2016), p. 215004.
- [6] Jacobsen, D.M., and R.S. Canon, Contain this, unleashing Docker for HPC. *Proceedings of the Cray User Group*, 2015.
- [7] Abbott, B. P., et al., GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral. *Physics Letters B*, 119 (2017), DOI:10.1103/PhysRevLett.119.161101.

Levin, Deborah

- [1] Jambunathan, R., and D. Levin, CHAOS: An octree-based PIC-DSMC code for modeling of electron kinetic properties in a plasma plume using MPI-CUDA parallelization. *Journal of Computational Physics*, 373 (2018), pp. 571–604.
- [2] Jambunathan, R., and D. Levin, Comparison of Plasma Plume Characteristics Obtained Using PIC–DSMC Approach With Boltzmann Approximations. *45th International Conference on Plasma Science* (Denver, Colo., June 24–28, 2018).

Liu, Yi-Hsin

- [1] Bowers, K.J., et al., 0.374 Pflop/s Trillion-Particle Kinetic Modeling of Laser Plasma Interaction on Road-runner. *Proc. 2008 ACM/IEEE Conf. Supercomputing* (Austin, Tex., November 15–21, 2008).

Noble, Scott

- [1] Abbott, B.P., et al., GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral. *Phys. Rev. Lett.*, 119 (2017), DOI:10.1103/PhysRevLett.119.161101.
- [2] Abbott, B.P., et al., Multi-messenger Observations of a Binary Neutron Star Merger. *Astrophysical Journal Letters*, 848:2 (2017), DOI: 10.3847/2041-8213/aa91c9.
- [3] Abbott, B.P., et al., Gravitational Waves and Gamma-Rays from a Binary Neutron Star Merger: GW170817 and GRB 170817A. *The Astrophysical Journal Letters*, 848:2 (2017), DOI:10.3847/2041-8213/aa920c.
- [4] Bowen, D., et al., Quasi-periodic Behavior of Mini-disks in Binary Black Holes Approaching Merger. *Astrophysical Journal*, 853 (2018), DOI:10.3847/2041-8213/aaa756.

Norman, Michael

- [1] Tytler, D., et al., The effect of large-scale power on simulated spectra of the Ly α forest. *MNRAS*, 393:723 (2009), DOI:10.1111/j.1365-2966.2008.14196.x.
- [2] Bryan, G., et al., ENZO: An Adaptive Mesh Refinement Code for Astrophysics. *Astrophysical Journal Supp.*, 211:19 (2014), DOI:10.1088/0067-0049/211/2/19.
- [3] Burstedde, C., L. 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.
- [4] Cello Project Homepage: <http://cello-project.org>.

O'Shea, Brian

- [1] Bryan, G., et al., Enzo: An Adaptive Mesh Refinement Code for Astrophysics. *The Astrophysical Journal*, 211 (2014), p. 19.
- [2] Turk, M., et al., yt: A Multi-code Analysis Toolkit for Astrophysical Simulation Data. *The Astrophysical Journal Supplements*, 192 (2011), p. 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).
- [2] Gurnett, D.A., et al., Precursors to Interstellar Shocks of Solar Origin. *Astrophys. J.*, 809:121 (2015).
- [3] McComas, D.J., et al., Seven Years of Imaging the Global Heliosphere with IBEX. *Astrophys. J. Suppl.*, 229:41 (2017).

Quinn, Thomas

- [1] Kale, L.V., and S. Krishnan, Charm++: Parallel Programming with Message-Driven Objects. in *Parallel Programming using C++*, (1996), pp. 175–213.
- [2] Menon, H., et al., Adaptive techniques for clustered N-body cosmological simulations. *Computational Astrophysics and Cosmology*, 2 (2015), p. 1.
- [3] Tremmel, M., F. 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] Chen, C.H.K., and S. Boldyrev, Nature of Kinetic Scale Turbulence in the Earth's Magnetosheath. *Astrophys J.*, 842 (2017), p.122.
- [2] Passot, T., P.L. Sulem, and E. Tassi, Gyrofluid modeling and phenomenology of low- β_e Alfvén wave turbulence. *Phys. Plasmas*, 25 (2018), p. 042107.
- [3] Roytershteyn, V., et al., Numerical Study of Inertial Kinetic–Alfvén Turbulence. *Astrophysical Journal*, submitted (2018).
- [4] Roytershteyn, V., and G.L. Delzanno, Spectral Approach to Simulations of Kinetic Phenomena in Plasmas Based on Hermite Decomposition in the Velocity Space. *Frontiers in Astronomy and Space Sciences*, submitted (2018).

Schive, Hsi-Yu

- [1] Smith, B., et al., GRACKLE: a chemistry and cooling library for astrophysics. *Monthly Notices of the Royal Astronomical Society*, 466:2 (2017), DOI:10.1093/mnras/stw3291.
- [2] Bryan, G., et al., ENZO: An Adaptive Mesh Refinement Code for Astrophysics. *The Astrophysical Journal Supplement Series*, 211:2 (2014), DOI:10.1088/0067-0049/211/2/19.
- [3] Fryxell, B., et al., FLASH: An Adaptive Mesh Hydrodynamics Code for Modeling Astrophysical Thermonuclear Flashes. *The Astrophysical Journal Supplement Series*, 131 (1999), DOI:10.1086/317361.

Shapiro, Stuart

- [1] Gold, R., et al., Accretion disks around binary black holes of unequal mass: General relativistic magnetohydrodynamic simulations near decoupling. *Phys. Rev. D*, 89 (2014), p. 064060.
- [2] Zilhao, M., et al., Resolving the relative influence of strong field spacetime dynamics and MHD on circumbinary disk physics. *Phys. Rev. D*, 91 (2015), p. 024034.
- [3] Graham, M., et al., A systematic search for close supermassive black hole binaries in the Catalina Real-time Transient Survey. *Mon. Not. R. Astron. Soc.*, 453 (2015), p. 1562.
- [4] Charisi, M., et al., A population of short-period variable quasars from PTF as supermassive black hole binary candidates. *Mon. Not. R. Astron. Soc.*, 463 (2016), p. 2145.
- [5] D'Orazio, D., et al., A reduced orbital period for the supermassive black hole binary candidate in the quasar PG 1302-102? *Mon. Not. Roy. Astron. Soc.*, 452 (2015), p. 2540.
- [6] Etienne, Z., Y. Liu, and S. Shapiro, Relativistic magnetohydrodynamics in dynamical spacetimes: A new adaptive mesh refinement implementation. *Phys. Rev. D*, 82 (2010), p. 084031.
- [7] Khan, A., et al., Disks Around Merging Binary Black Holes: From GW150914 to Supermassive Black Holes. *Phys. Rev. D*, 97 (2018), p. 044036.
- [8] Farris, B., et al., Binary Black-Hole Mergers in Magnetized Disks: Simulations in Full General Relativity. *Phys. Rev. Lett.*, 109 (2012), p. 221102.
- [9] Gold, R., et al., Accretion disks around binary black holes of unequal mass: General relativistic MHD simulations of postdecoupling and merger. *Phys. Rev. D*, 90 (2014), p. 104030.

Tchekhovskoy, Alexander

- [1] Bardeen, J.M., and J.A. Petterson, The Lense–Thirring Effect and Accretion Disks around Kerr Black Holes. *Astrophysical Journal Letters*, 195 (1975), p. L65.
- [2] Liska, M., et al., Formation of precessing jets by tilted black hole discs in 3D general relativistic MHD simulations. *Monthly Notices of the Royal Astronomical Society*, 474 (2018), pp. L81–L85.

Teukolsky, Saul A.

- [1] <http://www.black-holes.org/SpEC.html>
- [2] Hannam, M., et al., The Samurai Project: Verifying the consistency of black-hole-binary waveforms for gravitational-wave detection. *Phys. Rev. D*, 79 (2009), DOI:10.1103/PhysRevD.79.084025.

Woodward, Paul R.

- [1] Woodward, P.R., F. Herwig, and T. Wetherbee, Simulating Stellar Hydrodynamics at Extreme Scale. *Computing in Science & Engineering*. In press (2018).
- [2] 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.
- [3] 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.
- [4] Herwig, F., et al., Global Non-Spherical Oscillations in 3-D 4 π Simulations of the H-Ingestion Flash. *Astrophys. J. Letters*, 792 (2014), p. L3.
- [5] 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.
- [6] Ritter, C., et al., Convective-reactive nucleosynthesis of K, Sc, Cl, and p-process isotopes in O-C shell mergers. *Monthly Notices Royal Astron. Soc.*, 474 (2018), pp. L1–L6.
- [7] Andrassy, R., F. Herwig, P.R. Woodward, and C. Ritter, 3D Hydrodynamic Simulations of C Ingestion into a Convective O Shell. In preparation for submission to MNRAS (2018).

PHYSICS & ENGINEERING**Acharya, Sumanta**

- [1] Varty, J.W., and E.E. Ames, Experimental Heat Transfer Distributions Over an Aft Loaded Vane With a Large Leading Edge at Very High Turbulence Levels. *ASME 2016 International Mechanical Engineering Congress and Exposition* (ASME, Phoenix, Ariz.), p. V008T10A040.
- [2] Jarrin, N., S. Benhamadouche, D. Laurence, and R. Prosser, A Synthetic-Eddy-Method for Generating Inflow Conditions for Large-Eddy Simulations. *Int. J. Heat Fluid Flow*, 27:4 (2006), pp. 585–593.

Aluru, Narayana (1)

- [1] Bi, L., S. Boulfrad, and E. Traversa, Steam electrolysis by solid oxide electrolysis cells (SOECs) with proton-conducting oxides. *Chem. Soc. Rev.*, 43 (2014), pp. 8255–8270.
- [2] Laguna-Bercero, M. A., Recent advances in high temperature electrolysis using solid oxide fuel cells: A review. *J. Power Sources*, 203 (2012), pp. 4–16.
- [3] Badwal, S. P. S., S. Giddey, and C. Munnings, Hydrogen production via solid electrolytic routes. *WIREs Energy Environ.*, 2 (2013), pp. 473–487.

- [4] Kröger, F.A., and H.J.Vink, Relations between the concentrations of imperfections in crystalline solids. *Solid State Phys.*, 3 (1956), pp. 307–435.
- [5] Norby, T., et al., Hydrogen in oxides. *Dalton Trans.*, 19 (2004), pp. 3012–3018.
- [6] Kreuer, K.D., et al., Proton conducting alkaline earth zirconates and titanates for high drain electrochemical applications. *Solid State Ionics*, 145 (2001), pp. 295–306.
- [7] Løken, A.; T.S. Bjørheim, and R. Haugsrud, The pivotal role of the dopant choice on the thermodynamics of hydration and associations in proton conducting BaCe_{0.9}X_{0.1}O_{3- δ} (X = Sc, Ga, Y, In, Gd and Er). *J. Mater. Chem. A*, 3 (2015), pp. 23289–23298.
- [8] Bjørheim, T.S., A. Løken, and R. Haugsrud, On the relationship between chemical expansion and hydration thermodynamics of proton conducting perovskites. *J. Mater. Chem. A*, 4 (2016), pp. 5917–5924.
- [9] Kresse, G., and J. Furthmüller, Efficiency of *ab-initio* total energy calculations for metals and semi-conductors using a plane-wave basis set. *Comput. Mater. Sci.*, 6 (1996), pp. 15–50.
- [10] Kresse, G., and D. Joubert, From ultrasoft pseudopotentials to the projector augmented-wave method. *Phys. Rev. B: Condens. Matter Mater. Phys.*, 59 (1999), p. 1758.
- [11] Kresse, G., and J. Furthmüller, Efficient iterative schemes for *ab initio* total-energy calculations using a plane-wave basis set. *Phys. Rev. B: Condens. Matter Mater. Phys.*, 54 (1996), p. 11169.
- [12] Perdew, J.P., K. Burke, and M. Ernzerhof, Generalized gradient approximation made simple. *Phys. Rev. Lett.*, 77 (1996), p. 3865.
- [13] Tang, W., E. Sanville, and G. Henkelman, A grid-based Bader analysis algorithm without lattice bias. *J. Phys.: Condens. Matter*, 21 (2009), p. 084204.
- [14] Henkelman, G., B.P. Uberuaga, and H. Jónsson, A climbing image nudged elastic band method for finding saddle points and minimum energy paths. *J. Chem. Phys.*, 113 (2000), pp. 9901–9904.

Aluru, Narayana (2)

- [1] Kondrat, S., P. Wu, R. Qiao, and A. Kornyshev, Accelerating Charging Dynamics in Sub-Nanometer Pores. *Nature Materials*, 13 (2014), pp. 387–394.
- [2] Fajardo, O.Y., et al., Electrolyte Lubricity with Ionic Liquid Nanoscale Films. *Sci. Rep.*, 5 (2015), DOI:10.1038/srep07698.
- [3] Wang, Y.-L., et al., Electrostatic Interactions in Soft Particle Systems: Mesoscale Simulations of Ionic Liquids. *Soft Matter*, 2018, DOI:10.1039/C8SM00387D.
- [4] Dong, K.; et al., Multiscale Studies on Ionic Liquids. *Chem. Rev.*, 117:10 (2017), pp. 6636–6695.
- [5] Dong, K.; S. Zhang, and J. Wang, Understanding the Hydrogen Bonds in Ionic Liquids and Their Roles in Properties and Reactions. *Chem. Commun.*, 52:41 (2016), pp. 6744–6764.

Balajewicz, Maciej

- [1] Bright, L., G. Lin, and J.N. Kutz, Compressive sensing based machine learning strategy for characterizing the flow around a cylinder with limited pressure measurements. *Phys. Fluids*, 25 (2013), DOI:10.1063/1.4836815.
- [2] San, O., and R. Maulik, Neural network closures for nonlinear model order reduction. arXiv preprint (2017), arXiv: 1705.08532.
- [3] Wang, Z., et al., Model identification of reduced order fluid dynamics systems using deep learning. *Int. J. Numer. Methods Fluids*, (2017), DOI: 10.1002/fld.4416.
- [4] Otto, S., and C. Rowley, Linearly-Recurrent Auto-encoder Networks for Learning Dynamics. arXiv preprint (2017), arXiv: 1712.01378.
- [5] Abadi, M., et al., TensorFlow: large-scale machine learning on heterogeneous distributed systems. White paper (2015), arXiv: 1603.04467.

