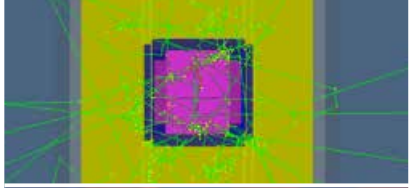

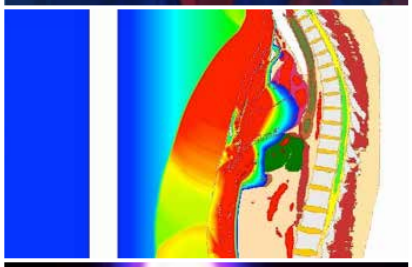

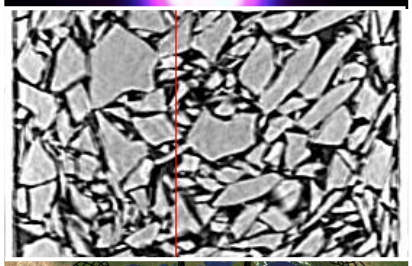

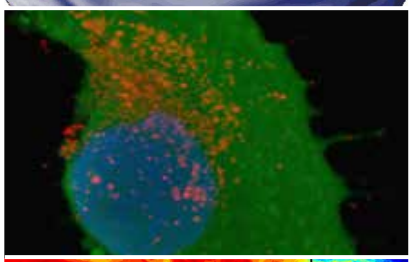
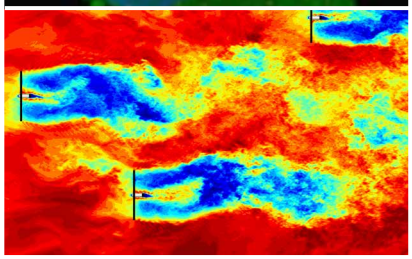
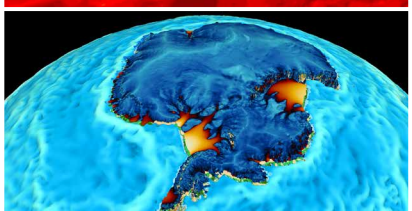

SANDIA MISSION

COMPUTING

2015 HPC ANNUAL REPORT

Sandia National Laboratories

SNLSimMagic©

AUGMENTED REALITY APP

Scanning this code with an iPhone or iPad will provide access to SNLSimMagic; an augmented reality iOS application that can be downloaded to the device. You can also download directly onto your device from the Apple App store. Readers with the application can use their mobile devices to scan images in this document that show the AR icon, and an associated movie clip will be played on their device. SNLSimMagic was developed at Sandia National Laboratories.

Laboratory Directed Research and Development

The Laboratory Directed Research and Development (LDRD) program, authorized by U.S. Congress in 1991, enables DOE laboratories to devote a small portion of their research funding to high-risk and potentially high-payoff research. Many successes at Sandia can be traced back to investments in LDRD — these investments have a history of enabling significant payoffs for long-running DOE and NNSA missions and for providing anticipatory new technologies that ultimately become critical to future missions. Look for the LDRD logo to learn more about projects with origins in LDRD.

Credits

Editor: Barbara Jennings

Contributing Writers: John Noe, Andrea Penner, Laura Sowko

Design: Stacey Long, Michael Vittitow

Video: Regina Valenzuela

Photography: Norman Johnson

SNLSimMagic: Gabriel Arrillaga, Barbara Jennings, Chadwick Johnson, Trace Norris, Matt Smith, Charlie Snider

TABLE OF CONTENTS

Providing Essential Infrastructure and Services

Center 9300 is very proud to provide a wide range of computing and networking resources ranging from desktop support to high performance computing that enable Sandia's national security mission needs. Our goal is to provide these services in the most efficient, secure, user friendly, and agile means necessary. High performance computing is one example of how we are delivering on this goal. Thanks to investments from NNSA's Advanced Simulation and Computing program, coupled with Sandia's Institutional Computing program, a wide variety of computing resources are available and used to support national security needs.

This year we are expanding our high performance computing infrastructure with the procurement of institutional resources for information computing or analytics. Big data analytics has been the buzzword in the computing industry for over a decade now and industries in finance, food and drug, business, and security organizations have been leading the charge. Much like how scientific computing has changed the way science and engineering are performed, data analytics is changing just about every aspect of our lives and more especially our ability to contribute to the national security mission of Sandia.

This third edition of our HPC Annual Report continues our commitment to communicate the details and impact of Sandia's large-scale computing resources that support missions work across Sandia. The Laboratories' and NNSA's strategic investments in high performance computing demonstrate our commitment to keeping mission-related research on the cutting edge of science, at pace with technology advances, and at a scale that meets the toughest demands of the nation we so proudly serve.

DAVID WHITE
Director, Computing and Network Services

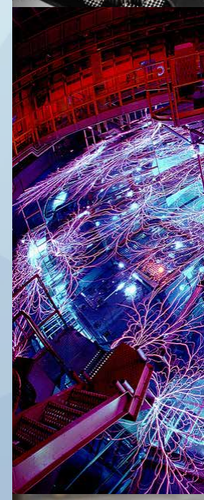

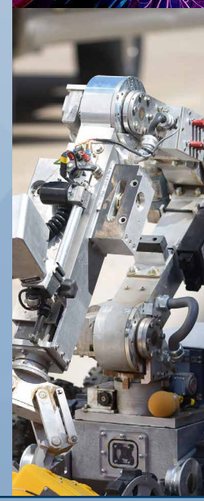

Sandia Creates Solutions for Security Challenges

It is gratifying to see the breadth of Sandia's missions being enabled by our high performance computing capabilities. The projects described in this report illustrate how central computational simulation has become to all fields of science and engineering. Computation, when coupled appropriately with physical experiments, enables greater depth of understanding, enhanced agility, and reduced cost. Sophisticated computational modeling relies on a broad ecosystem that includes mathematics, algorithms, advanced application codes, and deep scientific and engineering expertise, and Sandia sustains all of these capabilities. But a critical piece of the ecosystem is the availability of high performance computing platforms. Thanks to the investments made by NNSA's Advanced Simulation and Computing program and Sandia's Institutional Computing Program, we have these machines and the specialized expertise to run them and to support users.

Sandia is committed to sustaining and continuing to grow this critical capability that underpins so many of our national security mission areas. This capability enables new opportunities for discovery and innovation and new opportunities to contribute to the solution of national challenges. It also allows us to contribute to the broader national need for advanced computing. In recent Congressional testimony Norm Augustine, former chairman of Lockheed Martin, asserted that for the United States to compete successfully with other countries we need to make greater use of computing. As a Federally Funded Research and Development Center, and as a national laboratory, Sandia will continue to push the boundaries of science and technology in our commitment to exceptional service in the national interest. As part of this commitment, we will continue to provide national leadership in the development and application of high performance computing.

BRUCE HENDRICKSON
Director, Center for Computing Research

Sandia Unites Mission Capabilities

In parallel with Sandia National Laboratories having two major locations (NM and CA), along with a number of smaller facilities across the nation, so too is the distribution of scientific, engineering and computing resources. As a part of Sandia's Institutional Computing Program, CA site-based Sandia computer scientists and engineers have been providing mission and research staff with local CA resident expertise on computing options while also focusing on two growing high performance computing research problems. The first is how to increase system resilience to failure, as machines grow larger, more complex and heterogeneous. The second is how to ensure that computer hardware and configurations are optimized for specialized data analytical mission needs within the overall Sandia computing environment, including the HPC sub-environment.

All of these activities support the larger Sandia effort in accelerating development and integration of high performance computing into national security missions. Sandia continues to both promote national R&D objectives, including the recent Presidential Executive Order establishing the National Strategic Computing Initiative and work to ensure that the full range of computing services and capabilities are available for all mission responsibilities, from national security to energy to homeland defense.

JAMES COSTA
Senior Manager
Computational Science and Analysis

DEFENSE SYSTEMS & ASSESSMENTS

The Defense Systems and Assessments (DSA) Program Management Unit (PMU) exists to research and develop technologies to enable future force capabilities. In addition to developing these capabilities across multiple domains (air, land, maritime, space, and cyber) we exploit these innovations and transition technology-enabled capabilities to the current force through our cooperative agreements with industrial partners. Current and future national security threats remain many and varied, including those related to nuclear non-proliferation, cyber security, nuclear command and control, missile defense, reducing the vulnerability and increasing the resiliency of trusted systems, and preventing technological surprise. The DS&A PMU works strategically across these and many other threat areas with all levels of government to solve the nation's highest priority national security issues.

DAHLON CHU, *Deputy for Technology and Programs*

ENERGY & CLIMATE

Energy & Climate (EC) PMU research programs address the nation's most daunting stationary power, transportation energy, and climate science and technology challenges within the national security context—leveraging and enhancing key competencies associated with Sandia's nuclear weapons (NW) mission to amplify our contributions to broader national security.

EC research furthers Sandia engineering excellence with an emphasis on connecting deep science to engineering solutions—marshalling the efforts of world-class scientists and engineers to create an energy future that is sustainable by driving the development and deployment of energy sources that are safer, cleaner, more economical and efficient, and less dependent on scarce natural resources and by assuring a reliable and resilient energy infrastructure. Any sustainable energy future requires understanding and ensuring that the Earth's climate system supports the nation's energy systems while also mitigating the impact these energy systems have on the Earth. The PMU seeks to create this energy future—informed by a science-based understanding of the complex interdependencies between energy and climate.

A key element of EC success is how our programs fully integrate high performance computing (HPC) within our research activities. Whether it be to co-evolve biofuels and advanced engine designs; study fundamental materials/structure interaction to optimize energy storage, photovoltaics, or high-voltage power electronics; or develop advanced global climate models, laboratory experiments and fieldwork to provide a deeper understanding of physical systems that is subsequently incorporated into improved HPC modeling and simulation codes. Simulation results and analysis techniques, such as uncertainty quantification, often suggest where further empirical research can foster a still deeper understanding that will improve the next generation of the model. We highlight some of the EC PMU's current set of HPC projects in this report.

MARCEY HOOVER, *Chief Operating Officer for Energy and Climate Programs*

INTERNATIONAL, HOMELAND & NUCLEAR SECURITY

The International, Homeland, and Nuclear Security (IHNS) PMU oversees a broad portfolio of Sandia's programs in areas ranging from global nuclear security to critical asset protection. We use science and technology, innovative research, and global engagement to counter threats, reduce dangers, and respond to disasters.

The PMU draws on the skills of scientists and engineers from across Sandia. Our programs focus on protecting U.S. government installations, safeguarding nuclear weapons and materials, facilitating nonproliferation activities, securing infrastructures, countering chemical and biological dangers, and reducing the risk of terrorist threats. We conduct research in risk and threat analysis, monitoring and detection, decontamination and recovery, and situational awareness. We develop technologies for verifying arms control agreements, neutralizing dangerous materials, detecting intruders, and strengthening resiliency.