Ceperley, David

- [1] McMahon, J.M., M.A. Morales, C. Pierleoni, and D.M. Ceperley, The properties of hydrogen and helium under extreme conditions. *Rev. Mod. Phys.*, 84 (2012), pp. 1607–1653.
- [2] Morales, M.A., C. Pierleoni, E. Schwegler, and D.M. Ceperley, Evidence for a first-order liquid–liquid transition in high-pressure hydrogen from *ab initio* simulations. *PNAS*, 107 (2010), pp. 12799–12803.
- [3] Knudson, M.D., et al., Direct observation of an abrupt insulator-to-metal transition in dense liquid deuterium. *Science*, 348 (2015), pp. 1455–1460.
- [4] Zaghoo, M., A. Salamat, and I. Silvera, Evidence of a first-order phase transition to metallic hydrogen. *Phys. Rev. B*, 93 (2016), p. 155128.
- [5] Pierleoni, C., et al., Liquid–liquid phase transition in hydrogen by coupled electron–ion Monte Carlo simulations. *PNAS*, 113 (2016), pp. 4953–4957.
- [6] Dias, R.P., and I.F. Silvera, Observation of the Wigner–Huntington transition to metallic hydrogen. *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.*, 148:10 (2017), DOI:10.1063/1.5001387.

Clark, Bryan

- [1] Chertkov, E., and B. Clark, A computational inverse method for constructing spaces of quantum models from wave functions. *arXiv:1802.01590* (2018).
- [2] Changlani, H., et al., Macroscopically Degenerate Exactly Solvable Point in the Spin-1/2 Kagome Quantum Antiferromagnet. *Physical Review Letters*, 120:11 (2018), p. 117202.
- [3] Yan, S., D. Huse, and S. White, Spin-liquid ground state of the S = 1/2 Kagome Heisenberg antiferromagnet. *Science*, 332:6034 (2011), pp. 1173–1176.
- [4] Balents, L., Spin liquids in frustrated magnets. *Nature*, 464:7286 (2010), p. 199.
- [5] Savary, L., and L. Balents, Quantum spin liquids: a review. *Reports on Progress in Physics*, 80:1 (2016), p. 016502.

- [6] Pal, A., and D. Huse, Many-body localization phase transition. *Physical Review B*, 82:17 (2010), p. 174411.
- [7] Basko, D.M., I.L. Aleiner, and B.L. Altshuler, Metal-insulator transition in a weakly interacting many-electron system with localized single-particle states. *Annals of Physics*, 321:5 (2006), pp. 1126–1205.
- [8] Yu, X., D. Pekker, and B.K. Clark, Finding matrix product state representations of highly excited eigenstates of many-body localized Hamiltonians. *Physical Review Letters*, 118:1 (2017) p. 017201.
- [9] Läuchli, A.M., J. Sudan, and R. Moessner, The $S=1/2$ Kagome Heisenberg Antiferromagnet Revisited. *arXiv:1611.06990* (2016).
- [10] Essafi, K., O. Benton, and L. Jaubert, A kagome map of spin liquids from XXZ to Dzyaloshinskii–Moriya ferromagnet. *Nature Communications*, 7 (2016), p. 10297.

Draayer, Jerry

- [1] Dytrych, T., et al., Collective Modes in Light Nuclei from First Principles. *Phys. Rev. Lett.*, 111 (2013), p. 252501.
- [2] Launey, K.D., et al., Symmetry-guided large-scale shell-model theory. *Prog. Part. Nucl. Phys.*, 89 (2016), 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] Langr, D., T. Dytrych, K.D. Launey, and J.P. Draayer, Accelerating Many-Nucleon Basis Generation for High Performance Computing Enabled *Ab Initio* Nuclear Structure Studies. *Intl. J. High Performance Computing Applications*, submitted (2018).
- [5] 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), DOI:10.1016/j.cpc.2016.06.006.
- [6] Draayer, J.P., et al., Symmetry-adapted no-core shell model – the *ab initio* picture of nuclear collectivity. *Emergent Phenomena in Atomic Nuclei from Large-scale Modeling: A Symmetry-guided Perspective*, ed. K.D. Launey, World Scientific Publishing Co., in press (2017); ISBN: 978-981-3146-04-4.

Duan, Lian

- [1] Jiang, G.S., and C.W. Shu, Efficient Implementation of Weighted ENO Schemes. *Journal of Computational Physics*, 126:1 (1996), pp. 202–228.
- [2] Taylor E.M., M. Wu, and M.P. Martin, 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.
- [3] Williamson, J., Low-storage Runge-Kutta schemes. *Journal of Computational Physics*, 35:1 (1980), pp. 48–56.
- [4] Xu, S., and M.P. Martin, Assessment of inflow boundary conditions for compressible turbulent boundary layers. *Physics of Fluids*, 16:7 (2004), pp. 2623–2639.

- [5] Touber, E., and N.D. Sandham, Oblique shock impinging on a turbulent boundary layer: Low-frequency mechanisms. *38th Fluid Dynamics Conference and Exhibit*, AIAA Paper 2008–4170 (AIAA, Seattle, Wash., June 23–26, 2008).
- [6] Laufer, J., Some statistical properties of the pressure field radiated by a turbulent boundary layer. *Physics of Fluids*, 7:8 (1964), pp. 1191–1197.

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] Fischer, P.F., J.W. Lottes, and S.G. Kerkemeier, Nek5000 web page, <https://nek5000.mcs.anl.gov> (2008).
- [3] Deville, M.O., P.F. Fischer, and E.H. Mund, *High-order methods for incompressible fluid flow*. Cambridge University Press (Cambridge, U.K., 2002).
- [4] Mittal, K., and P.F. Fischer, Mesh smoothing for spectral element method. *J. Scientific Computing*, accepted for publication (2018).
- [5] Merrill, B., et al., A spectrally accurate method for overlapping grid solution of incompressible Navier–Stokes equations. *J. Comp. Physics*, 307 (2016), pp. 60–93.
- [6] Mittal, K., S. Dutta, and P.F. Fischer, Nonconforming Schwarz–Spectral element methods for incompressible flow. *Parallel CFD2018* (Indianapolis, Ind., May 14–17, 2018).

Garcia, Marcelo

- [1] Mier, J.M., *Experimental analysis of flow and turbulence characteristics in oscillatory boundary layers from LDV measurements*. Doctoral dissertation, University of Illinois at Urbana-Champaign (2014).
- [2] Fischer, P.F., J.W. Lottes, and S.G. Kerkemeier, nek5000 website, <http://nek5000.mcs.anl.gov> (2018).
- [3] Cantero, M.I., S. Balachandar, and M.H. García, An Eulerian–Eulerian model for gravity currents driven by inertial particles. *Intl J. Multiph. Flow*, 34 (2008), pp. 484–501.
- [4] García, M.H., ed., Chapter 2: Sediment Transport and Morphodynamics. *Sedimentation Engineering. ASCE Manuals and Reports on Engineering Practice No. 110* (2008), ASCE/EWRI, Reston, Va.
- [5] Pedocchi, F., M.I. Cantero, and M.H. García, 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*, 4 (2010), p. 115104.

- [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.S. 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.

Huff, Kathryn

- [1] Lindsay, A., and K. Huff, Moltres, Finite Element Based Simulation of Molten Salt Reactors. *The Journal of Open Source Software*, 3:21 (2018), DOI:10.21105/joss.00298.
- [2] Lindsay, A., G. Ridley, A. Rykhlevskii, and K. Huff, Introduction to Moltres: An Application for Simulation of Molten Salt Reactors. *Annals of Nuclear Energy*, 114 (2018), pp. 530–540.
- [3] Rykhlevskii, A., J.W. Bae, and K. Huff, Arfc/Saltproc: Code for Online Reprocessing Simulation of Molten Salt Reactor with External Depletion Solver SERPENT. *Zenodo*, (2018), DOI:10.5281/zenodo.1196455.
- [4] 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*, 239:10 (2009), pp. 1768–1778.
- [5] 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*, 22:3–4 (2006), pp. 237–254.
- [6] Balay, S., et al., *PETSc Users Manual* (2015), <http://www.mcs.anl.gov/petsc>.
- [7] Rykhlevskii, A., A. Lindsay, and K.D. Huff, Full-Core Analysis of Thorium-Fueled Molten Salt Breeder Reactor Using the SERPENT 2 Monte Carlo Code. *Transactions of the American Nuclear Society*, (2017).
- [8] Rykhlevskii, A., A. Lindsay, and K.D. Huff, Online Reprocessing Simulation for Thorium-Fueled Molten Salt Breeder Reactor. *Transactions of the American Nuclear Society*, (2017).

Ismail-Beigi, Sohrab

- [1] Jain, N., et al., OpenAtom: Scalable *Ab-Initio* Molecular Dynamics with Diverse Capabilities. *High Performance Computing*, eds. J.M. Kunkel, P. Balaji, and J. Dongarra, (Springer International Publishing, 2016), pp. 139–158).
- [2] Liu, J., et al., Adsorption and Diffusion of Hydrogen in a New Metal–Organic Framework Material: [Zn(bdc)(ted)_{0.5}]. *J. Phys. Chem. C*, 112 (2008), pp. 2911–2917.

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), p. 395502.

- [2] Hesjedal, T., and Y. Chen, Topological Insulators Engineered Heterostructures. *Nat. Mater.*, 16 (2017), pp. 3–4.
- [3] Dumett Torres, D., P. Banerjee, S. Pamidighantam, and P.K. Jain, A Non-Natural Wurtzite Polymorph of HgSe: A 3D Topological Insulator. *Chem. Mater.*, 29:15 (2017), pp. 6356–6366.

Jasiuk, Iwona

- [1] Bakir, M., and I. Jasiuk, Novel Metal–Carbon Nanomaterials: A Review on Covetics. *Advanced Materials Letters*, 8:9 (2017), pp. 884–890.
- [2] Plimpton, S., Fast Parallel Algorithms for Short-Range Molecular Dynamics. *J. Comp. Phys.*, 117 (1995), pp. 1–19.
- [3] 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).
- [4] ABAQUS Documentation. Dassault Systèmes, Providence, R.I., 2018.

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] 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.
- [4] Coleman, A.J., et al., Acoustic cavitation generated by an extracorporeal shockwave lithotripter. *Ultrasound Med. Bio.*, 13 (1987), pp. 69–76.
- [5] 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.
- [6] 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.
- [7] Philipp, A., and W. Lauterborn, Cavitation erosion by single laser-produced bubbles. *J. Fluid Mech.*, 361 (1998), pp. 75–116.
- [8] Johnsen, E., and T. Colonius, Numerical simulations of non-spherical bubble collapse. *J. Fluid Mech.*, 29 (2009), pp. 231–261.

Johnson, Harley

- [1] Ghale, P., and H.T. Johnson, A sparse matrix-vector multiplication based algorithm for accurate density matrix computations on systems of millions of atoms. *Computer Physics Communications*, 227 (2018), pp. 17–26.
- [2] Jay, L., H. Kim, Y. Saad, and J. Chelikowsky, Electronic structure calculations for plane-wave codes without diagonalization. *Computer Physics Communications*, 118:1 (1999), pp. 21–30.

- [3] Niklasson, A., C. Tymczak, and M. Challacombe, Trace resetting density matrix purification in O(N) self-consistent-field theory. *The Journal of Chemical Physics*, 118:19 (2003), pp. 8611–8620.

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>.

Leburton, Jean-Pierre

- [1] Kulis, M., and M. Esteller, DNA Methylation and Cancer. *Epigenetics and Cancer, Part A*, 70 (2010), pp. 27–56.
- [2] Laszlo, A.H., et al., Detection and Mapping of 5-Methylcytosine and 5-Hydroxymethylcytosine with Nanopore MspA. *Proc. Natl. Acad. Sci. U.S.A.*, 110 (2013), pp. 18904–18909.
- [3] Shim, J., et al., Nanopore-Based Assay for Detection of Methylation in Double-Stranded DNA Fragments. *ACS Nano*, 9 (2015), pp. 290–300.
- [4] Humphrey, W., A. Dalke, and K. Schulten, VMD—Visual Molecular Dynamics. *J. Mol. Graph.*, 14 (1996), pp. 33–38.
- [5] MacKerell, A. D., et al., All-Atom Empirical Potential for Molecular Modeling and Dynamics Studies of Proteins. *J. Phys. Chem. B*, 102 (1998), pp. 3586–3616.
- [6] Foloppe, N., and A.D. MacKerell, Jr., All-Atom Empirical Force Field for Nucleic Acids: I. Parameter Optimization Based on Small Molecule and Condensed Phase Macromolecular Target Data. *J. Comput. Chem.* 21 (2000), pp. 86–104.
- [7] Aksimentiev, A., J.B. Heng, G. Timp, and K. Schulten, Microscopic Kinetics of DNA Translocation through Synthetic Nanopores. *Biophys. J.*, 87 (2004), pp. 2086–2097.
- [8] Girdhar, A., C. Sathe, K. Schulten, and J.-P. Leburton, Graphene Quantum Point Contact Transistor for DNA Sensing. *Proc. Natl. Acad. Sci. U.S.A.* 110 (2013), pp. 16748–16753.
- [9] Sarathy, A., and J.-P. Leburton, Electronic Conductance Model in Constricted MoS₂ with Nanopores. *Appl. Phys. Lett.*, 108 (2016).
- [10] Qiu, H., A. Sarathy, K. Schulten, and J.-P. Leburton, Detection and Mapping of DNA Methylation with 2D Material Nanopores. *npj 2D Mater. Appl.*, 1 (2017), p. 3.
- [11] Sarathy, A., H. Qiu, and J.-P. Leburton, Graphene Nanopores for Electronic Recognition of DNA Methylation. *J. Phys. Chem. B*, 121 (2017), pp. 3757–3763.

- [12] Gilboa, T., et al., Single-Molecule DNA Methylation Quantification Using Electro-Optical Sensing in Solid-State Nanopores. *ACS Nano*, 10 (2016), pp. 8861–8870.
- [13] Du, Q., P.-L. Luu, C. Stirzaker, and S.J. Clark, Methyl-CpG-Binding Domain Proteins: Readers of the Epigenome. *Epigenomics*, 7 (2015), pp. 1051–1073.
- [14] Ratel, D., J.L. Ravanat, F. Berger, and D. Wion, N6-Methyladenine: The Other Methylated Base of DNA. *BioEssays*, 28 (2006), pp. 309–315.

Levin, Deborah

- [1] Sawant, S., O. Tumuklu, R. Jambunath, and D. Levin, Application of Adaptively Refined Unstructured Grids in DSMC to Shock Wave Simulations. *Computers and Fluids*, 170 (2018), pp. 197–212.

Mackenzie, Paul

- [1] Bai, Z., N.H. Christ, and C.T. Sachrajda, The K_L-K_S Mass Difference. *EPJ Web Conference*, 175 (2018), p. 13017.
- [2] Bazavov, A., et al., Up-, down-, strange-, charm-, and bottom-quark masses from four-flavor lattice QCD. Submitted to *Phys. Rev. D.*, (2018), arXiv:1802.04248.
- [3] Bazavov, A., et al., B- and D-meson leptonic decay constants from four-flavor lattice QCD. Submitted to *Phys. Rev. D.*, (2017), arXiv:1712.09262.

Mori, Warren

- [1] Xu, X.L., et al., High quality electron bunch generation using a longitudinal density-tailored plasma-based accelerator in the three-dimensional blowout regime. *Phys. Rev. Accel. Beams*, 20 (2017), p. 111303.
- [2] Wen, H., Petascale kinetic simulations of laser plasma interactions relevant to inertial fusion—controlling laser plasma interactions with laser bandwidth. *The 3rd International Conference on Matter and Radiation at Extremes*, Qingdao, China, May 2018.

Neubauer, Mark

- [1] ATLAS Collaboration, Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC. *Phys. Lett. B*, 716 (2012), pp. 1–29.
- [2] CMS Collaboration, Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC. *Phys. Lett. B*, 716 (2012), p. 30.

Raman, Venkat

- [1] Cocks, P.A., A.T. Holley, and B.A. Rankin, High fidelity simulations of a non-premixed rotating detonation engine. *54th AIAA Aerospace Sciences Meeting* (2016), p. 0125.
- [2] Rankin, B.A., et al., Imaging of OH* chemiluminescence in an optically accessible nonpremixed rotating detonation engine. *53rd AIAA Aerospace Sciences Meeting* (2015), p.1604.

- [3] Jasak, H., A. Jemcov, and Z. Tukovic, OpenFOAM: A C++ library for complex physics simulations. *International workshop on coupled methods in numerical dynamics*, 1000 (2007), pp. 1–20.

Riedl, Caroline

- [1] The DOE/NSF Nuclear Science Advisory Committee Working Group, *The Frontiers of Nuclear Science: A Long-Range Plan*. arXiv:0809.3137 (2007).
- [2] COMPASS Collaboration, First measurement of transverse-spin-dependent azimuthal asymmetries in the Drell–Yan process. *PRL*, 119 (2017), 112002.
- [3] Riahi, H., et al., FTS3: Quantitative Monitoring. *Journal of Physics*, Conference Series 664 (2015), 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] Klimentov, A., et al., Beyond ATLAS: A Scalable Workload Management System for Data Intensive Science. *International Symposium on Grid and Clouds Academia Sinica* (Taipei Taiwan, March 2014).
- [6] Böhlen, T.T., et al., The FLUKA Code: Developments and Challenges for High Energy and Medical Applications. *Nuclear Data Sheets*, 120 (2014), pp. 211–214.
- [7] Sjoestrand, T., et al., An Introduction to PYTHIA 8.2. *Computer Physics Communications*, 191 (2015), pp. 159–177.

Schleife, Andre (1)

- [1] Sun, Y., Q. Wu, and G. Shi, Graphene based new energy materials. *Energy & Environmental Science*, 4:4 (2011), DOI:10.1039/C0EE00683A.
- [2] Surwade, S.P., et al., Water desalination using nanoporous single-layer graphene. *Nature Nanotechnology*, 10:5 (2015), DOI:10.1038/nnano.2015.37.
- [3] Bangert, U., et al., Ion Implantation of Graphene—Toward IC Compatible Technologies. *Nano Letters*, 13:10 (2013), DOI:10.1021/nl402812y.
- [4] Vicarelli, L., et al., Controlling Defects in Graphene for Optimizing the Electrical Properties of Graphene Nanodevices. *ACS Nano*, 9:4 (2015), DOI:10.1021/acsnano.5b01762.
- [5] Schleife, A., et al., Plane-wave pseudopotential implementation of explicit integrators for time-dependent Kohn–Sham equations in large-scale simulations. *Journal of Chemical Physics*, 137:22 (2012), DOI:10.1063/1.4758792.
- [6] Ullrich, C., *Time-Dependent Density-Functional Theory: Concepts and Applications*. Oxford University Press (2011), DOI:10.1093/acprof:oso/9780199563029.001.0001.

Schleife, Andre (2)

- [1] Kresse, G., and J. Furthmüller, Efficient iterative schemes for *ab initio* total-energy calculations using a plane-wave basis set. *Phys. Rev. B*, 54 (1996), p. 11169.
- [2] Gajdoš, M., et al., Linear optical properties in the projector-augmented wave methodology. *Phys. Rev. B*, 73 (2006), p. 045112.

- [3] Umetsu, Y., K. Fukamichi, and A. Sakuma, Electrical and Magnetic Properties, and Electronic Structures of Pseudo-Gap-Type Antiferromagnetic L_{1₀}-Type MnPt Alloys. *Materials Transactions*, 47 (2006), p. 2.

Taha, Ahmed (1)

Idkaidek, Ashraf

- [1] Idkaidek, A., and I. Jasiuk, Toward high-speed 3D nonlinear soft tissue deformation simulations using Abaqus software. *J Robot Surg*, 9:4 (2015), pp. 299–310.
- [2] Idkaidek, A., and I. Jasiuk, Cortical bone fracture analysis using XFEM—case study. *Int J Numer Method Biomed Eng*, 33:4 (2017), DOI:10.1002/cnm.2809.
- [3] Idkaidek, A., S. Koric, and I. Jasiuk, Fracture analysis of multi-osteon cortical bone using XFEM. *Computational Mechanics*, 62:2 (2018), pp. 171–184.

Sabet, Fereshteh A.

- [1] Niebur, G.L., et al., High-resolution finite element models with tissue strength asymmetry accurately predict failure of trabecular bone. *Journal of Biomechanics*, 33:12 (2000), pp. 1575–1583.
- [2] Bayraktar, H.H., and T.M. Keaveny, A computational investigation of the nonlinear behavior of human trabecular bone. *Transactions of the 12th annual Pre-ORS Symposium on Computational Methods in Orthopaedic Biomechanics 2* (2004).
- [3] Sabet, F.A., O. Jin, S. Koric, and I. Jasiuk, Nonlinear micro-CT based FE modeling of trabecular bone—Sensitivity of apparent response to tissue constitutive law and bone volume fraction. *International journal for numerical methods in biomedical engineering*, 34:4 (2017), p. e2941.

Taha, Ahmed (2)

- [1] Govert, S., et al., Turbulent combustion modelling of a confined premixed methane/air jet flame including heat loss effects using tabulated chemistry. *Appl. Energ.*, 156 (2015), pp. 804–815.
- [2] Govert, S., et al., Heat loss prediction of a confined premixed jet flame using a conjugate heat transfer approach. *Int. J. Heat Mass Transfer*, 107 (2017), pp. 882–894.
- [3] Govert, S., et al., The Effect of Partial Premixing and Heat Loss on the Reacting Flow Field Prediction of a Swirl Stabilized Gas Turbine Model Combustor. *Flow Turbulence and Combustion*, 100:2, pp. 503–534.
- [4] Mira, D., et al., Numerical investigation of a lean premixed swirl-stabilized hydrogen combustor and operational conditions close to flashback. *ASME Turbo Expo 2018*, GT2018–76229.

Thomas, Brian

- [1] World Steel Association, *Steel Statistical Yearbook* (2017).
- [2] Yang, H., S.P. Vanka, and B.G. Thomas, A Hybrid Eulerian–Eulerian Discrete-Phase Model of Turbulent Bubbly Flow. *Journal of Fluids Engineering*, 140:10 (2018), pp. 101202–101212.

- [3] Yang, H., S.P. Vanka, and B.G. Thomas, Modeling of Argon Gas Behavior in Continuous Casting of Steel. *Proceedings of TMS 2018* (Phoenix, Ariz., March 11–15, 2018).
- [4] Jin, K., B.G. Thomas, and X. Ruan, Modeling of Measurements of Multiphase Flow and Bubble Entrapment in Steel Continuous Casting. *Metallurgical and Materials Transaction B*, 47B (2016), pp. 548–565.
- [5] Jin, K., S.P. Vanka, and B.G. Thomas, Large Eddy Simulations of Electromagnetic Braking Effects on Argon Bubble Transport and Capture in a Steel Continuous Casting Mold. *Metallurgical and Materials Transaction B*, 49B (2018), pp. 1360–1377.
- [6] Vanka, S.P., A.F. Shinn, and K.C. Sahu, Computational Fluid Dynamics Using Graphics Processing Units: Challenges and Opportunities. *Proceedings of the ASME 2011 IMECE Conference* (Denver, Colo., 2011).
- [7] Timmel, K., et al., Visualization of liquid metal two-phase flows in a physical model of the continuous casting process of steel. *Metallurgical and Materials Transaction B*, 46B (2014), pp. 700–710.

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., Flow and transport in regions with aquatic vegetation. *Ann. Rev. Fluid Mech.*, 44 (2012), pp. 123–142.
- [4] Tinoco, R., and G. Coco, A laboratory study on sediment resuspension within arrays of rigid cylinders. *Adv. Water Res.*, 92 (2016), pp. 1–9.
- [5] Fischer, P.F., J.W. Lottes, and S.G. Kerkemeier, Nek 5000 webpage (2008), <http://nek5000.mcs.anl.gov>.
- [6] Deville, M., P. Fischer, and E. Mund, *High-Order Methods for Incompressible Fluid Flow*. Cambridge University Press, Cambridge, England (2002).
- [7] Dutta, S., M.I. Cantero, and M.H. Garcia, Effect of self-stratification on sediment diffusivity in channel flows and boundary layers: a study using direct numerical simulations. *Earth Surface Dynamics*, 2:2 (2014), p. 419.
- [8] Jeong, J., and F. Hussain, On the identification of a vortex. *Journal of Fluid Mechanics*, 285 (1995), pp. 69–94.

Yeung, Pui-kuen

- [1] Frisch, U., *Turbulence: The Legacy of A.N. Kolmogorov*. Cambridge University Press (Cambridge, UK, 1995).
- [2] Yakhot, V., and K.R. Sreenivasan, Anomalous scaling of structure functions and dynamic constraints on turbulence simulations. *Journal of Statistical Physics*, 121 (2005), pp. 823–841.

- [3] Yeung, P.K., K.R. Sreenivasan, and S.B. Pope, Effects of finite spatial and temporal resolution on extreme events in direct numerical simulations of incompressible isotropic turbulence. *Physical Review Fluids*, 3 (2018), p. 064603.
- [4] Yeung, P.K., X.M. Zhai, and K.R. Sreenivasan, Extreme events in computational turbulence. *Proceedings of the National Academy of Sciences*, 112 (2015), pp. 12633–12638.
- [5] Yeung, P.K., D.A. Donzis, and K.R. Sreenivasan, Dissipation, enstrophy and pressure statistics in turbulence simulations at high Reynolds numbers. *Journal of Fluid Mechanics*, 700 (2012), pp. 5–15.
- [6] Sreenivasan, K.R., Fractals and multifractals in fluid turbulence. *Annual Review of Fluid Mechanics*, 23 (1991), pp. 539–600.
- [7] Buaria, D., A. Pumir, E. Bodenschatz, and P.K. Yeung, Extreme velocity gradients in turbulent flows. Submitted to *Physical Review X* (2018).

COMPUTER SCIENCE & ENGINEERING

Campbell, Roy H

- [1] Abadi, M., et al., TensorFlow: A System for Large-Scale Machine Learning. *Proceedings of the 12th USENIX Symposium on Operating Systems Design and Implementation* (USENIX, Savannah, Ga., November 2–4, 2016), Vol. 16, pp. 265–283.
- [2] Iandola, F.N., et al., FireCaffe: Near-linear acceleration of deep neural network training on compute clusters. *Proceedings of the 2016 IEEE Conference on Computer Vision and Pattern Recognition*, (IEEE, Las Vegas, Nev., June 27–30, 2016), pp. 2592–2600.
- [3] Keuper, J., Distributed training of deep neuronal networks: Theoretical and practical limits of parallel scalability. CoRR abs/1609.06870 (2016).
- [4] Barnett, M., and L. Shuler, et al., Interprocessor collective communication library (InterCom). *Proceedings of the 1994 Scalable High-Performance Computing Conference* (IEEE), pp. 357–364.
- [5] Thakur, R., and W.D. Gropp, Improving the performance of collective operations in MPICH. *Proceedings of the European Parallel Virtual Machine/Messsage Passing Interface Users' Group Meeting* (2003), Springer, pp. 257–267.
- [6] Sergeev, A., and M.D. Balso, Horovod: fast and easy distributed deep learning in TensorFlow. CoRR abs/1802.05799 (2018).

Cox, Donna

- [1] Hitchcock, A., C.N. Hunter, and M. Sener, Determination of Cell Doubling Times from the Return-on-Investment Time of Photosynthetic Vesicles Based on Atomic Detail Structural Models. *J. Phys. Chem. B*, 121:15 (2017), DOI: 10.1021/acs.jpcc.7b06180.
- [2] Stone, J., et al., Atomic detail visualization of photosynthetic membranes with GPU-accelerated ray tracing. *Parallel Computing*, 55 (2016), DOI:10.1016/j.parco.2015.10.015.

- [3] Cartron, M., et al., Integration of energy and electron transfer processes in the photosynthetic membrane of *Rhodospirillum rubrum*. *Biochimica et Biophysica Acta (BBA)–Bioenergetics*, 1837:10 (2014), DOI:10.1016/j.bbabi.2014.02.003.
- [4] Salmon, J., R.M. Canup, Accretion of the Moon from non-canonical discs. *Philosophical Transactions of the Royal Society A*, 372:20130256 (2014), DOI:10.1098/rsta.2013.0256.
- [5] O'Shea, B., et al., Probing the Ultraviolet Luminosity Function of the Earliest Galaxies. *The Astrophysical Journal Letters*, 807:1 (2015), DOI:10.1088/2041-8205/807/1/L12.

Jha, Shantenu

- [1] Shaw, D.E., et al., Atomic-Level Characterization of the Structural Dynamics of Proteins. *Science*, 330:6002 (2010), pp. 341–346.