Our programs use Sandia's high performance computing resources for predictive modeling and simulation of interdependent systems, for modeling dynamic threats and forecasting adaptive behavior, and for enabling decision support and processing large cyber data streams. In this report, we highlight four advanced computation projects that illustrate the breadth of the IHNS mission space.

GARY LAUGHLIN, *Technical Deputy for International, Homeland, and Nuclear Security*

LABORATORY DIRECTED RESEARCH & DEVELOPMENT

As Sandia's sole source of discretionary research and development funds, the Laboratory Directed Research and Development (LDRD) Program functions as a catalyst for the genesis of innovative science and applied advancements in engineering and technology that serve the Department of Energy and other national security missions. In fiscal year 2015, following a competitive review process, the LDRD Program Office awarded approximately 380 new and continuing projects (totaling $145 million) to Sandia's scientists and engineers.

Research activities throughout the LDRD program are high-risk but driven by anticipated mission needs. High performance computing capabilities and computational science expertise are critical to the program's and Sandia's success. Throughout the program, HPC algorithms and expertise are being developed and deployed to a variety of scientifically challenging problems. Fourteen percent of all LDRD projects utilize HPC systems—in 2015, over 75 million processor hours were clocked by LDRD projects. The LDRD projects featured on the following pages exemplify the important role Sandia's HPC resources play in developing the technologies and capabilities that support our nuclear weapons and national security missions.

SHERYL MARTINEZ, *Program Manager for Laboratory Directed Research and Development*

NUCLEAR WEAPONS

Sandia's Nuclear Weapons Program Management Unit takes advantage of DOE's exascale computing initiative to further enhance stockpile stewardship. Sandia meets the challenges of its stockpile stewardship mission by using newly developed modeling and simulation capabilities to assess and monitor the stockpile through advanced experimental design and enhanced validation and verification.

Modeling and simulation capabilities developed by the National Nuclear Security Administration Advanced Simulation and Computing Program (ASC) provide a computational basis for assessing current and future nuclear weapons stockpile. Uses include design and qualification for weapon systems and the current stockpile systems annual assessment .

Capabilities developed by ASC include Sierra and RAMSES code suites used to simulate mechanical, electrical, and radiation environments and other software supporting verification and validation, model development, foundational algorithms for computational simulation and pre- and post-processing. These capabilities enable exploration of technical and scientific issues including materials aging, fracture, failure, friction, and response in extreme environments.

DOE's new initiative to deploy exascale computing by2023 presents opportunities to advance the engineering process based on a fully predictive modeling capability. With exascale computing, NW engineers can simulate the re-entry environment and evaluate a weapon's response in wide range of abnormal and hostile environments. Exascale computing also presents challenges. Analysts must re-examine current workflow including meshing, visualization, analytics, and uncertainty quantification. As computational capabilities evolve and increase, the workflow must become more integrated with the analysis codes themselves.

DAVID WOMBLE, *Senior Manager for Advanced Simulation and Computing*

HIGH PERFORMANCE COMPUTING

The recent Executive Order creating the National Strategic Computing Initiative (NSCI) recognizes the value of high performance computing for economic competitiveness and scientific discovery and commits to accelerate delivery of exascale computing. The HPC programs at Sandia—the NNSA ASC program and Sandia's Institutional HPC Program—are focused on ensuring that Sandia has the resources necessary to deliver computation in the national interest.

Our investments in the Institutional HPC Program have helped close the resource demand gap for Sandia's Mission Areas. Institutional HPC usage has tripled across all of Sandia's Program Management Units (PMUs). Our newest system, "Sky Bridge," released in March 2015, increased computing cycles available on the Sandia Restricted Network three-fold.

This year, with agreement of the Mission Computing Council, two large data-analytics systems were purchased for use in FY16. One of the systems, to be sited in SNL/CA, will focus on data analytics and Cloud/Emulytics© research. The other system, to be sited in SNL/NM, will handle data analytics and Cloud/Emulytics production work.

Next year, the ASC program investment will result in siting the first PetaFLOP systems at SNL. Current plans are to purchase an Institutional HPC system of similar size for use by all the SNL PMUs.

I hope you enjoy reading about the diverse projects supported on our computer systems, and that you will engage us and your Mission Computing Council representatives in meaningful dialogue to help guide priorities and future investments in Sandia's HPC resources and capabilities.

TOM KLITSNER, *Senior Manager for Computing Systems and Technology Integration*

SIMULATION OF BLAST AND BEHIND-ARMOR BLUNT TRAUMA TO LIFE-CRITICAL ORGANS IN THE HUMAN TORSO

PAUL TAYLOR, PI
CANDICE COOPER

PROCESSING HOURS
1,051,120

pataylo@sandia.gov

"This work is focused on providing the means to understand the mechanics of wound injury to the U.S. warfighter and developing a new approach with which to assess new personal armor designs for their protection."

– Paul Taylor

Researchers at Sandia National Laboratories have developed a high-fidelity virtual model of the human torso to investigate the details of life-threatening injury to the respiratory and cardiovascular systems as a result of blast exposure and behind-armor blunt trauma. This model is an extension of the Sandia virtual head-neck model developed previously to investigate the connection between blast exposure and traumatic brain injury. The Sandia human torso model possesses anatomically correct distributions of bone, cartilage, intervertebral disks, vasculature, blood, airways, lungs, heart, liver, stomach, kidneys, spleen, spinal cord, muscle, and fat/skin. The torso model is used with the Sandia wave physics code, CTH, to simulate blast loading and ballistic projectile impact to the torso, without and with protective armor, to investigate the details of injury to life-critical organs such as the lungs, airways, heart, blood vessels, and liver as a result of the intrathoracic pressure waves that are generated from a blast or impact. The intent of this work is to demonstrate the advantages of applying a modeling and simulation approach to the investigation of wound injury dynamics and to assess protective body armor for the U.S. warfighter under conditions of blast and ballistic projectile impact. A typical torso injury scenario simulation requires anywhere from 448 to 960 cpu-cores, running the calculation for 30 to 60 cpu-hours, and generating over 500 GB of raw data that is post-processed at a later time for injury investigation and/or protective armor assessment.

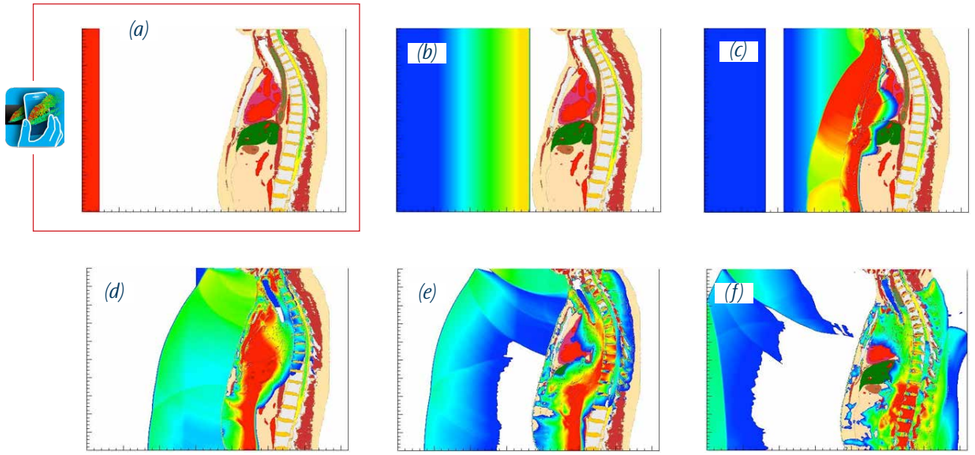

Frontal blast wave exposure of Sandia Human Torso Model. (a) Initial setup. Blast wave is generated from reservoir (denoted in red) of high pressure air on left side of image. (b) Blast wave propagating to right showing wave front just prior to contact with torso. (c) Blast wave interaction with torso has generated transmitted pressure wave within the thorax and reflected air wave propagating leftward towards blast source. (d)-(f) Time-evolution of wave transmission within thorax and reflected pressure wave in air. White space outside torso indicates regions experiencing pressures below 1 atmosphere in magnitude.

LDRD

OPTIMAL IMAGING FOR TREATY VERIFICATION

ERIK BRUBAKER, PI
NATHAN HILTON
WILLIAM JOHNSON
Sandia National Labs

MATT KUPINSKI
CHRIS MACGAHAN
University of Arizona

PROCESSING HOURS
900,806

ebrubak@sandia.gov

"It's exciting to help answer this hard but important question about how we might do verification measurements in future arms control treaties. Simulating those measurements while accounting for real-life variability makes for a big computational effort, and that's where HPC resources come in."