Koric, Seid

- [1] Gupta A. WSMP. Watson Sparse matrix package (Part-I: Direct solution of symmetric sparse systems). Technical Report RC 21866, IBM T. J. Watson Research Center. Yorkton Heights, NY; 2015.
- [2] Koric, S. and A. Gupta, "Sparse Matrix Factorization in the Implicit Finite Element Method on Petascale Architecture," *Computer Methods in Applied Mechanics and Engineering*, 2016 v.32,281-292, 2016.
- [3] HPC Wire <https://www.hpcwire.com/off-the-wire/hpcwire-reveals-winners-2017-readers-editors-choice-awards-sc17-conference-denver/>
- [4] Vázquez M. et al., Alya: Multiphysics engineering simulation towards exascale. *J. Comput. Sci.*, 14:15-27, 2016
- [5] Casoni E. et al., Alya: Computational Solid Mechanics for Supercomputers. *Arch Computat Methods Eng* DOI 10.1007/s11831-014-9126-8

Nagi, Rakesh

- [1] Poore, A.B., Multidimensional assignment formulation of data association problems arising from multitarget and multisensor tracking. *Computational Optimization and Applications*, 3:1 (1994), pp. 27–57.
- [2] Bandelt, H.J., Y. Crama, and F.C. Spieksma, Approximation algorithms for multi-dimensional assignment problems with decomposable costs. *Discrete Applied Mathematics*, 49 (1994), pp. 25–50.
- [3] Tauer, G., R. Nagi, and M. Sudit, The graph association problem: mathematical models and a Lagrangian heuristic. *Naval Research Logistics*, 60:3 (2013), pp. 251–268.
- [4] Tauer, G., and R. Nagi, A map-reduce Lagrangian heuristic for multidimensional assignment problems with decomposable costs. *Parallel Computing*, 39:11 (2013), pp. 653–668.
- [5] Date, K., and R. Nagi, GPU-accelerated Hungarian algorithms for the Linear Assignment Problem. *Parallel Computing*, 57 (2016), pp. 52–72.

Olson, Luke

- [1] Reisner, A., L. Olson, and D. Moulton, Scaling Structured Multigrid to 500K+ Cores through Coarse-Grid Redistribution. *SIAM Journal on Scientific Computing*, to appear (2018).

West, Matthew

- [1] LeCun, Y., F.J. Huang, and L. Bottou, Learning methods for generic object recognition with invariance to pose and lighting. *Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, Vol. 2, pp. II–104, (IEEE Computer Society, Washington, D.C., USA, June 27–July 2, 2004).
- [2] Graves, A., A.R. Mohamed, and G. Hinton, Speech recognition with deep recurrent neural networks. *Proceedings of the 2013 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 6645–6649, (IEEE, Vancouver, Canada, May 26–31, 2013).
- [3] Stolfo, S., et al., Credit card fraud detection using meta-learning: issues and initial results. In *AAAI-97 Workshop on Fraud Detection and Risk Management*, (AAAI, Providence, Rhode Island, July 27–28, 1997).
- [4] Guzella, T.S., and W.M. Caminhas, A review of machine learning approaches to spam filtering. *Expert Systems with Applications*, 36:7 (2009), pp. 10206–10222.
- [5] Lops, P., M.D. Gemmis, and G. Semeraro, Content-based recommender systems: state of the art and trends. *Recommender Systems Handbook*, Springer (2011), pp. 73–105.
- [6] Swersky, K., J. Snoek, and R.P. Adams, Multi-task Bayesian optimization. *Advances in Neural Information Processing Systems 26* (2013), pp. 2004–2012.
- [7] Rasmussen, C.E., and C.K.I. Williams, *Gaussian Processes for Machine Learning*, Vol. 1. MIT Press, Cambridge, Massachusetts (2006).
- [8] Snoek, J., H. Larochelle, and R.P. Adams, Practical Bayesian optimization of machine learning algorithms. In *Advances in Neural Information Processing Systems* (2012), pp. 2951–2959.
- [9] González, J., et al., Batch Bayesian optimization via local penalization. *Artificial Intelligence and Statistics* (2016), pp. 648–657.

BIOLOGY, CHEMISTRY, & HEALTH

Aksimentiev, Aleksei (1)

- [1] Ngo, T.T., et al., Asymmetric unwrapping of nucleosomes under tension directed by DNA local flexibility. *Cell*, 160:6 (2015), pp. 1135–1144.
- [2] Tillo, D., and T.R. Hughes, G+C content dominates intrinsic nucleosome occupancy. *BMC bioinformatics*, 10:1 (2009), p. 442.
- [3] Clapier, C.R., and B.R. Cairns, The biology of chromatin remodeling complexes. *Annual review of biochemistry*, 78 (2009), pp. 273–304.

- [4] Kalé, L., et al., NAMD2: greater scalability for parallel molecular dynamics. *Journal of Computational Physics*, 151:1 (1999), pp. 283–312.
- [5] Phillips, J.C., et al., Scalable molecular dynamics with NAMD. *Journal of computational chemistry*, 26:16 (2005), pp. 1781–1802.
- [6] Tjaden, G., and G.M. Coruzzi. A novel AT-rich DNA binding protein that combines an HMG I-like DNA binding domain with a putative transcription domain. *The Plant Cell*, 6:1 (1994), pp. 107–118.

Aksimentiev, Aleksei (2)

- [1] Pomorski, T., J.C. Holthuis, A. Herrmann, and G. van Meer, Tracking down lipid flippases and their biological functions. *Journal of cell science*, 117 (2004), pp. 805–813.
- [2] Sahu, S.K., S.N. Gummadi, N. Manoj, and G.K. Aradhya, Phospholipid scramblases: an overview. *Archives of biochemistry and biophysics*, 462 (2007), pp. 103–114.
- [3] Kalé, L., et al., NAMD2: greater scalability for parallel molecular dynamics. *Journal of Computational Physics*, 151 (1999), pp. 283–312.
- [4] Phillips, J.C., et al., Scalable molecular dynamics with NAMD. *Journal of computational chemistry*, 26 (2005), pp. 1781–1802.
- [5] 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 of the U.S.A.*, 110 (2013), pp. 20099–20104.
- [6] Göpfrich, K., et al., Large-Conductance Transmembrane Porin Made from DNA Origami. *ACS Nano*, 10 (2016), pp. 8207–8214.
- [7] Göpfrich, K., et al., Ion Channels Made from a Single Membrane-Spanning DNA Duplex. *Nano Lett.*, 16 (2016), pp. 4665–4669.
- [8] Yoo, J., and A. Aksimentiev, Molecular Dynamics of Membrane-Spanning DNA Channels: Conductance Mechanism, Electro-Osmotic Transport and Mechanical Gating. *J. Phys. Chem. Lett.*, 6:23 (2015), pp. 4680–4687.

Aluru, Narayana

- [1] Wells, D.B., et al., Assessing Graphene Nanopores for Sequencing DNA. *Nano letters*, 12:8 (2012), pp. 4117–4123.
- [2] Branton, D., et al., The Potential and Challenges of Nanopore Sequencing. *Nat. Biotech.*, 26 (2008), pp. 1146–1153.

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] Papaleo, E., et al., The role of protein loops and linkers in conformational dynamics and allostery. *Chemical Reviews*, 116 (2016), pp. 6391–6423.
- [2] Liberles, D.A., et al., The interface of protein structure, protein biophysics, and molecular evolution. *Protein Science*, 21 (2012), pp. 769–85.
- [3] Marsh, J.A., and S.A. Teichmann, Protein flexibility facilitates quaternary structure assembly and evolution. *PLoS Biology*, 12 (2014), DOI:10.1371/journal.pbio.1001870.
- [4] Goldman, A.D., J.A. Baross, and R. Samudrala, The enzymatic and metabolic capabilities of early life. *PLoS One*, 7 (2012), DOI:10.1371/journal.pone.0039912.
- [5] O'Leary, N.A., et al., Reference sequence (RefSeq) database at NCBI: current status, taxonomic expansion, and functional annotation. *Nucleic Acids Research*, 44 (2016), DOI:10.1093/nar/gkv1189.
- [6] Kim, K.M., and G. Caetano-Anolles, The evolutionary history of protein fold families and proteomes confirms that the archaeal ancestor is more ancient than the ancestors of other superkingdoms. *BMC Evolutionary Biology*, 12 (2012), DOI:10.1186/1471-2148-12-13.
- [7] Eddy, S.R., Accelerated profile HMM searches. *PLoS Computational Biology*, 7 (2011), DOI:10.1371/journal.pcbi.1002195.
- [8] Bonet, J., et al., ArchDB 2014: structural classification of loops in proteins. *Nucleic Acids Research*, 42 (2014), DOI:10.1093/nar/gkt1189.
- [9] Kim, K.M., et al., Protein domain structure uncovers the origin of aerobic metabolism and the rise of planetary oxygen. *Structure*, 20 (2012), pp. 67–76.
- [10] Madej, T., et al., MMDB and VAST+: tracking structural similarities between macromolecular complexes. *Nucleic Acids Research*, 42 (2014), pp. D297–D303.
- [11] Barnoud, J., et al., PBxplore: a tool to analyze local protein structure and deformability with Protein Blocks. *PeerJ*, 5 (2017), DOI:10.7717/peerj.4013.
- [12] Gu, Y., D.-W. Li, and R. Brüschweiler, Decoding the mobility and time scales of protein loops. *Journal of Chemical Theory and Computation*, 11 (2015), pp. 1308–1314.
- [13] Vogel, C., Function annotation of SCOP domain superfamilies: <http://supfam.org/SUPERFAMILY/function.html>.

Carnevale, Vincenzo

- [1] Institute of Medicine Report: *Relieving Pain in America—A Blueprint for Transforming Prevention, Care, Education and Research*. The National Academies Press (2011).
- [2] Trevisani, M., and R. Gatti, TRPV1 Antagonists as Analgesic Agents. *The Open Pain Journal*, 6 (2013), pp. 108–118.
- [3] Lee, Y.C., N.J. Nassikas, and D.J. Clauw, The role of the central nervous system in the generation and maintenance of chronic pain in rheumatoid arthritis, osteoarthritis and fibromyalgia. *Arthritis Research and Therapy*, 13 (2011), p. 211.

- [4] Baliki, M.N., P.Y. Geha, A.V. Apkarian, and D.R. Chialvo, Beyond Feeling: Chronic Pain Hurts the Brain, Disrupting the Default-Mode Network Dynamics. *The Journal of Neuroscience*, 28:6 (2008), pp. 1398–1403.
- [5] Interagency Pain Research Coordinating Committee, *National Pain Strategy: A comprehensive population health-level strategy for pain*. NIH publications (2015).
- [6] Jorge, L.L., C.C. Feres, and V.E.P. Teles, Topical preparations for pain relief: efficacy and patient adherence. *Journal of Pain Research*, 4 (2011), pp. 11–24.
- [7] Herndon, C.M., et al., Management of Chronic Nonmalignant Pain with Nonsteroidal Anti-inflammatory Drugs. *Pharmacotherapy*, 28:6 (2008), pp. 788–805.
- [8] American Society of Regional Anesthesia and Pain, Treatment options for chronic pain (2016).
- [9] Nilius, B., and A. Szallasi, Transient Receptor Potential Channels as Drug Targets: From the Science of Basic Research to the Art of Medicine. *Pharmacol Rev*, 66 (2014), pp. 676–814.
- [10] Brito, R., et al., TRPV1: A Potential Drug Target for Treating Various Diseases. *Cells*, 3 (2014), pp. 517–545.
- [11] Carnevale, V., and T. Rohacs, TRPV1: A Target for Rational Drug Design. *Pharmaceuticals*, 9 (2016), p. 52.
- [12] Cao, E., M. Liao, Y. Cheng, and D. Julius, TRPV1 structures in distinct conformations reveal activation mechanisms. *Nature*, 504:7478 (2013), pp. 113–1188.
- [13] 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.
- [14] Elokely, K, et al., Understanding TRPV1 activation by ligands: Insights from the binding modes of capsaicin and resiniferatoxin. *Proc Natl Acad Sci U.S.A.*, 113:2 (2016), pp. E137–E145.
- [15] Mackerell, Jr., A.D., M. Feig, and C.L. Brooks III, 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.
- [16] Jorgensen, W.L., et al., Comparison of simple potential functions for simulating liquid water. *J Chem Phys.*, 79:2 (1983), p. 926.
- [17] 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.
- [18] Phillips, J.C., et al., Scalable molecular dynamics with NAMD. *J Comput Chem*, 26:16 (2005), pp. 1781–1802.
- [19] 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.

- [20] Laio, A., et al., Assessing the Accuracy of Metadynamics. *J Phys Chem B*, 109:14 (2005), pp. 6714–6721.
- [21] 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.
- [22] Steinberg, X., et al., Conformational dynamics in TRPV1 channels reported by an encoded coumarin amino acid. *eLife*, 6 (2017).
- [23] Kasimova, et al., A consistent picture of TRPV1 activation emerges from molecular simulations and experiments. *bioRxiv*, (2018), DOI:10.1101/310151.
- [24] Schmiede, P., M. Fine, G. Blobel, and X. Li, Human TRPML1 channel structures in open and closed conformations. *Nature*, 550:7676 (2017), p. 366.
- [25] Woll, K.A., et al., Sites Contributing to TRPA1 Activation by the Anesthetic Propofol Identified by Photoaffinity Labeling. *Biophysical Journal*, 113:10 (2017), pp. 2168–2172.
- [26] Ton, H.T., et al., Identification of a putative binding site critical for general anesthetic activation of TRPA1. *PNAS*, 114:14 (2017), pp. 3762–3767.
- [27] Raju, S.G., et al., Exploring volatile general anesthetic binding to a closed membrane-bound bacterial voltage-gated sodium channel via computation. *PLoS computational biology*, 9:6 (2013), p. e1003090.
- [28] Kinde, M.N., et al., Fluorine-19 NMR and computational quantification of isoflurane binding to the voltage-gated sodium channel NaChBac. *PNAS*, 113:48 (2016), pp. 13762–13767.

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 U.S.A.*, 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).
- [5] Broide, M.L., et al., Binary-liquid phase separation of lens protein solutions. *Proc Natl Acad Sci U.S.A.*, 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.
- [7] Phillips, J.C., et al., Scalable molecular dynamics with NAMD. *J. Comput. Chem.*, 26 (2005), pp. 1781–1802.

- [8] Edgar, R.C., MUSCLE: multiple sequence alignment with high accuracy and high throughput. *Nucleic Acids Res.*, 32:5 (2004), pp. 1792–1797.