– Erik Brubaker

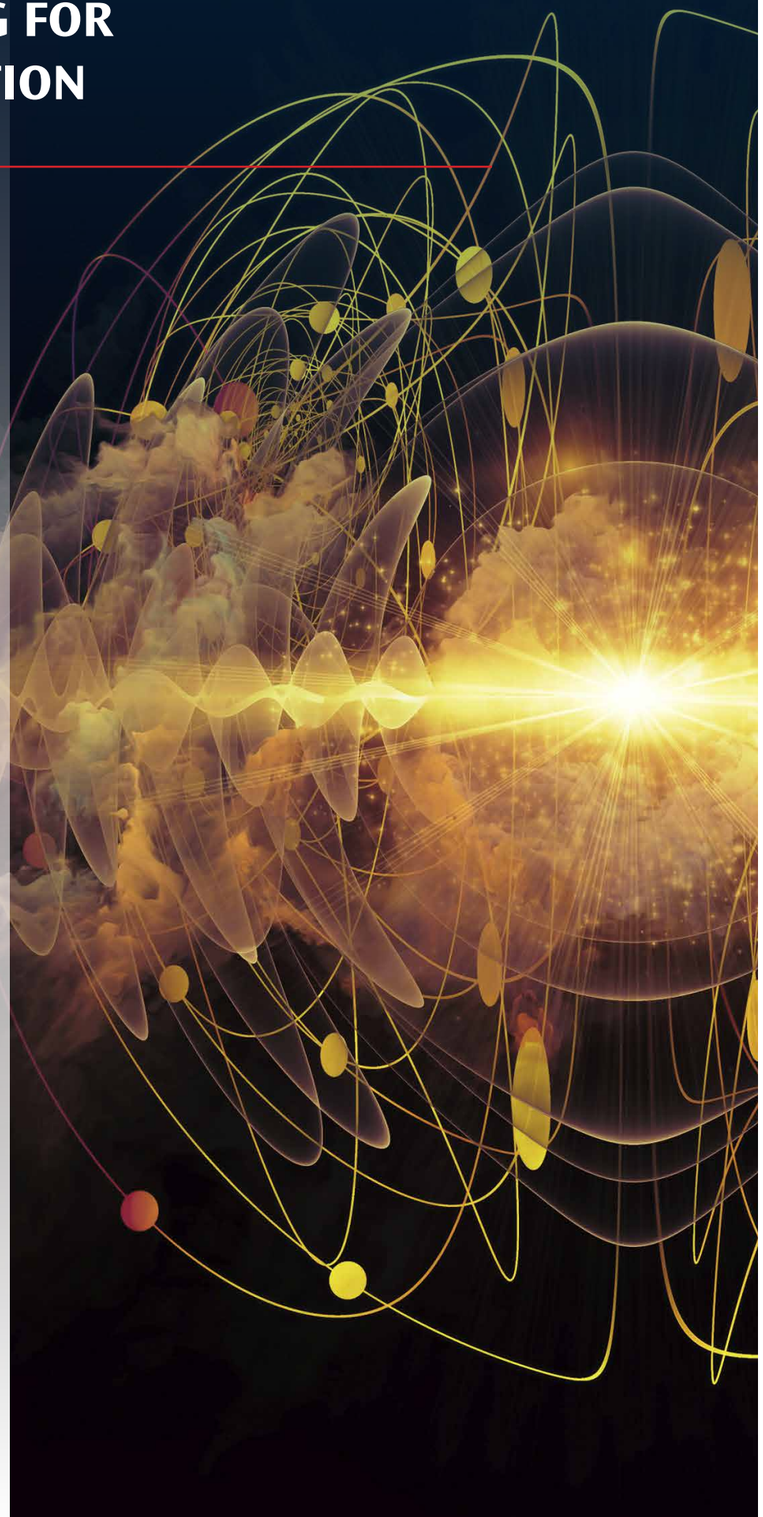

Some proposed verification methods for nuclear-arms-control treaties require the use of an information barrier. This can be either a hardware or software barrier that prevents the transmission of sensitive data to inspectors. Our focus is on the development of mathematical observer models to perform treaty verification tasks using radiation imaging systems while avoiding the need for these barriers. To test these observer models, our team simulated photon and neutron data using the GEANT4 toolkit, which uses Monte Carlo techniques to simulate radiation transport from a source object to the imaging system's detector array. We modeled plutonium and uranium inspection objects developed by Idaho National Laboratory (INL) along with a fast-neutron coded aperture detector.

Because of shielding in these objects, as few as one in a billion emitted photons may be detected by the imager. We increased the photon detection probability by imposing a minimum energy threshold of 100 keV and incorporating a linear energy bias to our source sampling. However, simulating neutron transport is slow, due in part to the need to propagate low-energy neutrons through the geometry. To overcome these obstacles, we migrated to the multithreaded build of GEANT4 and ran the simulations on Sandia's Glory high performance computing cluster.

We have also begun extensive studies on the impact of nuisance parameters, such as object orientation and location on these observer models. Initial studies have confirmed that lack of knowledge of these nuisance parameters can degrade observer performance. To avoid repeatedly simulating transport through the object, the simulation was split in two; one simulation to find the flux out of the object and another, that reads in the flux data and transports particles to the detector. These are computationally demanding simulations, and without the use of HPC resources, this work would not have been possible.

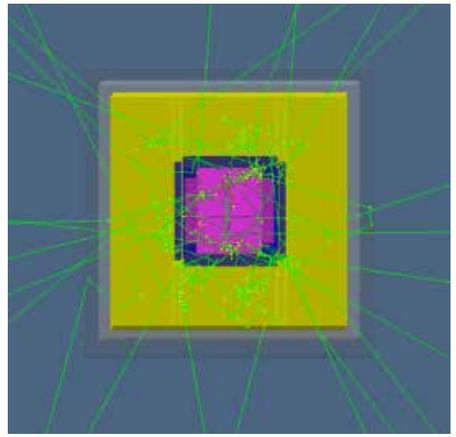

Our simulations track radiation as it is generated, scattered, and absorbed by an Idaho National Laboratory inspection object and then detected by an imaging system. The inspection object shown above contains plutonium shielded by depleted uranium inside a thick polyethylene case. Neutrons are emitted through spontaneous fission and are moderated through elastic scattering, primarily within the polyethylene.

CHARACTERIZING QUANTUM DEVICES USING MODEL SELECTION

ROBIN BLUME-KOHOUT, PI
TRAVIS SCHOLTEN

PROCESSING HOURS
659,107

rjblume@sandia.gov

"Monte Carlo simulations on Red Sky allowed us to explore the enormous space of test cases for our algorithm, and thus to figure out how model selection could enhance quantum tomography."

– Travis Scholten

Quantum information processors are a new kind of computing device with the potential to solve certain problems very quickly. The quantum bits, or "qubits," that enable such processors are being developed at Sandia and around the world. In support of this effort, Sandia's Center for Computing Research is designing methods and algorithms to characterize and debug the behavior of experimental qubits, such as trapped ions and electron spins in silicon.

This is a challenge, because qubits are hard to characterize for exactly the same reason that they are powerful. Their behavior (and misbehavior) is far richer and more complicated than that of classical bits! The most reliable models of quantum errors have infinitely many parameters, and we never have enough data to fit all of them. Until now, researchers simply threw out most of these parameters, resulting in a crude approximation that makes the problem tractable, but isn't reliable or accurate enough for Sandia's experimental work.

Our new techniques use statistical model selection to identify "significant" error parameters on the fly, applying a data-driven information criterion to choose a Hilbert space that describes the data well without overfitting. To find a reliable criterion, we needed to evaluate many candidates on an enormous corpus of test case,for which we relied on Red Sky's ability to run parallel Monte Carlo simulations. Our computations demonstrate how classic model selection methods have to be modified for quantum devices. These results will enable new techniques that identify the effective dimension of as-built qubits on the fly, and then use this information to reliably pin down error behaviors and produce successively better and more useful qubits.

This project highlights the close and productive relationship between Sandia and the University of New Mexico, where Travis Scholten is working towards a Ph.D. in physics. Both the intellectual environment of UNM and the high-powered supercomputing available at Sandia were critical ingredients in achieving our research goals.

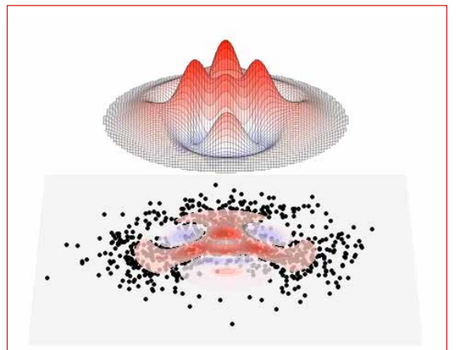

The black dots represent 1,000 numerically generated data points. The contour map created from the model illustrates that the estimate of the quantum state fits the data well. The estimate has a high degree of accuracy, as indicated by the fact the fidelity with the true state is close to 1.

MODELING OF ARCTIC STORMS WITH A VARIABLE HIGH-RESOLUTION GENERAL CIRCULATION MODEL

MARK TAYLOR, PI
ERIKA ROESLER
PETER BOSLER

PROCESSING HOURS
416,078

elroesl@sandia.gov

"Our team is composed of atmospheric scientists with backgrounds in physics, math, and computer science. We are driven by the need to develop the best climate model for predicting the future climate of this country."

– Erika Roesler

The Department of Energy's (DOE) Biological and Environmental Research (BER) project, "Water Cycle and Climate Extremes Modeling" is improving our understanding and modeling of regional details of the Earth's water cycle. Large uncertainties exist in the ability of today's climate models to simulate regional water cycle extremes such as storms, droughts, and floods. Increased model resolution is a key tool for improved modeling of extremes. However, global high-resolution simulations are computationally expensive and usually performed at leadership computing facilities. Work at Sandia has led to dramatic improvements in scalability of DOE's atmosphere component model, making these high-resolution simulations possible, and at the same time also supporting a variable resolution capability where high resolution is used over fixed locations on the globe. This new capability allows scientists to gain insight into high-resolution model behavior over a smaller domain at a fraction of the computational cost.

Our team is using variable resolution to investigate storms in the Arctic. Strong storms in the Arctic are believed to further break-up sea ice. The sea ice extent affects storm tracks, and accurate predictions of the sea ice extent is desired by many agencies.

Currently, it is unknown if or how powerful storms will change in the future climate. To help answer this question, our team developed a variable resolution region in the Arctic to compare the number of storms and their properties in the low, high, and variable resolution simulations, using Red Sky and Sky Bridge for the simulations, data processing, and analysis. We developed a new storm-detection tool used to identify and track storms. As the resolution of the model is increased, it can capture small-scale storms that are not possible to resolve at low resolution, establishing our new variable resolution capability as an effective tool that will allow the climate modeling community achieve a better understanding of Arctic storms.

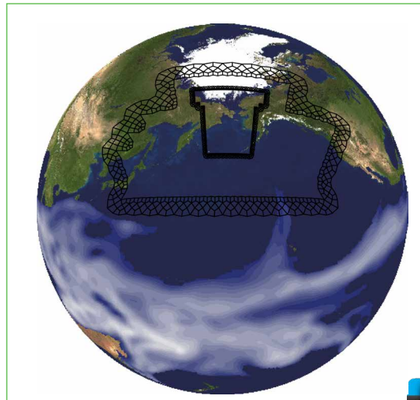

This figure shows the northern Pacific Ocean basin. The colored contours represent the total amount of water in the atmosphere at a given time and location as simulated by the atmospheric model. The lighter, whiter color indicates the presence of a higher water concentration, which also indicates where precipitation is likely to occur. The model was run in a variable resolution configuration using the spectral element atmospheric dynamical core. Areas of higher-resolution in atmospheric models develop more realistic storms. This grid will be used to study Arctic storms, polar lows, and the energetics of storms hitting the Alaskan coast. The resolution in the smallest area outlined over the Bering Sea has an effective resolution of about 13 kilometers (1/8-degree). The area outlined in the northern Pacific Ocean basin stretching from the west's Sea of Okhotsk to the east's Gulf of Alaska has an effective resolution of about 27 kilometers (1/4-dgree). For the remainder of the globe, the effective resolution is about 110 kilometers (1-degree).

WIND TURBINE WAKES

CHRISTOPHER KELLEY, PI
DAVID MANIACI
BRIAN RESOR
Sandia National Labs

XIAOLEI YANG
DANIEL FOTI
FOTIS SOTIROPOULOS
University of Minnesota

PROCESSING HOURS
9,219,642

clkell@sandia.gov

"Having access to powerful computing resources at Sandia allows our Wind Energy Technologies group to conduct innovative research that advances our goal of reducing the cost of wind energy."

– Chris Kelley

The total energy produced by a wind farm depends on the complex interaction of many wind turbines operating in proximity with the turbulent atmosphere. Sometimes, the unsteady forces associated with wind negatively influence power production, causing damage and increasing the cost of producing energy associated with wind power. Wakes and the motion of air generated by rotating blades need to be better understood. Predicting wakes and other wind forces could lead to more effective wind turbine designs and farm layouts, thereby reducing the cost of energy, allowing the United States to increase the installed capacity of wind energy.

The Wind Energy Technologies Department at Sandia has collaborated with the University of Minnesota to simulate the interaction of multiple wind turbines. By combining the validated, large-eddy simulation code with Sandia's HPC capability, this consortium has improved its ability to predict unsteady forces and the electrical power generated by an array of wind turbines. The array of wind turbines simulated were specifically those at the Sandia Scaled Wind Farm Testbed (SWiFT) site, which aided the design of new wind turbine blades being manufactured as part of the National Rotor Testbed project with the Department of Energy.