Dill, Ken

- [1] Maier, J.A., et al., ff14SB: Improving the Accuracy of Protein Side Chain and Backbone Parameters from ff99SB. *J. Chem. Theory Comput.*, 11 (2015), pp. 3696–3713.
- [2] Perez, A., J.L. MacCallum, and K. Dill, Accelerating molecular simulations of proteins using Bayesian inference on weak information. *Proc. Natl. Acad. Sci. U.S.A.*, 112 (2015), pp. 11846–11851.
- [3] 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. U.S.A.*, 112 (2015), pp. 6985–6990.
- [4] Moulton, J.A., Decade of CASP: progress, bottlenecks and prognosis in protein structure prediction. *Curr. Opin. Struct. Biol.*, 15 (2005), pp. 285–289.
- [5] Perez, A., J.A. Morrone, E. Brini, J.L. MacCallum, and K. Dill, Blind protein structure prediction using accelerated free-energy simulations. *Science Advances*, 2 (2016), pp. e1601274–e1601274.
- [6] Perez, A., et al., MELD-Path Efficiently Computes Conformational Transitions, Including Multiple and Diverse Paths. *J. Chem. Theory Comput.*, 14 (2018), pp. 2109–2116.

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] Gazzola, M., L. Dudte, A. McCormick, and L. Mahadevan. Forward and inverse problems in the mechanics of soft filaments. Accepted in *Royal Society Open Source*, (2018), arXiv:1607.00430v2.
- [4] Pagan-Diaz, G.J., et al., Simulation and fabrication of stronger, larger and faster walking biohybrid machines. *Advanced Functional Materials*, (2018), DOI:10.1002/adfm.201801145.

Gropp, Bill

- [1] Nguyen, N., et al., TIPP: taxonomic identification and phylogenetic profiling. *Bioinformatics*, 30:24 (2014), DOI:10.1093/bioinformatics/btu721.
- [2] Mirarab, S., et al. PASTA: Ultra-Large Multiple Sequence Alignment for Nucleotide and Amino-Acid Sequences. *Journal of Computational Biology*, 22:5 (2015), pp. 377–386.
- [3] Kozlov, A.M., A.J. Aberer, and A. Stamatakis, ExaML version 3: a tool for phylogenomic analyses on supercomputers. *Bioinformatics*, 31:15 (2015), DOI:10.1093/bioinformatics/btv184.

- [4] Stamatakis, A., RAXML-VI-HPC: maximum likelihood-based phylogenetic analyses with thousands of taxa and mixed models. *Bioinformatics*, 22:21 (2016), pp. 2577–2579.
- [5] Nelesen, S., et al., DACTAL: Divide-And-Conquer Trees (almost) without Alignments. *Bioinformatics*, 28:12 (2012), pp. i274–i282.
- [6] Bansal, M.S., et al., Robinson-Foulds Supertrees. *Algorithms for Molecular Biology*, 5:1 (2010), DOI:10.1186/1748-7188-5-18.
- [7] Jiang, T., P. Kearney, and M. Li. A polynomial time approximation scheme for inferring evolutionary trees from quartet topologies and its application. *SIAM Journal on Computing*, 30:6 (2001), pp. 1942–1961.
- [8] Katoh, K., and D.M. Standley, MAFFT Multiple Sequence Alignment Software Version 7: Improvements in Performance and Usability. *Molecular Biology and Evolution*, 30:4 (2013), p. 772.
- [9] Simonsen, M., T. Mailund, and C.N.S. Pedersen, Rapid Neighbour-Joining. *Algorithms in Bioinformatics, WABI 2008, Lecture Notes in Computer Science*, 5251 (2008), Springer Berlin Heidelberg, pp.113–122.

Hirata, So

- [1] Hirata, S., A.E. Doran, P.J. Knowles, and J.V. Ortiz, One-particle many-body Green's function theory: Algebraic recursive definitions, linked-diagram theorem, irreducible-diagram theorem, and general-order algorithms. *J. Chem. Phys.*, 147 (2018), p. 044108.
- [2] Willow, S.Y., K.S. Kim, and S. Hirata, Stochastic evaluation of second-order Dyson self-energies. *J. Chem. Phys.*, 138 (2013), p. 164111.
- [3] Johnson, C.M., A.E. Doran, S.L. Ten-no, and S. Hirata, Monte Carlo explicitly correlated second-order many-body Green's function theory. To be submitted (2018).
- [4] Doran, A.E., and S. Hirata, Monte Carlo MP2 on many graphical processing units. *J. Chem. Theory Comput.*, 12 (2016), pp. 4821–4832.
- [5] 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).
- [6] Doran, A.E., and S. Hirata, Monte Carlo third-order many-body Green's function theory. In preparation (2018).
- [7] 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.*, 81:25 (1977), pp. 2340–2361.
- [2] Alfonsi, A., et al., Adaptive simulation of hybrid stochastic and deterministic models for biochemical systems. *ESIAM Proceeding*, 14 (2005), pp. 1–13.

- [3] Cao, Y., D. Gillespie, and L. Petzold, The slow-scale stochastic simulation algorithm. *J. Chem. Phys.*, 122:1 (2005), p. 0141116.
- [4] Jahnke, T., and M. Kreim, Error bound for piecewise deterministic processes modeling stochastic reaction systems. *SIAM Multiscale Model. Simul.*, 10 (2012), pp. 1119–1131.
- [5] Roberts, E., J. 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., et al., A Problem Solving Environment for Stochastic Biological Simulations. *PyHPC'13* (2013), DOI:10.13140/2.1.3207.7440.
- [7] Earnest, T., et al., Challenges of Integrating Stochastic Dynamics and Cryo-Electron Tomograms in Whole-Cell Simulations. *J. Phys. Chem. B*, 121:15 (2017), pp. 3871–3881.

Makri, Nancy

- [1] Feynman, R.P., Space-time approach to non-relativistic quantum mechanics. *Rev. Mod. Phys.*, 20 (1948), p. 367.
- [2] Lambert, R., and N. Makri, Quantum–classical path integral: Classical memory and weak quantum nonlocality. *J. Chem. Phys.*, 137 (2012), p. 22A552.
- [3] Lambert, R., and N. Makri, Quantum–classical path integral: Numerical formulation. *J. Chem. Phys.*, 137 (2012), p. 22A553.
- [4] Banerjee, T., and N. Makri, Quantum–classical path integral with self-consistent solvent propagators. *J. Phys. Chem.*, 117 (2013), p. 13357.
- [5] Makri, N., Quantum–classical path integral: A rigorous approach to condensed phase dynamics. *International Journal of Quantum Chemistry*, 115 (2015), p. 1209.
- [6] Walters, P. L., and N. Makri, Iterative quantum–classical path integral with dynamically consistent state hopping. *J. Chem. Phys.*, 144 (2016), p. 044108.
- [7] Makri, N., Exploiting classical decoherence in dissipative quantum dynamics: Memory, phonon emission, and the blip sum. *Chem. Phys. Lett.*, 593 (2014), p. 93.
- [8] Makri, N., Blip decomposition of the path integral: Exponential acceleration of real-time calculations for quantum dissipative systems. *J. Chem. Phys.*, 141 (2014), p. 134117.
- [9] Makri, N., Iterative blip-summed path integral for quantum dynamics in strongly dissipative environments. *J. Chem. Phys.*, 146 (2017), p. 134101.
- [10] Makri, N., Blip-summed quantum–classical path integral with cumulative quantum memory. *Faraday Disc.*, 195 (2016), pp. 81–92.
- [11] Bose, A., and N. Makri, Evaluation of the Wigner distribution via classical adiabatic switching. *J. Chem. Phys.*, 143 (2015), p. 114114.
- [12] Bose, A., and N. Makri, Path integral evaluation of the Wigner phase space distribution. *J. Chem. Phys.*, to be submitted.
- [13] Walters, P. L., and N. Makri, Quantum–classical path integral simulation of the ferrocene-ferrocenium charge transfer in liquid hexane. *J. Phys. Chem. Lett.*, 6 (2015), pp. 4959–4965.

- [14] Walters, P. L., A. Bose, and N. Makri, Solvent effects on the rate of the ferrocene-ferrocenium electron transfer reaction: results of accurate quantum–classical path integral simulations. *J. Phys. Chem.*, to be submitted.

Masud, Arif

- [1] Balay, S., W.D. Gropp, L. Curfman McInnes, and B.F. Smith, Efficient Management of Parallelism in Object Oriented Numerical Software Libraries. in *Modern Software Tools in Scientific Computing*, Birkhäuser Press (1997), pp. 163–202.

Perilla, Juan

- [1] Perilla, J.R., et al., All-atom molecular dynamics of virus capsids as drug targets. *Journal of Physical Chemistry Letters*, 7:10 (2016), pp. 1836–1844.
- [2] Phillips, J.C., et al., Scalable molecular dynamics with NAMD. *Journal of Computational Chemistry*, 26:16 (2005), pp. 1781–1802.
- [3] 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.
- [4] Goh, B.C., et al., Atomic modeling of an immature retroviral lattice using molecular dynamics and mutagenesis. *Structure*, 23:8 (2015), pp. 1414–1425.

Roux, Benoit

- [1] Deng, Y., and B. Roux, Computations of standard binding free energies with molecular dynamics simulations. *J. Phys. Chem. B*, 113:8 (2009), pp. 2234–2246.
- [2] Maragliano, L., E. Vanden-Eijnden, and B. Roux, Free Energy and Kinetics of Conformational Transitions from Voronoi Tessellated Milestoning with Restraining Potentials. *J. Chem. Theor. Comp.*, 5:10 (2009), pp. 2594–2589.
- [3] Jiang, W., et al., Calculation of free energy landscape in multidimensions with Hamiltonian-exchange umbrella sampling on petascale supercomputer. *J. Chem. Theor. Comp.*, 8:11 (2012), pp. 4672–4680.
- [4] Jiang, W., et al., Generalized Scalable Multiple Copy Algorithms for Molecular Dynamics Simulations in NAMD. *Comput. Phys. Commun.*, 185:3 (2014), pp. 908–916.
- [5] Meng, Y., and B. Roux, Locking the Active Conformation of c-Src Kinase through the Phosphorylation of the Activation Loop. *J. Mol. Biol.*, 426:2 (2014), pp. 423–435.
- [6] Ostmeier, J., et al., Recovery from slow inactivation in K⁺ channels is controlled by water molecules. *Nature*, 501 (2013), pp. 121–124.

Shukla, Diwakar (1)

- [1] Zhao, L.-H., et al., Crystal structures of two phytohormone signal-transducing α/β hydrolases: karrikin-signaling KAI2 and strigolactone-signaling DWARF14. *Cell research*, 23 (2013), pp. 436–439.
- [2] Yao, R., et al., DWARF14 is a non-canonical hormone receptor for strigolactone. *Nature* (2016).

Shukla, Diwakar (2)

- [1] Heaviest pumpkin. *Guinness World Records*, <http://www.guinnessworldrecords.com/world-records/heaviest-pumpkin>.
- [2] Savage, J.A., D.F. Haines, and N.M. Holbrook, The making of giant pumpkins: how selective breeding changed the phloem of *Cucurbita maxima* from source to sink. *Plant Cell Environ.*, 38 (2015), pp. 1543–1554.
- [3] Chen, L.-Q., SWEET sugar transporters for phloem transport and pathogen nutrition. *New Phytol.*, 201 (2014), pp. 1150–1155.
- [4] Case, D., et al., Amber 14 (2014).
- [5] Roe, D.R., and T.E. Cheatham, PTRAJ and CPPTRAJ: Software for Processing and Analysis of Molecular Dynamics Trajectory Data. *J. Chem. Theory Comput.*, 9 (2013), pp. 3084–3095.
- [6] Nguyen, H., D.A. Case, and A.S. Rose, NGLview-interactive molecular graphics for Jupyter notebooks. *Bioinforma. Oxf. Engl.*, 34 (2018), pp. 1241–1242.
- [7] McGibbon, R.T., et al., MDTraj: A Modern Open Library for the Analysis of Molecular Dynamics Trajectories. *Biophys. J.*, 109 (2015), pp. 1528–1532.
- [8] Beauchamp, K.A., et al., MSMBuild2: Modeling Conformational Dynamics on the Picosecond to Millisecond Scale. *J. Chem. Theory Comput.*, 7 (2011), pp. 3412–3419.
- [9] Metzner, P., D. Schütte, and E. Vanden-Eijnden, Transition Path Theory for Markov Jump Processes. *Multiscale Model. Simul.*, 7 (2009), pp. 1192–1219.
- [10] Humphrey, W., A. Dalke, and K. Schulten, VMD: visual molecular dynamics. *J. Mol. Graph.*, 14 (1996), pp. 33–38.
- [11] *The PyMOL Molecular Graphics System*, Version 1.8. Schrödinger, LLC.
- [12] Tao, Y., et al., Structure of a eukaryotic SWEET transporter in a homotrimeric complex. *Nature*, 527 (2015), pp. 259–263.

Soltesz, Ivan

- [1] Carnevale, N., and M. Hines, *The NEURON Book*. Cambridge University Press (Cambridge, UK), 2006.
- [2] Hines, M., and N. Carnevale, Computer modeling methods for neurons. In: *The Handbook of Brain Theory and Neural Networks*, edited by M. Arbib, MIT Press (Cambridge, Mass.), 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.
- [4] Bezaire, M., et al., Interneuronal mechanisms of hippocampal theta oscillations in a full-scale model of the rodent CA1 circuit. *Elife* (2016), DOI:10.7554/eLife.18566.
- [5] Schneider, C.J., H. Cuntz, and I. Soltesz, Linking macroscopic with microscopic neuroanatomy using synthetic neuronal populations. *PLoS Comput Biol.*, 10:10 (2014), DOI:10.1371/journal.pcbi.1003921.

- [6] Senzai, Y., and G. Buzsaki, Physiological Properties and Behavioral Correlates of Hippocampal Granule Cells and Mossy Cells. *Neuron*, 93:3 (2017), pp. 691–704.e5.
- [7] Morrison, A., A. Aertsen, and M. Diesmann, Spike-timing-dependent plasticity in balanced random networks. *Neural Comput.*, 19:6 (2007), pp. 1437–67.

Srinivasan, Ashok

- [1] Helbing, D., and P. Molnar, Social force model for pedestrian dynamics. *Physical review E*, 51:5 (1995), p. 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.
- [4] Chunduri, S., et al., Parallel Low Discrepancy Parameter Sweep for Public Health Policy. *IEEE/ACM International Symposium on Cluster, Cloud, and Grid Computing* (2018).
- [5] Suchard, M.A., et al., Bayesian phylogenetic and phylodynamic data integration using BEAST 1.10. *Virus Evolution*, 4 (2018), DOI:10.1093/ve/voy016.

Sutton, Brad

- [1] Johnson, C., et al., Magnetic Resonance Elastography of the Brain Using Multishot Spiral Readouts with Self-Navigated Motion Correction. *Magn Reson Med*, 70:2 (2013), DOI:10.1002/mrm.24473.
- [2] Johnson, C., et al., 3D Multislab, Multishot Acquisition for Fast, Whole-Brain MR Elastography with High Signal-to-Noise Efficiency. *Magn Reson Med*, 71:2 (2014), DOI:10.1002/mrm.25065.
- [3] Johnson, C., et al., Brain MR elastography with multiband excitation and nonlinear motion-induced phase error correction. *Proceedings of the 24th Annual Meeting of the International Society for Magnetic Resonance in Medicine* (Singapore, 2016), p. 1951.
- [4] McGarry, M., et al., Multiresolution MR Elastography Using Nonlinear Inversion. *Med Phys*, 39:10 (2012), DOI:10.1118/1.4754649.
- [5] McGarry, M., et al., Including Spatial Information in Nonlinear Inversion MR Elastography Using Soft Prior Regularization. *IEEE T Med Imaging*, 32:10 (2013), DOI:10.1109/TMI.2013.2268978.
- [6] Johnson, C., et al., Local Mechanical Properties of White Matter Structures in the Human Brain. *NeuroImage*, 79 (2013), DOI:10.1016/j.neuroimage.2013.04.089.
- [7] Anderson, A., et al., Observation of Direction-Dependent Mechanical Properties in the Human Brain with Multi-Excitation MR Elastography. *J Mech Behav Biomed*, 59 (2016), DOI:10.1016/j.jmbm.2016.03.005.
- [8] Johnson, C., et al., Viscoelasticity of Subcortical Gray Matter Structures. *Hum Brain Mapp*, 37:12 (2016), DOI:10.1002/hbm.23314.

- [9] Toledo, J., et al., Contribution of cerebrovascular disease in autopsy confirmed neurodegenerative disease cases in the National Alzheimer's Coordinating Centre. *Brain*, 136:9 (2013), pp. 2697–2706.
- [10] Richard, E., et al., Morphometric Changes in the Cortical Microvascular Network in Alzheimer's Disease. *Journal of Alzheimer's Disease*, 22:3 (2010), pp. 811–818.
- [11] Fang, H.C.H., Observations on Aging Characteristics of Cerebral Blood Vessels, Macroscopic and Microscopic Features. *Neurobiology of Aging*, 1st ed. (Raven Press, New York, New York, 1976), pp. 155–166.
- [12] Cerjanic, A., et al., PowerGrid: An open source library for accelerated iterative magnetic resonance image reconstruction. *Proceedings of the 24th Annual Meeting of the International Society for Magnetic Resonance in Medicine* (Singapore, 2016).

Tagkopoulos, Ilias

- [1] World Health Organization, *Antimicrobial resistance: global report on surveillance* (2014).
- [2] King, R.D., et al., The automation of science. *Science*, 324:5923 (2009), pp. 85–89.
- [3] King, R.D., et al., Functional genomic hypothesis generation and experimentation by a robot scientist. *Nature*, 427:6971 (2004), pp. 247–252.
- [4] Waltz, D., and B.G. Buchanan, Computer science: Automating science. *Science*, 324:5923 (2009), pp. 43–44.
- [5] Suchanek, F.M., G. Kasneci, and G. Weikum, Yago: a core of semantic knowledge. *Proceedings of the 16th international conference on World Wide Web*, (ACM, 2007).
- [6] Carlson, A., et al., Toward an architecture for never-ending language learning. in *AAAI* (Atlanta, Ga., 2010).
- [7] Bollacker, K., et al., Freebase: a collaboratively created graph database for structuring human knowledge. *Proceedings of the 2008 ACM SIGMOD international conference on management of data* (ADM, 2008).
- [8] Singhal, A., Introducing the knowledge graph: things, not strings. Official Google blog, 2012.
- [9] Ernst, P., A. Siu, and G. Weikum, KnowLife: a versatile approach for constructing a large knowledge graph for biomedical sciences. *BMC Bioinformatics*, 16 (2015), p. 157.
- [10] Liekens, A.M., et al., BioGraph: unsupervised biomedical knowledge discovery via automated hypothesis generation. *Genome biology*, 12:6 (2011), p. R57.
- [11] Palmieri, V., et al., The graphene oxide contradictory effects against human pathogens. *Nanotechnology*, 28:15 (2017), p. 152001.
- [12] Tamae, C., et al., Determination of antibiotic hypersensitivity among 4,000 single-gene-knockout mutants of *Escherichia coli*. *J Bacteriol*, 190:17 (2008), pp. 5981–5988.
- [13] Gebser, M., et al., Repair and Prediction (under Inconsistency) in Large Biological Networks with Answer Set Programming. in *KR* (2010).

- [14] Li, Y., et al., A survey on truth discovery. *ACM SigKDD Explorations Newsletter*, 17:2 (2016), pp. 1–16.
- [15] Li, X., et al., Truth finding on the deep web: Is the problem solved? in *Proceedings of the VLDB Endowment* (2012).
- [16] Nickel, M., et al., A review of relational machine learning for knowledge graphs. *Proceedings of the IEEE*, 104:1 (2016), pp. 11–33.
- [17] Melas, I.N., et al., Detecting and removing inconsistencies between experimental data and signaling network topologies using integer linear programming on interaction graphs. *PLoS computational biology*, 9:9 (2013), p. e1003204.
- [18] Soo, V.W., P. Hanson-Manful, and W.M. Patrick, Artificial gene amplification reveals an abundance of promiscuous resistance determinants in *Escherichia coli*. *Proc Natl Acad Sci U.S.A.*, 108:4 (2011), pp. 1484–1489.
- [19] Monnet, D.L., and N. Frimodt-Møller, Antimicrobial-drug use and methicillin-resistant *Staphylococcus aureus*. *Emerging Infectious Diseases*, 7:1 (2001), p. 161.
- [20] Dong, X., et al., Knowledge vault: A web-scale approach to probabilistic knowledge fusion. *Proceedings of the 20th ACM SIGKDD international conference on knowledge discovery and data mining* (ACM, 2014).

Tajkhorshid, Emad (1)

- [1] Bevers, E.M., and P.L. Williamson, Getting to the outer leaflet: physiology of phosphatidylserine exposure at the plasma membrane. *Physiological Reviews*, 96:2 (2016), pp. 605–645.
- [2] Nagata, S., et al., Exposure of phosphatidylserine on the cell surface. *Cell Death and Differentiation*, 23 (2016), pp. 952–961.
- [3] Pomorski, T.G., and A.K. Menon, Lipid somersaults: uncovering the mechanisms of protein-mediated lipid flipping. *Progress in Lipid Research*, 64 (2016), pp. 69–84.
- [4] Brunner, J.D., et al., X-ray structure of a calcium-activated TMEM16 lipid scramblase. *Nature*, 516 (2014), pp. 207–212.
- [5] Paulino, C., et al., Activation mechanism of the calcium-activated chloride channel TMEM16A revealed by cryo-EM. *Nature*, 552 (2017), pp. 421–425.
- [6] Dang, S., et al., Cryo-EM structures of the TME-16A calcium-activated chloride channel. *Nature*, 552 (2017), pp. 426–429.
- [7] Phillips, J.C., et al., Scalable molecular dynamics with NAMD. *Journal of Computational Chemistry*, 26:16 (2005), pp. 1781–1802.

Tajkhorshid, Emad (2)

- [1] Matilla, M.A., and T. Krell, The effect of bacterial chemotaxis on host infection and pathogenicity. *FEMS microbiology reviews*, 42:1 (2017), p. fux052.
- [2] O'Neill, J., Antimicrobial resistance: tackling a crisis for the health and wealth of nations. *The Review on Antimicrobial Resistance*, 20 (2014), pp. 1–16.

- [3] Wadhams, G.H., and J.P. Armitage, Making sense of it all: bacterial chemotaxis. *Nat Rev Mol Cell Biol.*, 5:12 (2004), pp.1024–1037.
- [4] Parkinson, J.S., G.L. Hazelbauer, and J.J. Falke, Signaling and sensory adaptation in *Escherichia coli* chemoreceptors: 2015 update. *Trends Microbiol.*, 23:5 (2015), pp. 257–266.
- [5] Falke, J.J., and K.N. Piasta, Architecture and signal transduction mechanism of the bacterial chemosensory array: Progress, controversies, and challenges. *Curr. Opin. Struct. Biol.*, 29 (2014), pp. 85–94.
- [6] Phillips, J.C., et al., Scalable molecular dynamics with NAMD. *J. Comput. Chem.*, 26:16 (2005), pp. 1781–1802.
- [7] Cassidy, C.K., et al., CryoEM-based hybrid modeling approaches for structure determination. *Curr. Opin. Microbiol.*, 43 (2018), pp. 14–23.
- [8] Stone, J.E., et al., GPU-accelerated molecular dynamics clustering analysis with OpenACC. *In: Parallel Programming with OpenACC.*, Elsevier (2017), pp. 215–240.
- [9] McGreevy, R., et al., Advances in the molecular dynamics flexible fitting method for cryo-EM modeling. *Methods.*, 100 (2016), pp. 50–60.
- [10] Bernardi, R.C., M. C.R. Melo, and K. Schulten, Enhanced sampling techniques in molecular dynamics simulations of biological systems. *Biochem. Biophys. Acta–Gen. Subj.*, 1850 (2015), pp. 872–877.

Voth, Gregory

- [1] Grime, J., et al., Coarse-grained simulation reveals key features of HIV-1 capsid self-assembly. *Nat. Commun.*, 7 (2016), p. 11568.
- [2] Grime, J., and G.A. Voth, Highly scalable and memory efficient ultra-coarse-grained molecular dynamics simulations. *J. Chem. Theory. Comput.*, 10 (2014), pp. 423–431.
- [3] Pak, A., et al., Immature HIV-1 lattice assembly dynamics are regulated by scaffolding from nucleic acid and the plasma membrane. *Proc. Natl. Acad. Sci. U.S.A.*, 114:47 (2017), pp. E10056–E10065.

Warnow, Tandy

- [1] Mirarab, S., et al., PASTA: Ultra-Large Multiple Sequence Alignment for Nucleotide and Amino-Acid Sequences. *Journal of Computational Biology*, 22:5 (2014), pp. 377–386.
- [2] Nguyen, N., et al., Ultra-large alignments using phylogeny aware profiles. *Genome Biology* 16:124 (2015), DOI:10.1186/s13059-015-0688-z.
- [3] Redelings, B.D., and M.A. Suchard, Joint Bayesian estimation of alignment and phylogeny. *Systematic Biology*, 54:3 (2005), pp. 401–418.
- [4] Nute, M., and T. Warnow, Scaling statistical multiple sequence alignment to large datasets. *BMC Genomics* 17: Supplement 10 (2016), DOI:10.1186/s12864-016-3101-8.
- [5] Nute, M., E. Saleh, and T. Warnow, Benchmarking statistical alignment. *bioRxiv* (2018), DOI:10.1101/304659.

- [6] Mirarab, S., et al., PASTA github site (<https://github.com/smirarab/pasta>), accessed May 13, 2018.

GEOSCIENCE

Corcoran, Jennifer

- [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 Sensing*, 52:2 (2015), pp. 198–217.
- [2] Huberty, B., Vegetation Surface Canopy Digital Elevation Models. *DOI Remote Sensing Activities 2017*, U.S. Department of Interior.
- [3] Huberty, B., et al., Wetland Identification and Change Detection using a Multi-sensor, Multi-frequency Remote Sensing Approach. *Earth Observation Summit* (Montreal, Canada, 2017).

Di Girolamo, Larry

- [1] <https://science.nasa.gov/earth-science/earth-science-data/data-processing-levels-for-eosdis-data-products>
- [2] Nuclear Research Council, *Earth Science and Applications from Space: National Imperatives for the Next Decade and Beyond*. The National Academies Press, Washington D.C. (2007).
- [3] Wolfe, R., and H. Ramapriyan, Scaling the pipe: NASA EOS Terra data systems at 10. *IEEE International. Geoscience and Remote Sensing Symposium (IGARSS), 2010 IEEE International* (Honolulu, Hawaii), pp.1300–1303.
- [4] Di Girolamo, L., et al., The Terra data fusion project: an update. *American Geophysical Union 2017 Fall Meeting* (New Orleans, La., December 10–15, 2017).

Guan, Kaiyu

- [1] Herman, J., and W. Usher, SALib: an open-source Python library for sensitivity analysis. *The Journal of Open Source Software*, 2 (2017).
- [2] Hurrell, J.W., et al., The Community Earth System Model: A Framework for Collaborative Research. *Bulletin of the American Meteorological Society*, 94 (2013), pp. 1339–1360.
- [3] Laloy, E., and J.A. Vrugt, High-dimensional posterior exploration of hydrologic models using multiple-try DREAM (ZS) and high-performance computing. *Water Resources Research*, 48 (2012).
- [4] Lawrence, D.M., Parameterization improvements and functional and structural advances in version 4 of the Community Land Model. *Journal of Advances in Modeling Earth Systems*, 3 (2011), p. 03001.
- [5] Lobell, D.B., et al., The critical role of extreme heat for maize production in the United States. *Nature Clim. Change*, 3 (2013), pp. 497–501.
- [6] Long, S.P., et al., Food for thought: lower-than-expected crop yield stimulation with rising CO₂ concentrations. *Science*, 312 (2006), pp. 1918–1921.

- [7] Rosenzweig, C., et al., Assessing agricultural risks of climate change in the 21st century in a global gridded crop model intercomparison. *Proceedings of the National Academy of Sciences*, 111 (2014), pp. 3268–3273.
- [8] Schlenker, W., and M.J. Roberts, Nonlinear temperature effects indicate severe damages to U.S. crop yields under climate change. *Proceedings of the National Academy of Sciences*, 106 (2009), pp. 15594–15598.
- [9] Vrugt, J.A., Markov chain Monte Carlo Simulation Using the DREAM Software Package: Theory, Concepts, and MATLAB Implementation. *Environ. Model. Softw.*, (2016), pp. 273–316.