Three specific areas of research were addressed with different simulation cases. First, researchers investigated the persistence of the wake for unique distributions of force along the wind turbine blades. Unique force distributions produced unique wake contours; however, turbulence removed any differences beyond four rotor diameters downstream. The next simulation investigated the effect of scale on wake persistence and showed that two wind turbines of the same design but different scales will produce different wakes. Finally, the simulations of the experimental SWiFT facility quantified the increased power and force fluctuations of downwind turbines. The research findings have already influenced the aerodynamic design of new wind turbine blades being designed at Sandia.

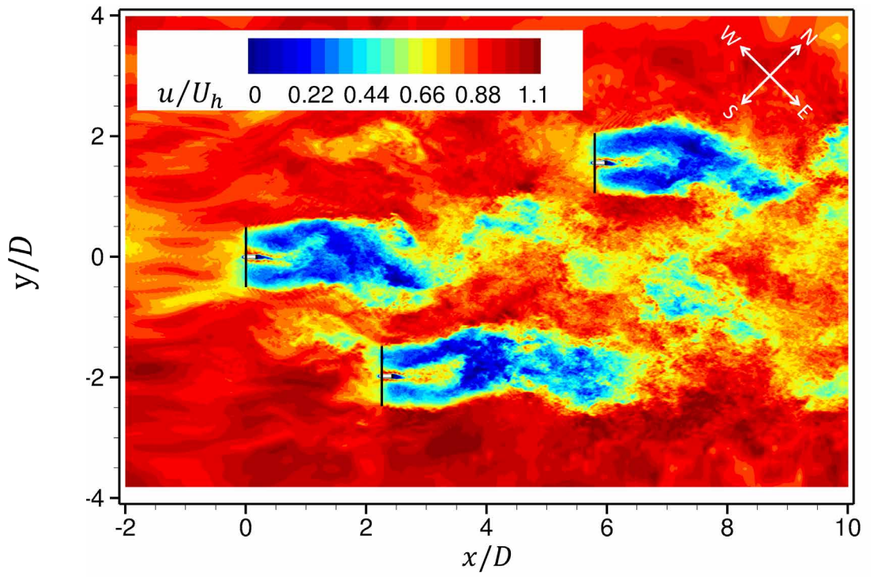

Velocity contour of a southwest wind for three wind turbines closely spaced.

TRANSPORT IN IONIC POLYMERS

GARY GREST, PI
Sandia National Labs

ANUPRIYA AGRAWAL
DVORA PERAHIA
Clemson University

PROCESSING HOURS
20,000,000

gsgrest@sandia.gov

"This study was possible only because of the exceptional high performance computing resources available at Sandia that enabled us to research the dynamics of our system."

– Gary Grest

The association of ionic groups controls the function of several classes of materials with potential uses from clean energy and water purification to sensors and drug delivery. One such class is ionomers, macromolecules that contain ionizable groups. These ionic groups associate into random clusters and control the function of the material. The distribution and morphology of the clusters impacts the conductivity and electrolytic transport of the membrane. Manipulating the association of ionic clusters into well-defined morphologies will offer a new level of control of these materials.

Using high performance computers, Sandia researchers and scientists from Clemson University conducted large-scale molecular dynamics simulations. The collaborators followed the systems for hundreds of nanoseconds to probe the relationship between ionic cluster size and morphology and polymer mobility in ionomers. Million-atom simulations of melts of randomly sulfonated polystyrene, as a model system, have revealed for the first time the evolution of ionic clusters with the strength of the electrostatic interactions. Further, they showed how these changes in cluster morphology impact polymer mobility. For low dielectric media, the ionic groups form ladder-like clusters (as shown in below), but transform to more spherical clusters as the strength of the electrostatic interaction is reduced, resulting in higher mobility of the polymer, where controlling the mobility of the polymers is a key to the design of new classes of membranes for improved performance.

Non-equilibrium molecular dynamics simulations were also carried out to determine how the mobility and shear viscosity of the polymer depends on the fraction of sulfonated groups and valency of the counterion. These simulations, in agreement with the experiments, found that addition of only a small fraction of ionizable groups significantly increases the viscosity of the system. These results show that slight changes in ionic content can serve as an excellent tool to tune desirable properties for membranes, opening the way to new smart materials.

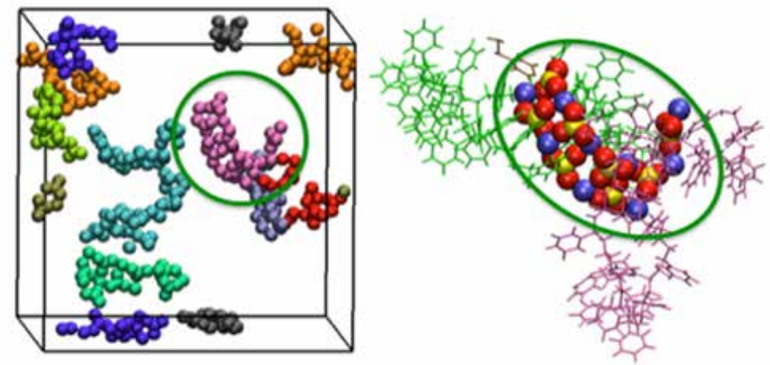

Left panel shows the snapshot of different ionic clusters for 10% sulfonated polystyrene. Different colors represent distinct clusters. Zoomed in cluster at right shows individual atoms S (yellow), O (red) and Na (blue).

RAPID DEVELOPMENT OF AN ICE SHEET CLIMATE APPLICATION USING THE COMPONENTS-BASED APPROACH

ANDREW SALINGER, PI
IRINA TEZAUR
MAURO PEREGO
RAYMOND TUMINARO
Sandia National Labs

STEPHEN PRICE
Los Alamos National Labs

PROCESSING HOURS
1,000,000

agsalin@sandia.gov

"As computational scientists with expertise in math and algorithms, it is challenging to get deep enough into a new science application area to make an impact. This team has made a sustained effort in learning about ice sheets and building relationships with climate scientists, and has been rewarded seeing our code on the critical path of DOE's climate science program."

– Andy Salinger

In the third year of a five-year jointly funded project, Sandia's Albany/FELIX simulation code for ice sheet dynamics is being developed to be run in two modes: (1) as a stand-alone model for scientific investigations, and (2) as part of the land-ice component of coupled climate simulations in the Department of Energy's (DOE's) earth system model. The land-ice component simulates changes to the Greenland and Antarctic ice sheets, including their contributions to global sea-level rise. Using high performance computing (HPC), Sandia has recently completed a controlled mesh convergence study to compute the rate at which the Greenland ice sheet is flowing. The study demonstrated that the HPC code is capable of working accurately, efficiently, and reliably on larger scales (over 1 billion unknowns). The study also identified, for the first time, the minimum number of vertical levels needed in the mesh to maintain accuracy.

The rapid development of this impressive simulation capability was enabled by our components-based approach to computational science. Under our "Agile Components" strategy, small teams of experts develop independent math libraries, which are designed to be interoperable through software interfaces, and maintained and deployed for subsequent application by following software engineering best practices. The Albany code benefited from dozens of previously developed capabilities, and in turn, has improved the foundational code base for other applications.

Funded by the DOE Office of Science Scientific Discovery through Advanced Computing program (SciDAC), this project is a collaborative effort between the Climate (BER) and Applied Math (ASCR) programs. Partnering with Los Alamos National Laboratory, Sandia National Laboratories analysts have recently integrated the Albany/FELIX code into LANL's MPAS-LI code, the land ice component of DOE's Accelerated Climate Model for Energy (ACME) earth system model. In ACME, the Albany code will be coupled with the atmospheric, ocean, land, and sea ice components, and used in climate projections in support of DOE's energy and security missions. Demonstrating the capability for interagency computing, this work primarily uses the computing systems as National Energy Research Scientific Computing Center (Hopper) and Oakridge Leadership Computing Facility (Titan).

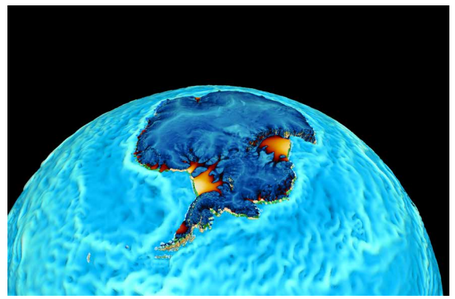

Visualization of the computed surface velocity of the Antarctic ice sheet, which clearly shows the fast-moving ice shelfs (red-orange colors). For visualization, the vertical dimension of the ice sheet is stretched and the figure is superimposed on an ocean model simulation result, foreshadowing our ongoing work in integrating the ice sheet model into global earth system models. (Graphic by P. Wolfram [LANL]).

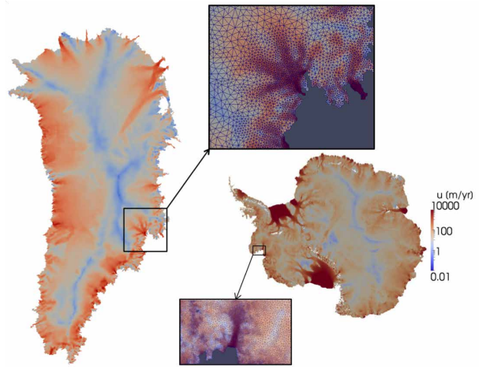

Computed surface velocities for both the Greenland and Antarctic ice sheets. The insets show how our unstructured-grid approach allows for enhanced resolution in the critical locations of fast-moving ice streams at greatly reduced computational costs over uniform grid approaches.

FUNCTIONAL POLYMERS FOR WATER DESALINATION

SUSAN REMPE, PI
MANGESH CHAUDHARI
JEFF BRINKER
Sandia National Labs

DAVID ROGERS
University of South Florida
(Formerly of Sandia)

YING-BING JIANG
SHAORONG YANG
University of New Mexico

PROCESSING HOURS
1,915,031

slrempe@sandia.gov

"We are excited to extend our research in the coming years to develop ion-selective channels for diverse applications such as electrodialysis, battery electrolytes, and separating toxic metals from water."

– Susan Rempe

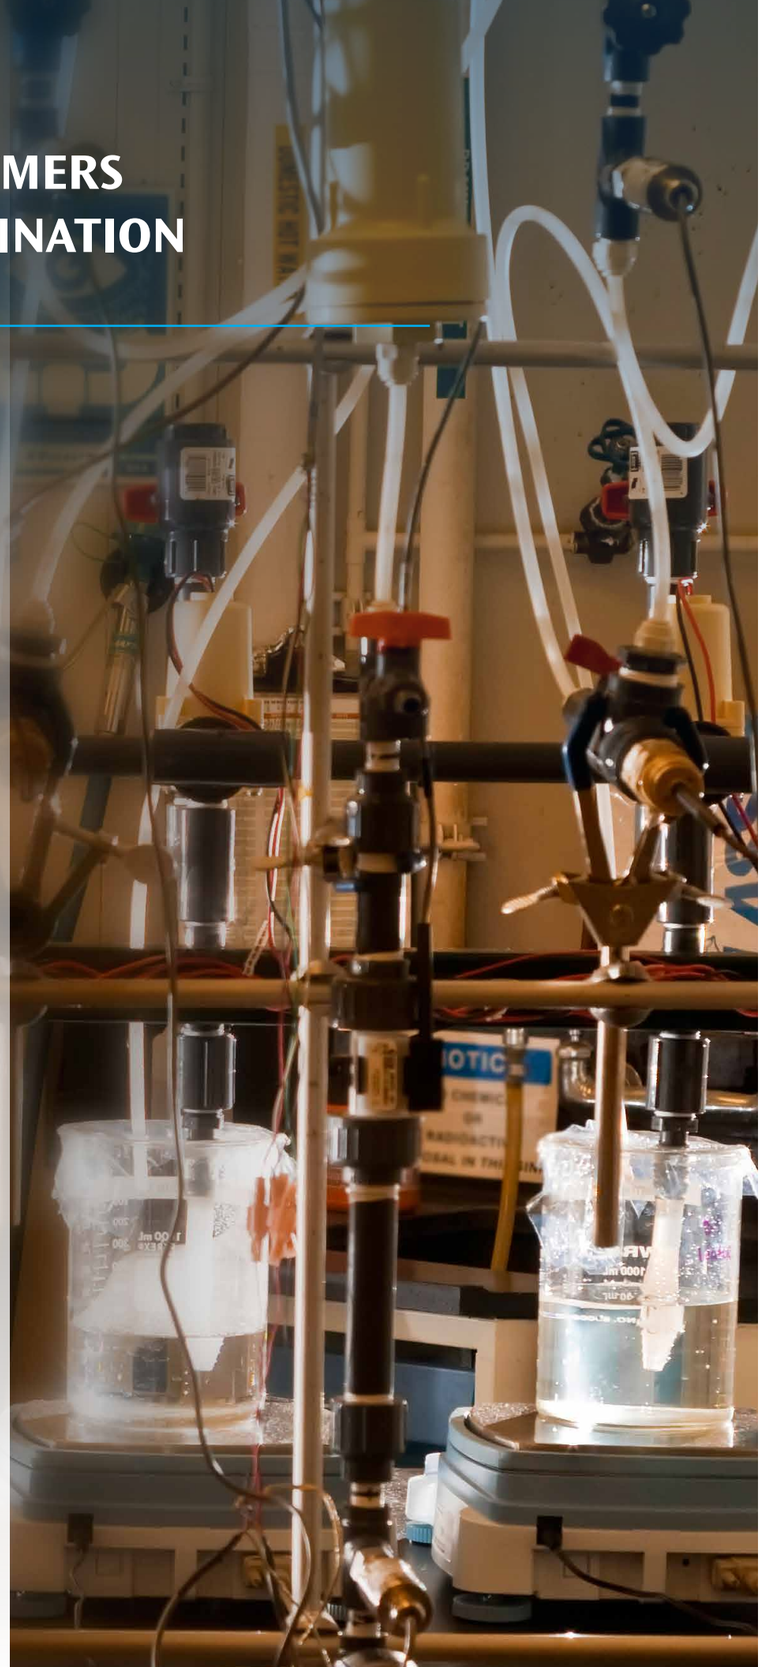

Water, energy, and food are interconnected and present a technological challenge that is hard to over-estimate. As early as 2003, Nature Magazine reported that more than 1 billion people lacked access to clean water and predicted that over the next two decades average per person water supply would be reduced by a third. More recently, a brighter future could rely on advanced water desalination techniques that use less energy to produce clean water at a lower cost, such as wind-powered reverse osmosis

Currently, polymers form the active component of reverse osmosis membranes used to separate water from salty solutions. Compared with cellular membranes, synthetic membranes lag behind in terms of water selectivity and permeation rate. To address this problem, researchers from Sandia and the University of New Mexico (UNM) partnered to develop synthetic membranes that mimic nature, called Biomimetic Membranes.

Membrane design and fabrication involved both theoretical and experimental efforts. To understand the relationship between molecular structure and selective water permeation, Dr. Susan Rempe and co-workers from Sandia applied molecular simulation techniques to two systems: 1) aquaporin Z, a water-selective channel protein in cellular membranes; and 2) synthetic nano-pores "decorated" with functional polymers to mimic the biological channel behavior. Colleagues at UNM used an advanced atomic layer deposition (ALD) technique to fabricate synthetic membranes. To estimate the work needed to move water and ions through the permeation pathways, the simulation specialists computed free energy profiles. Both the synthetic and biological channels show lower free energy barriers for water compared with competing ions like sodium and chloride. A systematic arrangement of functional polymers, achieved using ALD, improved water permeation rates by at least five-fold when compared with commercial membranes. An assessment of local coordination structure of water and ions along the permeation pathway in biological and synthetic membranes provides clues to the observed differences in behavior.

The researchers performed simulations using the high performance Red Mesa cluster at Sandia. The Gromacs molecular dynamics simulation package was used to compute the ion and water free energy profiles. Electronic structure calculations were also performed to obtain more details about the ion hydration structures and free energies.

The first implementation of the water-selective membranes won an R&D 100 award. Licensing and commercialization are being pursued.

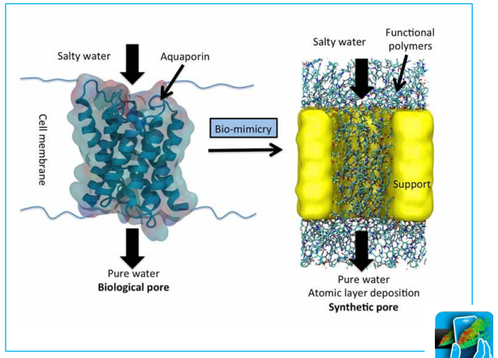

The biological aquaporin water-selective channel (left) is mimicked in the synthetic system using an advanced atomic layer deposition technique to decorate silica supports with functional polymers (right).

PHYSICS ANALOGUES FOR HYDRODYNAMIC EVENTS

DARYL DAGEL, PI
MATT HOPKINS
OLGA SPAHN
Sandia National Labs

RICK RUSSO
XIANGLEI MAO
Lawrence Berkeley Laboratory

CLARE KIMBLIN
RUSTY TRAINHAM
Special Technologies Laboratory

PROCESSING HOURS
202,202

djdagel@sandia.gov

"If the physics scales between macro and micro, the flexibility of doing experiments on 'micro-scale makes the potential connection to real-world explosions powerful. I find it challenging to visualize the evolution for the plasma and shockwave in the lab and translate that to the simulation in a way that captures the essential physics."

– Daryl Dagel

Some of the same physical interactions of matter and energy that occur in large explosive tests are also present in microscale tests, and the two sets of behaviors can be linked in a quantitatively scalable manner. For example, spectral optical emission is a function of atomic and molecular electronic state transitions. If the same environment that is present at the large scale can be produced for the same species at the microscale, then the same optical signatures are produced.

Setting up even a single series of large-scale experiments requires significant space (e.g., a desert) and involves many people: designers, diagnostics specialists, and safety personnel. Even one experiment takes months of planning and coordination. "Success" is achieved by producing a small number of tests. Replication is expensive and the entire series can be impacted by something entirely outside the controlled conditions, such as an unanticipated weather event.

To produce a microscale environment analogous to a large-scale event, we employ pulsed nanosecond and femtosecond lasers to ablate surface material. Compared with large-scale experiments, those done at the microscale are less expensive, extremely repeatable, simpler to coordinate, inherently safer, and take place in a more controlled environment.

We use theoretical models to bridge the two scales, such as Taylor-Sedov theory, describing the rate of shock front expansion, or rescaling the equations of hydrodynamics to recast a problem in one scale to another. For example, microseconds and meters at the macroscale may be equivalent to nanoseconds and microns in the microscale. The question persists: Can a microscale experiment produce a strong enough physics analogue to mimic that of a large-scale experiment?

The laser pulses are of sufficient energy and duration that thermal effects cause the creation of multiple phases (solid, liquid, gas, plasma), as happens in large-scale events. A difficulty for nanosecond pulses is that the incoming laser energy interacts with the outflowing gas and plasma, complicating the energy deposition profile to the surface. In contrast, a femtosecond laser pulse does not suffer from laser-plasma interaction, but the exact material emission mechanisms are not as well understood.

Using modeling and simulation allows us to explore the environments produced in both microscale and full scale experiments and estimate output sensitivities to known specific operational changes (e.g., pulse duration or material dependence). In fact, we often employ the same simulation tool, the shock-hydrodynamics CTH code, for both cases.

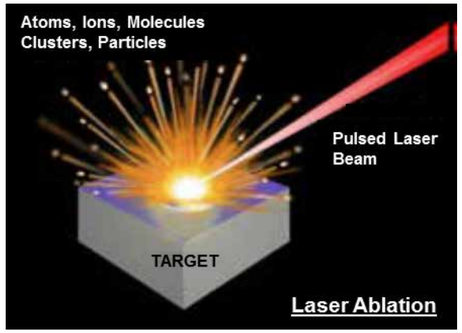

Depiction of laser ablation process. The incoming laser energy is deposited on the surface, resulting in an expanding gaseous shock. Other material is produced, depending on laser characteristics.

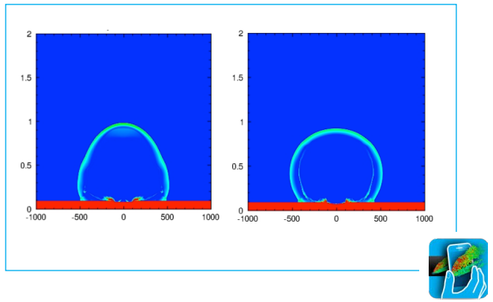

Density profile (shock front location) for a nanosecond laser ablation simulation using CTH code, and density profile (shock front location) for a ns laser ablation simulation in CTH with different temporal and spatial profiles than other image, but containing the same energy,

SIMULATION OF SURFACE RADIATIVE TRANSFER FOR THE SPECTRAL RECOGNITION OF SOLID TARGET COMPOUNDS

THOMAS REICHARDT
THOMAS KULP
PATTY HOUGH

PROCESSING HOURS
1,084,330

tareich@sandia.gov

"About 10 years ago I began participating on R&D teams addressing the challenges of detecting aerosols and solid deposits. Recognizing that numerical tools were being developed that, when coupled to parallel computing capabilities, could impact this area of research, I have made it the current focus of my career."

– Thomas Reichardt

Remote detection of a surface-bound chemical relies on the recognition of a pattern, or "signature," that is distinct from background materials. Such signatures are a function of a chemical's fundamental optical properties but also depend upon its specific morphology. Importantly, the same chemical can exhibit vastly different signatures depending on the size of particles composing the deposit.