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] Hu, J., L. Liu, A. Hermsillo, 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.
- [3] 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.
- [4] Hu, J., et al., Modification of the Western Gondwana craton by plume-lithosphere interaction. *Nature Geoscience*, (2018), DOI:10.1038/s41561-018-0064-1.
- [5] Zhou, Q., and L. Liu, A Hybrid Forward-Adjoint Data Assimilation Approach for Reconstructing the Temporal Evolution of Mantle Dynamics. *Geochem. Geophys. Geosyst.*, (2017), DOI:10.1002/2017GC007116.
- [6] Zhou, Q., L. Liu, and J. Hu, Western U.S. volcanism due to intruding oceanic mantle driven by ancient Farallon slabs. *Nature Geoscience*, (2018), DOI:10.1038/s41561-017-0035-y.

Morin, Paul

- [1] Noh, M.J., and I.M. Howat, Automated stereo-photogrammetric DEM generation at high latitudes: Surface Extraction from TIN-Based Search Minimization (SETSM) validation and demonstration over glaciated regions. *GIScience and Remote Sensing*, (2015), DOI:10.1080/15481603.2015.1008621.
- [2] Noh, M.J., and I.M. Howat, Surface Extraction from TIN based Search-space Minimization (SETSM) algorithm. *International Society for Photogrammetry and Remote Sensing (ISPRS) Journal of Photogrammetry and Remote Sensing*, 129 (2017), pp. 55–76.

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] Pekurovsky, D., P3DFFT: A framework for parallel computations of Fourier transforms in three dimensions. *SIAM J. Sci. Comput.*, 34:4 (2012), pp. C192–C209.
- [3] Fevrier, 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., 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.

Srifer, Ryan (1)

- [1] Emanuel, K., et al., Environmental Control of Tropical Cyclone Intensity. *J. Atmospheric Sci.*, 61:7, (2004), pp. 843–858.
- [2] Emanuel, K., Contribution of tropical cyclones to meridional heat transport by the oceans. *J. Geophys. Res. Atmospheres*, 106:D14 (2001), pp. 14771–14781.
- [3] Srifer, R. L., and M. Huber, Observational evidence for an ocean heat pump induced by tropical cyclones. *Nature*, 447:7144 (2007), pp. 577–580.
- [4] 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.

- [5] 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.
- [6] 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.

Striver, Ryan (2)

- [1] Gent, P., et al., The Community Climate System Model Version 4. *Journal of Climate*, 24 (2011), pp. 4973–4991.
- [2] 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.
- [3] Neale, R.B., et al., The Mean Climate of the Community Atmosphere Model (CAM4) in Forced SST and Fully Coupled Experiments. *Journal of Climate*, 26:14 (2013), pp. 5150–5168.

Striver, Ryan (3)

- [1] Shields, C.A., and J.T. Kiehl, Atmospheric River Landfall-Latitude Changes in Future Climate Simulations. *Geophys. Res. Lett.*, 43 (2016), DOI:10.1002/2016GL070470.
- [2] Shields, C.A., and J.T. Kiehl, Simulating the Pineapple Express in the half degree Community Climate System Model, CCSM4. *Geophys. Res. Lett.*, 43 (2016), DOI:10.1002/2016GL069476.
- [3] Shields, C.A., J.T. Kiehl, and G.A. Meehl, Future changes in regional precipitation simulated by a half-degree coupled climate model: Sensitivity to horizontal resolution. *J. Adv. Model. Earth Syst.*, 8 (2016), pp. 863–884.
- [4] Small, R.J., et al., A new synoptic scale resolving global climate simulation using the Community Earth System Model. *J. Adv. Modeling Earth Systems*, 6 (2014), DOI:10.1002/2014MS000363.

Vidale, John (Jordan, Thomas is former PI)

- [1] Dieterich, J.H., and K. Richards-Dinger, Earthquake recurrence in simulated fault systems. *Pure Appl. Geophysics*, 167 (2010), pp. 1087–1104.
- [2] Roten, D., et al., High-frequency nonlinear earthquake simulations on petascale heterogeneous supercomputers. *2016 ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis* (Salt Lake City, Utah).

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*, eds. J.D. Spengler and R. Wilson, Harvard University Physics Dept. (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), p. D9.
- [4] Skamarock, W., et al., A Description of the Advanced Research WRF Version 3. *NCAR Tech. Note NCAR/TN-475+STR* (2008).
- [5] *Preparation of Emissions Inventories for the Version 6.3, 2011 Emissions Modeling Platform Technical Support Document*. Environmental Protection Agency (2016).

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] Tilmes, S., et al., Description and evaluation of tropospheric chemistry and aerosols in the Community Earth System Model (CESM1. 2). *Geoscientific Model Development*, 8 (2015), pp.1395–1426.

SOCIAL SCIENCE, ECONOMICS & HUMANITIES

Cai, Yongyang

- [1] Cai, Y., et al., Climate Policy under Cooperation and Competition between Regions with Spatial Heat Transport. National Bureau of Economic Research Working Paper 24473 (2018).
- [2] Cai, Y., K.L. Judd, and T.S. Lontzek, The social cost of carbon with economic and climate risks. *Journal of Political Economy*, arXiv:1504.06909 (2015).
- [3] 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.
- [4] Cai, Y., et al., Modeling Uncertainty in Large Scale Multi Sectoral Land Use Problems. Working paper (2018).
- [5] Yeltekin, S., Y. Cai, and K.L. Judd, Computing equilibria of dynamic games. *Operations Research*, 65:2 (2017), pp. 337–356.

GRAD 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 PFLOTRAN: Evaluating the Parallel Performance of PFLOTRAN. *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).
- [2] Smetacek, V., and A. Zingone, Green and golden seaweed tides on the rise. *Nature*, 504 (2013), pp. 84–8.
- [3] Beron-Vera, F., et al., Dissipative inertial transport patterns near coherent Lagrangian eddies in the ocean. *Chaos*, 25 (2015), DOI:10.1063/1.4928693.
- [4] Chassignet, E., L. Smith, and G. Halliwell, North Atlantic Simulations with the Hybrid Coordinate Ocean Model (HYCOM): Impact of the Vertical Coordinate Choice, Reference Pressure, and Thermobaricity. *J Phys Oceanogr*, 33 (2003), pp. 2504–2526.
- [5] Fennel K., et al., Nitrogen cycling in the Middle Atlantic Bight: Results from a three-dimensional model and implications for the North Atlantic nitrogen budget. *Global Biogeochem Cycles*, 20 (2006), p. GB3007.
- [6] Gower J., E. Young, S. King, Satellite images suggest a new Sargassum source region in 2011. *Remote Sens Lett*, 4 (2013), pp. 764–773.

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 Supplement*, 204:2 (2013), p. 17
- [2] Salem, M., G.L. Bryan, and L. Corlies, Role of cosmic rays in the circumgalactic medium. *Monthly Notices of the Royal Astronomical Society*, 456:1 (2016), pp. 582–601
- [3] Bryan, G.L., et al., ENZO: An Adaptive Mesh Refinement Code for Astrophysics. *The Astrophysical Journal Supplement*, 211:2 (2014), p. 19
- [4] Boulares, A., and D. Cox, Galactic hydrostatic equilibrium with magnetic tension and cosmic ray diffusion. *The Astrophysical Journal*, 365:1 (1990), pp. 544–558

Clement, Matthew

- [1] Haisch, K., E. Lada, and C. Lada, Disk Frequencies and Lifetimes in Young Clusters. *ApJ*, 553:2 (2001), DOI:10.1086/320685.

- [2] Gomes, R., et al., Origin of the cataclysmic Late Heavy Bombardment period of the terrestrial planets. *Nature*, 435:7041 (2005), DOI:10.1038/nature03676.
- [3] Levison, H., et al., Origin of the structure of the Kuiper belt during a dynamical instability in the orbits of Uranus and Neptune. *Icarus*, 196:1 (2008), DOI:10.1016/j.icarus.2007.11.035.
- [4] Kaib, N., and J. Chambers, The fragility of the terrestrial planets during a giant-planet instability. *MNRAS*, 455:4 (2016), DOI:10.1093/mnras/stv2554.
- [5] Chambers, J., Late-stage planetary accretion including hit-and run collisions and fragmentation. *Icarus*, 224:1 (2013), DOI:10.1016/j.icarus.2013.02.015.
- [6] Grimm, S., and J. Stadel, The GENGA Code: Gravitational Encounters in N-body Simulations with GPU Acceleration. *ApJ*, 796:1 (2014), DOI:10.1088/0004-637X/796/1/23.
- [7] Chambers, J., Making More Terrestrial Planets. *Icarus*, 152:1 (2001), DOI:10.1006/icar.2001.6639.
- [8] Raymond, S., et al., Building the terrestrial planets: Constrained accretion in the inner Solar System. *Icarus*, 203:2 (2009), DOI:10.1016/j.icarus.2009.05.016.
- [9] Dauphas, N., and A. Pourmand, Hf-W-Th evidence for rapid growth of Mars and its staus as a planetary embryo. *Nature*, 473:7348 (2011), DOI:10.1038/nature10077.
- [10] Walsh, K., et al., A low mass for Mars from Jupiter's early gas-driven migration. *Nature*, 475:7355 (2011), DOI:10.1038/nature10201.
- [11] Chambers, J., A hybrid symplectic integrator that permits close encounters between massive bodies. *MNRAS*, 304:4 (1999), DOI:10.1046/j.1365-8711.1999.02379.x.

Cook, Salme

- [1] Ross, D.A., *Introduction to Oceanography*. Harper Collins College Publishers, New York, New York (1995).
- [2] Haidvogel, D.B., et al., Ocean forecasting in terrain-following coordinates: formulation and skill assessment of the Regional Ocean Modeling System. *Journal of Computational Physics*, 227 (2008), pp. 3595–3624.
- [3] Warner, J.C., et al., Development of a three-dimensional, regional, coupled wave, current, and sediment-transport model. *Computers and Geosciences*, 34 (2008), pp. 1284–1306.
- [4] Warner, J.C., et al., Development of a Coupled Ocean–Atmosphere–Wave–Sediment Transport (COAWST) Modeling System. *Ocean Modeling*, 35 (2010), pp. 230–244.

Foster, Lauren

- [1] U.S. Department of the Interior, *Colorado River Basin Water Supply and Demand Study*, 85 (2012).
- [2] Ikeda, K., et al., Simulation of seasonal snowfall over Colorado. *Atmos. Res.*, 97 (2010), pp. 462–477.
- [3] EDW Group, Elevation-dependent warming in mountain regions. *Nat. Clim. Chang.*, 5 (2015), pp. 424–430.