Through a concerted modeling and measurement effort, we are developing a numerically invertible solution to account for such morphological effects on surface-deposited chemical signatures. In parallel, we are fabricating and fully characterizing the particle size, shape, and packing density of deposits of materials of known refractive indices, and performing high-spectral-resolution laboratory reflectance measurements of these deposits. Tying these two efforts together is Sandia's Dakota package, used in optimizing agreement between modeling results and data.

These efforts are a component of NA-22's HARD Solids (Hyperspectral Advanced Research & Development for Solids) venture, with the goal of developing validated models for predicting and understanding the signatures of solid materials, with direct application to hyperspectral image detection. The project plan incorporates a spiral development, addressing increasing system complexity with each turn, such as considering intimate mixtures of more than one chemical, as well as chemical deposits of varying optical thickness.

So far our numerical approach has been applied solely to silica powders over the 3.6 to 16 um wavelength range, which provides a rigorous initial test. The imaginary refractive index of silica varies by almost five orders of magnitude over this wavelength range, while the particle size distributions include particles ranging from 1 um to 1 mm in effective diameter.

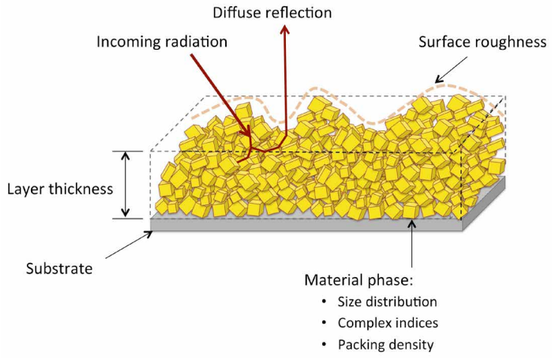

X-ray tomographic measurements of particulate media serving to validate radiative transfer models.

CALIBRATION OF ACCURATE AND AFFORDABLE ENGINEERING TURBULENCE MODELS

JEREMY TEMPLETON, PI
MYRA BLAYLOCK
STEFAN DOMINO
JOHN HEWSON
RAJ KUMAR
COSMIN SAFTA
ALESSIA STEWART

PROCESSING HOURS
7,048,278

jatempl@sandia.gov

"This project combines two of SNL's unique strengths, UQ and HPC, to tackle a long standing problem in physics: modeling turbulent flows. We've learned much about how to get more information about turbulent flows into the hands of engineers, and it's exciting to think about what they might be able to design in the future using our methods."

– Jeremy Templeton

Turbulent flows are encountered in a wide range of engineering applications at Sandia, including re-entry for nuclear weapons systems and combustion physics. Developing predictive simulation strategies for turbulence is challenging because of its multiscale nature in which eddies and whorls range in size from that of an aircraft down to microns. Our goal is to create a strategy that uses targeted, expensive calculations that account for most of these motions to calibrate inexpensive models that can be used by engineers to design and analyze real-world systems.

Using institutional high performance computing we developed methods building on Sandia's internationally recognized expertise in uncertainty quantification to statistically infer engineered models from data. By leveraging our knowledge of turbulence theory, we can use results from very expensive turbulence simulations that contain all of the fluid's motions and determine what answer an engineering model should provide. Bayesian calibration then compares what the model should produce with what it will produce with different model parameters and identifies the parameters that best reproduce the data.

Bayesian calibration approaches also provide uncertainty estimates of the parameters. This is important so that engineers can determine if the models are accurate enough to be useful for their applications. We use forward propagation of uncertainty to evaluate how accurately quantities of interest, like drag or fuel efficiency, are predicted given how well we know the model parameters.

Our team has learned two important lessons from this work. First, it is possible to use data from highly accurate and expensive simulations to calibrate engineering models of turbulent flows. Second, is that in addition to this fundamental information, the calibrated model must account for the acceptable cost of the engineered model and adapt to it. In the end, data analysis has yet to find a "universal" turbulence model, but we can use it to find the best model for a given cost.

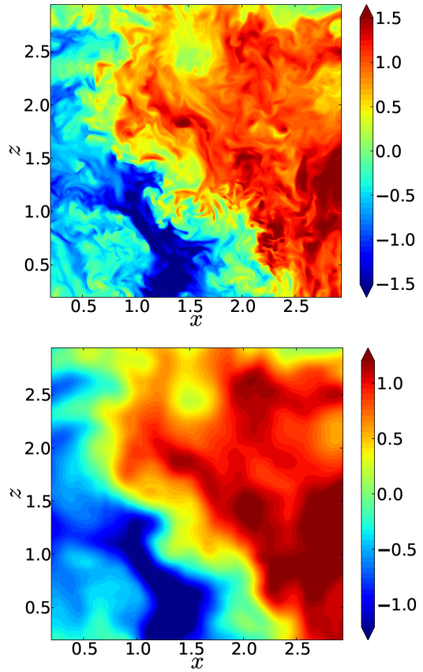

Top: Snapshot of a turbulent flow from a simulation resolving all the scales of motion. Bottom: What an engineering calculation should produce to be consistent.

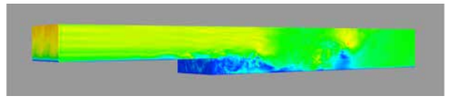

An engineering model of turbulent flow over a backward-facing step.

MOLECULAR MODELING OF NANOPARTICLE-CELLULAR INTERACTIONS TO OPTIMIZE NANOPARTICLE DELIVERY OF MEDICAL COUNTERMEASURES

CARLEE ASHLEY, PI
CHRISTINA TING, PI
ERIC CARNES
OSCAR NEGRETE

PROCESSING HOURS*
1,133,679

ceashle@sandia.gov

"Sandia is one of the few institutions with expertise in synthesis of systematically variable nanoparticles, as well as simulation and modeling. It's exciting to see our collaboration yield nanoparticle delivery technologies with such enormous potential in biological defense, as well as various medical applications."

– Carlee Ashley

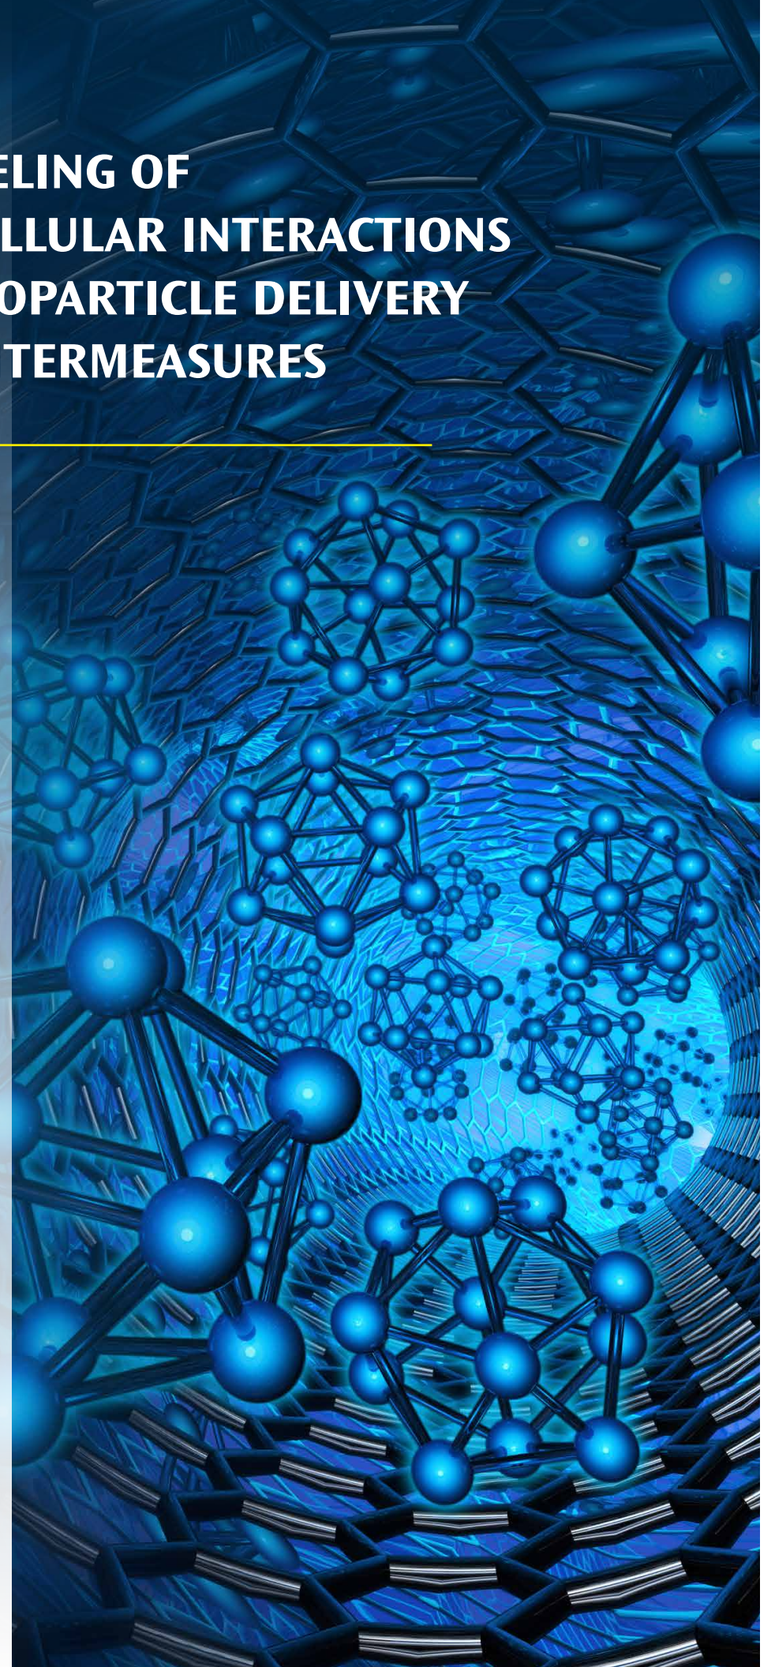

Deadly pathogens, like the Ebola virus, kill millions of people each year, and the accelerating prevalence of pathogens resistant to most, if not all, drugs threatens to bring about a "post-antibiotic" era. To address this potentially devastating problem, a team of researchers at Sandia are developing novel strategies for protecting warfighters and civilians against emerging infectious diseases, bioterrorism agents, and multi-drug-resistant pathogens. Their efforts primarily focus on loading nanoparticles with pathogen-fighting drugs and designing the particles to accumulate in organs and cells that are infected with the pathogen, with the overall goal of more effectively treating the disease while minimizing any side effects caused by the drug.

To effectively combat pathogens, however, nanoparticle delivery vehicles must effectively transverse complicated biological environments, such as the blood, and efficiently deliver their payload to sites of infection. In recent years, the team has employed molecular modeling, coupled with the transition path finding algorithms, to more rapidly design nanoparticles with the necessary properties. Specifically, they have developed models that help elucidate interactions between nanoparticles and blood proteins, non-target cells (e.g., healthy cells), and target cells (e.g., cells infected with a pathogen), the sum of which govern the ultimate efficacy of a nanoparticle-based therapeutic.

Recent results obtained from computer simulations run on Red Sky reveal favorable interactions between the interface of the nanoparticle and the cell membrane critical to overcoming the free-energy penalty for nanoparticle translocation through the hydrophobic core of the cellular membrane. Future work, which will be funded by a Grand Challenge Laboratory Research and Development project, will focus on understanding and optimizing endocytosis mechanisms with the ultimate goal of designing nanoparticles that deliver drugs to infected cells without negatively impacting healthy cells.

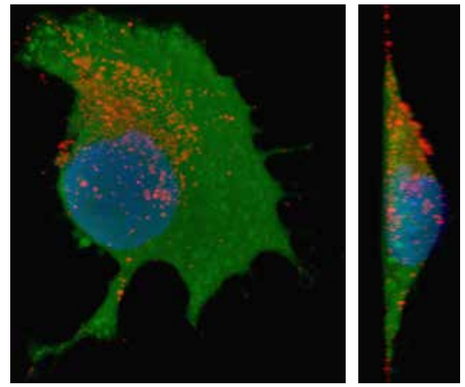

Front and side views of a human cell (cytosol = green, nucleus = blue) that has internalized ~3500 nanoparticles (red).

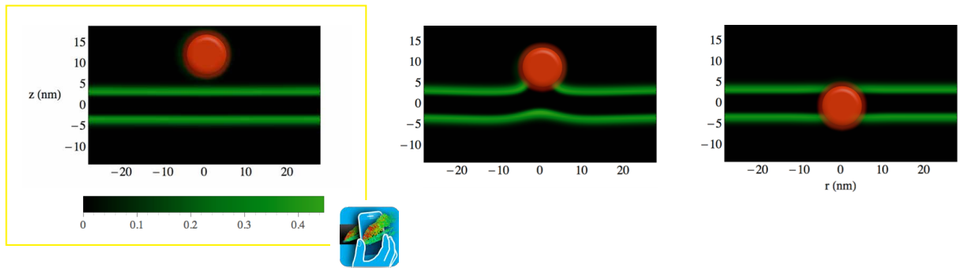

Nanoparticle (red) crossing the lipid bilayer. Contour plots show the hydrophilic density of the lipid bilayer (green) in cylindrical coordinates.

*The computer runs for this work were completed in FY14. Analysis was done FY15.

MOLECULAR SIMULATIONS OF FAILURE OF POLYMERIC STRUCTURAL ADHESIVE

DAVID REEDY, PI
BOB CHAMBERS
JAMIE KROPKA
LINDSEY HUGHES
MARK STAVIG
MARK STEVENS

PROCESSING HOURS
4,035,894

msteve@sandia.gov

"These simulations have produced several important results that reveal new insight into the connection between the nanoscale interfacial dynamics and the macroscopic mechanics. It is very exciting to see an opening into understanding this complex and hidden world."

– Mark Stevens

Adhesives are an important and common part of manufactured products. They are an essential component in today's advanced composite materials, as adhesives are used to bond the separate materials of the composite together. We have a great deal of experience and knowledge about adhesives, but our fundamental understanding at the molecular level is limited. Because measuring the processes at the buried interface between the adhesive and adherend is difficult, we have little knowledge about the interface. We have used simulations to reveal molecular processes at the interface during fracture. Understanding the origin of fracture in adhesives is a fundamental issue that needs to be better resolved for many engineering applications. For the basic geometry of an adhesive between two solid surfaces we found that the stress is largest in the corners and that cracks initiate in corners. We thus focused our simulations on the interfacial behavior at corners.

Simulations were run using LAMMPS molecular dynamic code on Red Mesa and Red Sky systems. These computer simulations reveal new insights into the deformations and bond breaking that occurs at corners. Our tensile pull simulations of a coarse-grained highly cross-linked polymer system with open sides show failure by crack initiation in the corners. The open sides result in a well-observed contraction of the sides. Moreover, the failure strain strongly decreases as the system size is increased. For the thin systems simulated, system size plays a large role. Extrapolating this result indicates that systems with a thickness of about a few microns will have failure strains similar to those observed in experiments. As a result, we have found a major source of the large difference in failure strains between simulation and experiment. At some size, the decrease in failure strain must halt. This suggests an important length scale in the polymer adhesive system that has not been characterized and drives the material properties. Future work will characterize the structure associated with this length scale.

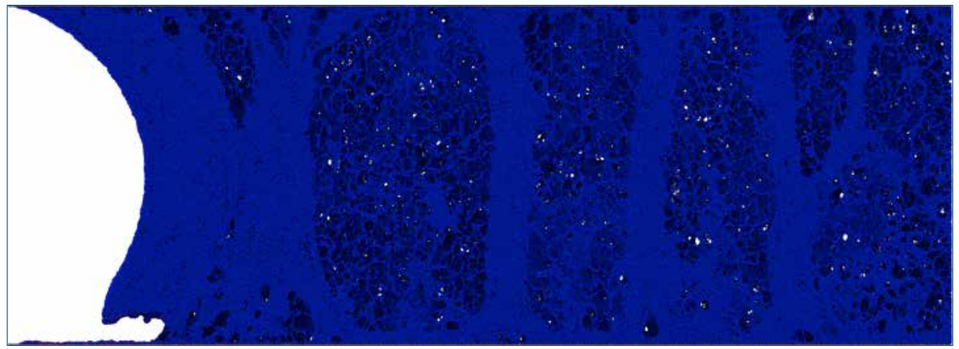

An image of a highly cross-linked polymer adhesive system shows a crack forming in lower left corner. The voiding due to cavitation in the system is also visible, but does not lead to a crack (Only the left half of the system is shown).

CONTROLLING OPTICAL EMISSION FROM MICRODEVICE DISCHARGE

RON MANGINELL, PI
MATT HOPKINS
MATT MOORMAN
CHRIS MOORE

PROCESSING HOURS
2,561,409

rpmangi@sandia.gov

"The field of microplasmas is fertile ground for fundamental physical studies and applications directly relevant to Sandia's NW and National/Homeland Security missions. The modeling and experimental teams are individually and jointly producing exciting results, both fundamental and applied, and we hope to build on this success to directly impact the labs' missions."

– Ron Manginell

The use of microscale discharge devices for producing light with very specific characteristics is becoming a foundational technology in many new areas, including as light sources, laser gain media, chemical and biological sensing, medical, nanosynthesis, and disinfection. Given Sandia's enduring interest and capabilities in developing microscale technology, discharge-based light emitting microdevices is a natural focus.

The basic design of a light emitting microdevice of the sort being investigated is based on sandwiching a dielectric (e.g., polyimide) between anode and electrode surfaces (e.g., aluminum and silicon), applying a voltage drop of sufficient magnitude to cause discharge, but not so large as to cause a runaway event (arc). This is often termed a "glow discharge" and shares some physics with the operation of a fluorescent lamp (although the gas itself produces the final photons, not a coated surface). As the discharge sets up, the voltage drop decreases and a steady operation can be reached.

Which wavelengths (colors) are generated depends on gas species, and which excited states are created and then undergo spontaneous emission of a photon. The creation of excited states depends in turn on gas pressure, electron energies, and other factors. By modifying the operating conditions we can "tune" the microdevice to prefer one or another of the excitations, and in combination with modifying the geometry of the device we hope to effectively select which photons are emitted from the device.

Although design and fabrication of microscale devices is becoming more efficient, it is very beneficial to produce and assess designs via simulation. In addition, simulations can give additional insights that may not be as easily uncovered through experiment because the full 3D transient simulation can be interrogated for information more simply than an experiment. The fully kinetic low temperature plasma modeling code Aleph was used to simulate the transient discharge, including plasma chemistry, surface response, and driving circuit behavior.

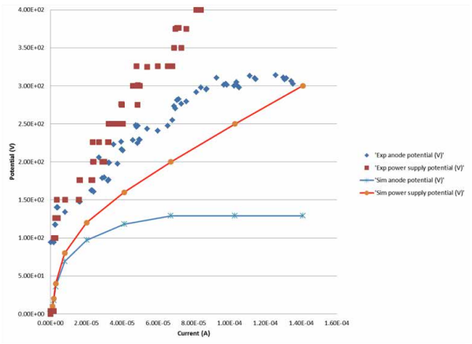

Comparison of experimental and simulation voltage vs. current operation conditions for a 655 Torr neon system. Note the levelling off as I/V approaches 0 in both cases. Simulation is a 50 µm diameter vs. 100 µm radius in the experiment.

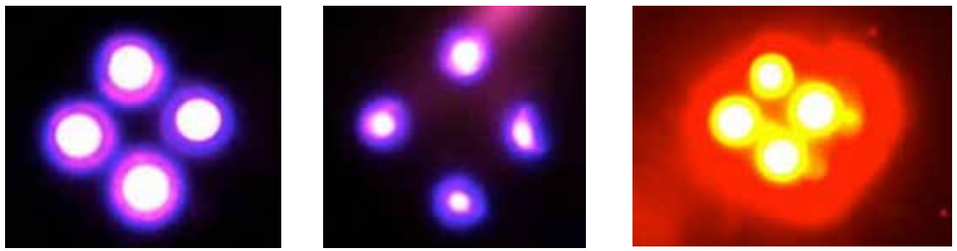

Images of 2x2 microscale discharge devices containing (from left to right) helium, argon, and neon. Images are looking down into cavities each 100 µm in diameter. Dielectric separators are 10 µm long.

KOKKOS PROGRAMMING MODEL AND LIBRARY

CARTER EDWARDS, PI
CHRISTIAN TROTT
DANIEL SUNDERLAND
MARK HOEMMEN
SIMON HAMMOND

hcedwar@sandia.gov

"Sandia's Center for Computing Research solves a portability barrier on high performance computing many-core heterogeneous architectures by developing and deploying the KOKKOS programming model and library."

– Carter Edwards

High performance computing is in the midst of a many-core revolution driven by rapidly evolving microchips with diverse and heterogeneous architectures, increasing concurrency (thread count), and decreasing memory per thread. Our HPC applications and domain libraries face a major challenge to effectively use these microchips with a single code base to be performance portable to these evolving architectures. Performance portability requires that our codes leverage increasingly complex on-chip concurrency and manage performance-critical constraints on their memory access patterns. Several open-collaboration efforts (e.g., OpenMP, OpenACC, and OpenCL) are defining programming models to enable portable concurrency; however, they fail to address architecture-dependent constraints on memory access patterns.

The KOKKOS performance portable programming model and library were developed to enable on-chip concurrency and manage architecture-dependent constraints on memory access patterns. The goal is to maximize the amount of application and domain library code that is portable across diverse architectures and still obtain the same (or nearly the same) performance as a variant of that code written specifically for that architecture. The KOKKOS development team has ongoing research and development (R&D) co-design collaborations with Sandia's and vendors' microchip and system engineers to rapidly incorporate evolving computer architecture features into KOKKOS, and to provide feedback on the usefulness of potential features. This team is also providing leadership for DOE laboratories' participation on the ISO/C++ language standard committee with the goal of subsuming KOKKOS features into a future C++ language standard.