- [4] Beniston, M., Climatic Change in Mountain Regions. *Clim. Chang.*, 59 (2003), pp. 5–31.
- [5] Leung, L.R. and Y. Qian, Sensitivity of Precipitation and Snowpack Simulations to Model Resolution in Regions of Complex Terrain. *J. Hydrometeorol.*, 4 (2003), pp. 1025–1043.
- [6] Wood, A.W., et al., Hydrologic implications of dynamical and statistical approaches to downscaling climate model outputs. *Clim. Change*, 62 (2004), pp. 189–216.
- [7] Maxwell, R.M. and N.L. Miller, Development of a Coupled Land Surface and Groundwater Model. *J. Hydrometeorol.*, 6 (2005), pp. 233–247.
- [8] Kollet, S.J., and R.M. Maxwell, Integrated surface-groundwater flow modeling. *Adv. Water Resour.*, 29 (2006), pp. 945–958.
- [9] Ray, A., J. Barsugli, and K. Averyt, Climate change in Colorado: a synthesis to support water resources management and adaptation. *Color. Water Conserv. Board* (2008), pp. 1–58.
- Goldsmith, Zachary**
- [1] Goldsmith, Z., et al., Characterization of NiFe oxyhydroxide electrocatalysts by integrated electronic structure calculations and spectroelectrochemistry. *Proc. Natl. Acad. Sci. U.S.A.*, 114:12 (2017), pp. 3050–3055.
- [2] Chen, J.Y.C., et al., *Operando* analysis of NiFe and Fe oxyhydroxide electrocatalysts for water oxidation: Detection of Fe⁴⁺ by Mössbauer spectroscopy. *J. Am. Chem. Soc.*, 137:48 (2015), pp. 15090–15093.
- [3] Stevens, M.B., et al., Reactive Fe-sites in Ni/Fe (oxy) hydroxide are responsible for exceptional oxygen electrocatalysis activity. *J. Am. Chem. Soc.*, 139:33 (2017), pp. 11361–11364.
- [4] Hunter, B.M., et al., Trapping an Iron(VI) Water-Splitting Intermediate in Nonaqueous Media. *Joule*, 2:4 (2018), pp. 747–763.
- [5] 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), pp. 395502–395520.
- Hays, Jennifer**
- [1] Peng, H., F. Long, and C. Ding, Feature selection based on mutual information criteria of max-dependency, max-relevance, and min-redundancy. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27:8 (2005), pp. 1226–1238.
- [2] Roux, B., and S. Islam, Restrained-Ensemble Molecular Dynamics Simulations Based on Distance Histograms from Double Electron–Electron Resonance Spectroscopy. *The Journal of Physical Chemistry B*, 117:17 (2013), pp. 4733–4739.
- 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] 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.
- [3] Hoogerbrugge, P.J., and J.M.V.A. Koelman, Simulating Microscopic Hydrodynamic Phenomena with Dissipative Particle Dynamics. *Europhys. Lett.*, 19:3 (1992), pp. 155–160.
- [4] Groot, R.D., and P.B. Warren, Dissipative particle dynamics: Bridging the gap between atomistic and mesoscopic simulation. *J. Chem. Phys.*, 107:11 (1997), pp. 4423–4435.
- [5] Anderson, J.A., C.D. Lorenz, and A. Travasset, General purpose molecular dynamics simulations fully implemented on graphics processing units. *J. Comput. Phys.*, 227 (2008), pp. 5342–5359.
- [6] Glaser, J., et al., Strong scaling of general-purpose molecular dynamics simulations on GPUs. *Comput. Phys. Commun.*, 192 (2015), pp. 97–107.
- [7] Phillips, C.L., J.A. Anderson, and S.C. Glotzer, Pseudo-random number generation for Brownian Dynamics and Dissipative Particle Dynamics simulations on GPU devices. *J. Comput. Phys.*, 230 (2011), pp. 7191–7201.
- Kruse, Ethan**
- [1] Thompson, S.E., et al., Planetary Candidates Observed by Kepler. VIII. A Fully Automated Catalog With Measured Completeness and Reliability Based on Data Release 25. *Astrophysical Journal Supplement*, 235:2 (2018), pp. 38–49.
- [2] Fulton, B.J., et al., The California-Kepler Survey. III. A Gap in the Radius Distribution of Small Planets. *Astronomical Journal*, 154 (2017), p.109.
- [3] Luger, R., et al., EVEREST: Pixel Level Decorrelation of K2 Light Curves. *Astronomical Journal*, 152 (2016), p. 100.
- [4] Carter, J.A., and Agol, E. The Quasiperiodic Automated Transit Search Algorithm. *Astrophysical Journal*, 765 (2013), p. 132.
- Mansbach, Rachael**
- [1] Wall, B., et al., Supramolecular polymorphism: Tunable electronic interactions within π -conjugated peptide nanostructures dictated by primary amino acid sequence. *Langmuir*, 30:20 (2014), pp. 5946–5956.
- [2] Anderson, J., C. Lorenz, and A. Travasset, General purpose molecular dynamics simulations fully implemented on graphics processing units. *Journal of Computational Physics*, 227:10 (2008), pp. 5342–5359.
- [3] Glaser, J., et al., Strong scaling of general-purpose molecular dynamic simulations on GPUs. *Computer Physics Communications*, 192 (2015), pp. 97–107.
- Reames, Larissa**
- [1] Skamarock, W., and J. Klemp, A time-split nonhydrostatic atmospheric model for weather research and forecasting applications. *J. Comput. Phys.*, 227:7 (2008), pp. 3465–3485.
- Richers, Sherwood**
- [1] Woosley, S. E., et al., The r-process and neutrino-heated supernova ejecta. *The Astrophysical Journal*, 443, DOI:10.1086/174638.
- [2] Fröhlich, C., et al., Neutrino-Induced Nucleosynthesis of A>64 Nuclei: The vp Process. *Physical Review Letters*, 96:14 (2006), DOI:10.1103/PhysRevLett.96.142502.
- [3] Lattimer, J. M., and D.N. Schramm, Black-hole-neutron-star collisions. *The Astrophysical Journal*, 192 (1974), DOI:10.1086/181612.
- [4] Bethe, H. and J.R. Wilson, Revival of a stalled supernova shock by neutrino heating. *The Astrophysical Journal*, 295 (1985), DOI:10.1086/163343.
- [5] Metzger, B. D., and R. Fernández, Red or blue? A potential kilonova imprint of the delay until black hole formation following a neutron star merger. *Monthly Notices of the Royal Astronomical Society*, 441:4 (2014), DOI:10.1093/mnras/stu802.
- [6] Richers, S., et al., Monte Carlo Neutrino Transport through Remnant Disks from Neutron Star Mergers. *The Astrophysical Journal*, 813:1 (2015), DOI:10.1088/0004-637X/813/1/38.
- [7] O'Connor, E., An Open-source Neutrino Radiation Hydrodynamics Code for Core-collapse Supernovae. *The Astrophysical Journal Supplement Series*, 219:2 (2015), DOI:10.1088/0067-0049/219/2/24.
- Seyler, Sean**
- [1] Kaiser, A., A. Snezhko, and I.S. Aranson, Flocking ferromagnetic colloids. *Sci. Adv.*, 3:2 (2017), DOI:10.1126/sciadv.1601469.
- [2] Woodhouse, F.G., and R.E. Goldstein, Spontaneous circulation of confined active suspensions. *Phys. Rev. Lett.*, 109:16 (2012), DOI:10.1103/PhysRevLett.109.168105.
- [3] Ben-Isaac, E., et al., Modeling the dynamics of a tracer particle in an elastic active gel. *Phys. Rev. E*, 92:1 (2015), DOI:10.1103/PhysRevE.92.012716.
- [4] Höfling, F., and T. Franosch, Anomalous transport in the crowded world of biological cells. *Rep. Prog. Phys.*, 76:4 (2013), DOI:10.1088/0034-4885/76/4/046602.
- [5] Kokot, G., et al., Active turbulence in a gas of self-assembled spinners. *Proc. Natl. Acad. Sci.*, (2017), DOI:10.1073/pnas.1710188114.
- [6] Vandebroek, H., and C. Vanderzande, The effect of active fluctuations on the dynamics of particles, motors and DNA-hairpins. *Soft Matter*, 13:11 (2017), DOI:10.1039/c6sm02568d.
- [7] Fodor, É., et al., Activity-driven fluctuations in living cells. *EPL*, 110:4 (2015), DOI:10.1209/0295-5075/110/48005.
- [8] Wijesinghe, H.S., and N.G. Hadjiconstantinou, Discussion of Hybrid Atomistic-Continuum Methods for Multiscale Hydrodynamics. *Int. J. Multiscale Com.*, 2:2 (2004), DOI:10.1615/IntJMultCompEng.v2.i2.20.
- [9] Mohamed, K.M., and A.A. Mohamad, A review of the development of hybrid atomistic-continuum methods for dense fluids. *Nanofluidics*, 8:3 (2010), DOI:10.1007/s10404-009-0529-z.
- [10] Donev, A., J. Bell, A. Garcia, and B. Alder, A Hybrid Particle-Continuum Method for Hydrodynamics of Complex Fluids. *Multiscale Model. Simul.*, 8:3 (2010), DOI:10.1137/090774501.
- [11] Landau, L.D., and E.M. Lifshitz, *Fluid Mechanics*, Vol. 6 (1959), pp. 523–529.
- [12] Fox, R.F., and G.E. Uhlenbeck, Contributions to Nonequilibrium Thermodynamics. II. Fluctuation Theory for the Boltzmann Equation. *The Physics of Fluids*, 13:12 (1970), pp. 2881–2890.
- [13] Bell, J.B., A.L. Garcia, and S.A. Williams, Numerical methods for the stochastic Landau–Lifshitz Navier–Stokes equations. *Phys. Rev. E*, 76:1 (2007), DOI:10.1103/PhysRevE.76.016708.
- [14] Grad, H., On the kinetic theory of rarefied gases. *Comm. Pure Appl. Math.*, 2 (1949), pp. 331–407.
- [15] Seyler, C.E., and M.R. Martin, Relaxation model for extended magnetohydrodynamics: Comparison to magnetohydrodynamics for dense Z-pinches. *Phys. Plasmas*, 18:1 (2011), DOI:10.1063/1.3543799.
- [16] Riedel, C., et al., The heat released during catalytic turnover enhances the diffusion of an enzyme. *Nature*, 517:7533 (1999), DOI:10.1038/nature14043.
- [17] Ikoma, A., et al., Fluctuating hydrodynamics for a rarefied gas based on extended thermodynamics. *Phys. Lett. A*, 375:27 (2011), DOI:10.1016/j.physleta.2011.05.057.
- [18] Torrilhon, M., *Regularization of Grad's 13-moment equations in kinetic gas theory*. Technical report: RWTH Aachen University, Germany, RTO-EN-AVT-194 (2011), DOI:10.14339/RTO-EN-AVT-194.
- [19] De Fabritiis, G., M. Serrano, R. Delgado-Buscalioni, and P.V. Coveney, Fluctuating hydrodynamic modeling of fluids at the nanoscale. *Phys. Rev. E*, 75:2 (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., et al., Centrifugal precipitation transport in tornadic supercells: An algorithm consistent for use with bulk microphysics schemes. *14th Intl. Conf. on Clouds and Precipitation* (Bologna, Italy, 2004), pp.1654–1656.
- [3] Thompson, R.L., et al., Close proximity soundings within supercell environments obtained from the rapid update cycle. *Wea. Forecasting*, 18 (2003), pp. 1243–1261.

Teich, Erin

- [1] Damasceno, P.F., M. Engel, and S.C. Glotzer, Predictive self-assembly of polyhedra into complex structures. *Science*, 337:6093 (2012), pp. 453–457.
- [2] Wuttig, M., and N. Yamada, Phase-change materials for rewriteable data storage. *Nature Materials*, 6:11 (2007), pp. 824–832.
- [3] Berthier, L., and G. Biroli, Theoretical perspective on the glass transition and amorphous materials. *Rev. Mod. Phys.*, 83 (2011), pp. 587–645.
- [4] Available at <https://glotzerlab.engin.umich.edu/hoomd-blue>.
- [5] Anderson, J.A., M.E. Irrgang, and S.C. Glotzer, Scalable Metropolis Monte Carlo for simulation of hard shapes. *Computer Physics Communications*, 204 (2016), pp. 21–30.
- [6] Chen, E.R., et al., Complexity in surfaces of densest packings for families of polyhedra. *Physical Review X*, 4:1 (2014), pp. 011024.
- [7] Klotsa, D., et al., Intermediate crystalline phases of colloids in shape space. In preparation (2018).
- [8] Kob, W., et al., Dynamical Heterogeneities in a Supercooled Lennard-Jones Liquid. *Physical Review Letters*, 79:15 (1997), pp. 2827–2830.
- [9] Glotzer, S.C., V.N. Novikov, and T.B. Schroder, Time-dependent, four-point density correlation function description of dynamical heterogeneity and decoupling in supercooled liquids. *The Journal of Chemical Physics*, 112:2 (2000), pp. 509–512.

Totorica, Sam

- [1] Zweibel, E.G., and M. Yamada, Magnetic Reconnection in Astrophysical and Laboratory Plasmas. *Annu. Rev. Astron. Astrophys.*, 47 (2009), pp. 291–332.
- [2] Fonseca, R.A., et al., Exploiting multi-scale parallelism for large scale numerical modelling of laser wakefield accelerators. *Plasma Phys. Controlled Fusion*, 55 (2013), p. 124011.
- [3] Abel, T., O. Hahn, and R. Kaehler, Tracing the dark matter sheet in phase space. *Mon. Not. R. Astron. Soc.*, 427 (2012), pp. 61–76.

INDEX

A

Acharya, Sumanta 162
 Agee, Elizabeth 270
 Aksimentiev, Aleksei 202, 204
 Aluru, Narayana 160, 164, 206

B

Balajewicz, Maciej 106
 Balsara, Dinshaw 56
 Bauer, Gregory 214
 Bernholc, Jerry 168
 Bodony, Daniel 174
 Brooks, Maureen 272
 Butsky, Iryna 274

C

Caetano-Anolles, Gustavo 208
 Cai, Yongyang 266
 Campbell, Roy H 180
 Carnevale, Vincenzo 210
 Carrasco Kind, Matias 58
 Ceperley, David 142
 Cheatham, Thomas 212
 Cheng, Chi-Hing 216
 Chew, Huck Beng 170
 Cho, Wendy 264
 Clark, Bryan 148
 Clement, Matthew 276
 Cook, Salme 278
 Corcoran, Jennifer 90
 Cox, Donna 182

D

Di Girolamo, Larry 86
 Di Matteo, Tiziana 30
 Dill, Ken 218
 Draayer, Jerry 136
 Duan, Lian 150

E

El-Kebir, Mohammed 220

F

Faucher-Giguere, Claude-Andre 68
 Feinberg, Evan 280
 Ferguson, Andrew 222
 Fischer, Paul 140
 Foster, Lauren 282

G

Garcia, Marcelo 158
 Gazzola, Mattia 224
 Gnedin, Nickolay 64
 Goldsmith, Zachary 284
 Gropp, Bill 184, 226
 Guan, Kaiyu 82
 Gupta, Indranil 186

H

Haas, Roland 26
 Hawley, John 52
 Hays, Jennifer 286
 Hime, Paul 316
 Hirata, So 228
 Hooberman, Benjamin 166
 Hopkins, Philip 66
 Howard, Michael 288
 Huerta Escudero, Eliu 22, 28
 Huff, Kathryn 104

I

Ilie, Raluca 60
 Ismail-Beigi, Sohrab 154
 Iyer, Ravishankar 188

J

Jain, Prashant 152
 Jasiuk, Iwona 138
 Jha, Shantenu 190
 Johnsen, Eric 156
 Johnson, Harley 146
 Jones, Alexandra 290
 Jordan, Thomas 98

K

Kirby, Andrew 292
 Klimeck, Gerhard 144
 Koric, Seid 172, 192
 Kruse, Ethan 294
 Kurchin, Rachel 296

L

Lasher-Trapp, Sonia 84
 Leburton, Jean-Pierre 132
 Levin, Deborah 42, 124
 Liu, Lijun 96
 Liu, Yi-Hsin 38
 Luthey-Schulten, Zaida 230

M

Mackenzie, Paul 130
 Mainzer, Liudmila 232
 Makri, Nancy 234
 Mansbach, Rachael 298
 Masud, Arif 236
 Mori, Warren 112
 Morin, Paul 88

N

Nagi, Rakesh 194
 Neubauer, Mark 128
 Noble, Scott 32
 Norman, Michael 24

O

O'Shea, Brian 50
 Olson, Luke 196
 Orf, Leigh G. 72

P

Payne, William 300
 Perilla, Juan 238
 Pogorelov, Nikolai 54

Q

Quinn, Thomas 44

R

Raman, Venkat 134
 Rani, Sarma 94
 Reames, Larissa 302
 Richers, Sherwood 304
 Riedl, Caroline 126
 Roux, Benoit 240
 Roytershteyn, Vadim 40

S

Schive, Hsi-Yu 46
 Schleife, Andre 108, 110
 Seyler, Sean 306
 Shapiro, Stuart 20
 Shukla, Diwakar 242, 244
 Sirignano, Justin 187
 Soltesz, Ivan 246
 Srinivasan, Ashok 248
 Sriver, Ryan 76, 78, 80
 Stenz, Ronald 308
 Sutton, Brad 250

T

Tagkopoulos, Ilias 252
 Taha, Ahmed 116, 122
 Tajkhorshid, Emad 254, 256
 Tchekhovskoy, Alexander 34
 Teich, Erin 310
 Teukolsky, Saul A. 48
 Thomas, Brian 118
 Tinoco Lopez, Rafael 120
 Toth, Gabor 62
 Totorica, Sam 312
 Trapp, Robert 100
 Turk, Matthew 49

V

Vidale, John 98
 Voth, Gregory 258

W

Wagner, Lucas 176
 Warnow, Tandy 260
 West, Matthew 74, 198
 Woodward, Paul R. 36
 Wuebbles, Donald J. 92

Y

Yeung, Pui-Kuen 114

Z

Zemskova, Varvara 314





I ILLINOIS
NCSA | National Center for
Supercomputing Applications

*Blue Waters is supported by the
National Science Foundation*