The Center for Computing Research Advanced Architecture Test Bed program has been critical to the success of KOKKOS. This program fields leading industry vendors' first-of-a-kind prototype computing architectures for early application porting and performance analysis across DOE's Sandia National Laboratories, Lawrence Livemore National Laboratory, and Los Alamos National Laboratory Tri-Lab complex. Test beds included hardware and software environment prototypes for the forthcoming Advanced Simulation and Computing Trinity and Sierra platforms. KOKKOS R&D used these test beds with a suite of mini-applications to continually verify performance portability of KOKKOS programming model concepts and library implementation.

Application and domain library projects are adopting KOKKOS to enable performance portability through the ongoing HPC many-core revolution. Within the DOE laboratories, these include Trillinos, Zoltan, LAMMPS, Albany, ASCR Multiphysics MHD, FASTMath SciDAC, Empress, SHIFT, and others. Outside of DOE, the University of Utah, King Abdullah University of Science and Technology, Helmholtz-Zentrum Dresden-Rossendorf German research laboratory, and Swiss Supercomputing Center have expressed interest in adopting KOKKOS for new projects.

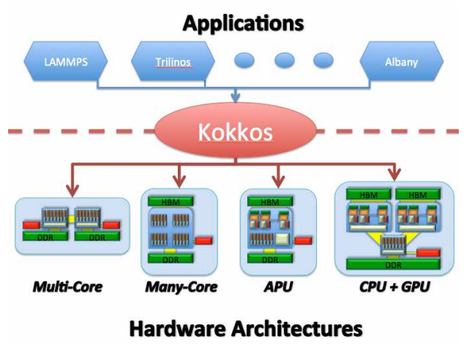

KOKKOS enables high performance computing applications to achieve performance portability across heterogeneous many-core architectures.

TIME-DEPENDENT DENSITY FUNCTIONAL THEORY

RUDOLPH MAGYAR, PI
ANDREW BACZEWSKI
LUKE SHULENBURGER
MICHAEL DESJARLAIS

PROCESSING HOURS
14,143,782

rjmagya@sandia.gov

"These simulations are computationally expensive, but the insight gleaned by boldly describing the physics in very general terms allows us to probe models that are otherwise hard to test."

– Rudy Magyar

This work represents the first application of time-dependent density functional theory (TDDFT) to warm dense matter. Many ensemble-averaged quantities can be computed in Born Oppenheimer density functional theory (BODFT) in a linear response framework, but TDDFT allows us to simulate the response explicitly, to test assumptions in BODFT, and to access physics where explicit representation of the electron dynamics is essential for capturing the key physics.

Using TDDFT, Sandia scientists have advanced the capability for simulating non-equilibrium electron-ion dynamics. While these dynamics are important in many areas of advanced engineering, the Sandia team focuses on matter at extreme pressures and temperatures occurring, for example, in the core of exo-planets and inertial confinement fusion. Application of these new tools may advance X-ray Thomson scattering (XRTS) capabilities, which are an important experimental probe of the temperature, density, and ionization state of opaque plasmas.

This work establishes a fundamental method to explore dynamic response and transport properties in situations where reference data is unavailable and experiments are costly and time consuming. Scientists can use this tool to find XRTS cross-sections and stopping powers under a wide range of conditions. This capability provides both significant theoretical and practical utility. It allows exploration of many of the underlying assumptions and inconsistencies of traditional plasma models and enables the critical parameters used in modeling inertial confinement fusion to be computed from first principles. Further, these simulations are computationally intensive, posing new challenges and opportunities for advancement of high performance computing. Aspects of this work, for example, suggest a way to deliver more scalable density functional molecular dynamics simulations.

This work conclusively shows the effect of modeling the bound and free electrons independently and choice of ionization model has on predictive simulations of dynamic materials properties that are not directly accessible in Born Oppenheimer formulations. Stopping power simulations are in good agreement with existing literature and the high degree of controllability in our simulations allows us to probe the underlying causes of the variability in experimental and tabulated data.

Follow-on work simulating the stopping powers of various ions in warm dense matter is planned within the National Nuclear Security Administration's Science Campaigns. One short-term goal is to improve the accuracy of forces on the ions and to develop a mathematically robust time integration framework for coupled electron-nuclear motion. A long-term target is to explore the use of this capability to model polymers and other materials addressed with atto-scale laser pulses on ambient materials such as polymers pulses by the Linac Coherent Light Source.

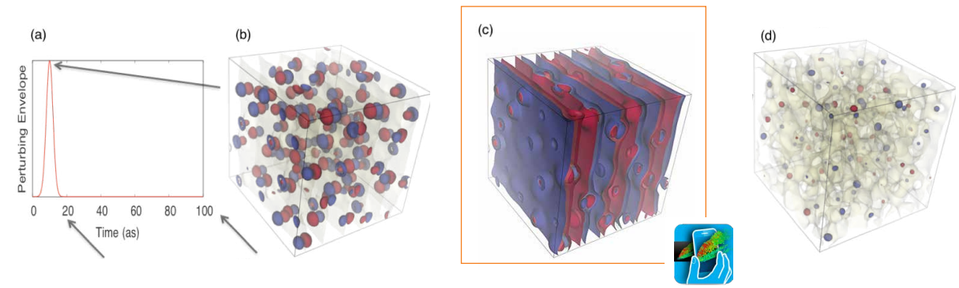

The electronic charge density response of warm dense Beryllium (3x compressed, 150kK). (a) Illustrates the envelope of the probe pulse that hits the system and (b-d) illustrate the density response at the peak of the pulse (b), 10 as after the pulse (c), and 90 as after the pulse (d). Red (blue) isosurfaces bound positive (negative) regions of density response, and the light yellow isosurfaces coincide with the nodal surface.

UNCERTAINTY QUANTIFICATION OF FLOW PREDICTIONS

SOPHIA LEFANTZI, PI
SRINIVASAN ARUNAJATESAN
LAWRENCE DECHANT
JAIDEEP RAY

PROCESSING HOURS
31,943,468

slefant@sandia.gov

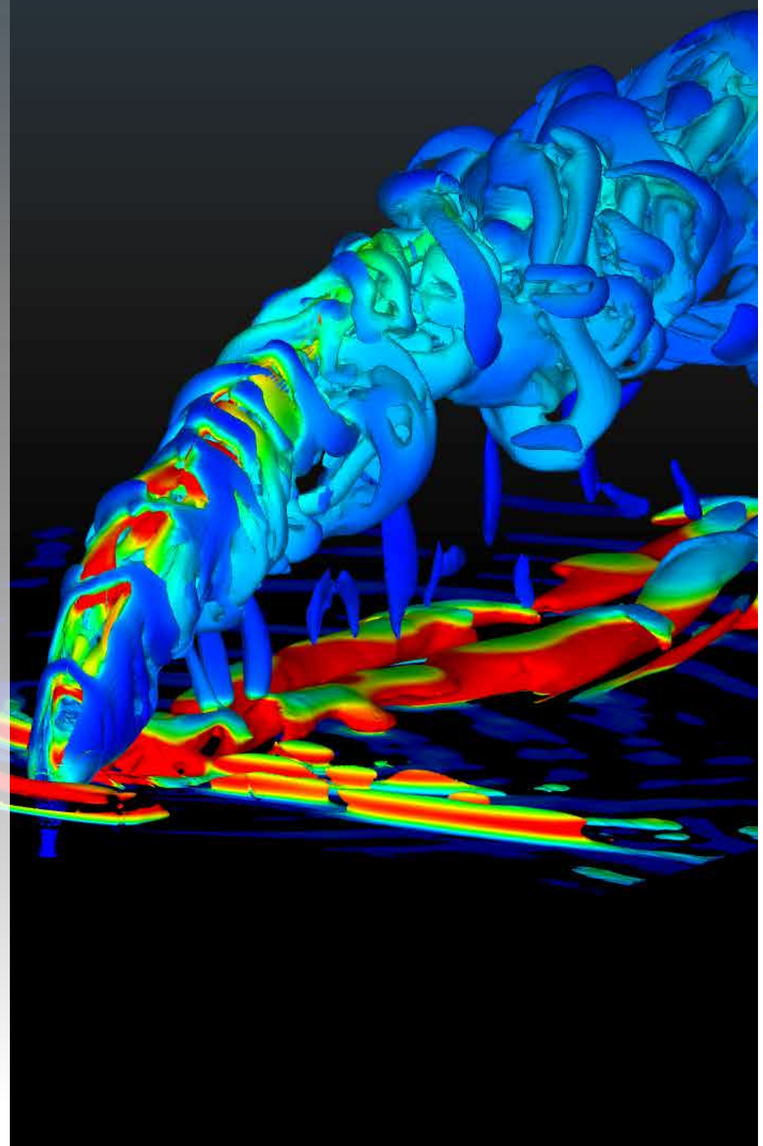

"I am excited to work with brilliant scientists who are world class experts in their respective fields. We tackle hard engineering problems to improve our stockpile reliability."

– Sophia Lefantzi

The team created fast and reliable models of fluid jets in cross-flow (JICF). Applications of JICF include high-speed flight control, gas turbine combustor cooling, and fuel injection. In the cases studied, JICF idealize the turbulent aerodynamic processes that are associated with flight control. The end customers are tasked with Life Extension Programs (LEPs) and stewardship of the U.S. stockpile. They use computationally expensive 3D Reynolds Averaged Navier Stokes (RANS) models to simulate the flight of re-entry vehicles and bombs.

A novel Bayesian calibration method has been developed to estimate turbulence model parameters embedded inside computationally expensive engineering models, such as JICF. Typically, turbulence model parameters are selected from published problems, but they are rarely predictive for real-life applications. Here, the turbulence model parameters are estimated from data similar to the final use-case as probability density functions instead of as single values. The new Bayesian method quantifies the uncertainty due to limited calibration data and explains the inability of models to reproduce observations exactly.

The team has demonstrated the first successful application of the Bayesian turbulence parameter estimation to engineering models of this complexity and scale. The new method can improve a broad range of JICF models. By reducing simulation errors substantially, the method tightens limits on product deployment uncertainties and informs the experimental data and mathematical techniques for further improvement. Ultimately, the improved models will enable accurate simulations over a broad range of flight specifications.

The Bayesian method helps our customers improve the accuracy of their 3D RANS models for LEP applications. DOE's Advanced Simulation and Computing (ASC) Verification and Validation program funds this work. The project relies on the Tri-Lab shared computing capability. Simulations were performed in the National Nuclear Security Administration's ASC petascale system, Sequoia, located at Lawrence Livermore National Laboratory.

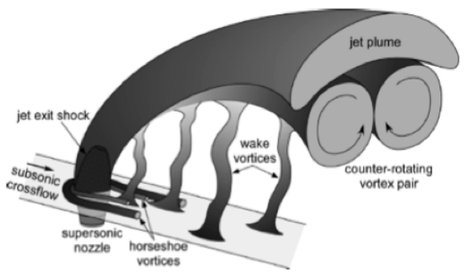

Counter-rotating vortex pair in a jet in cross-flow interaction.

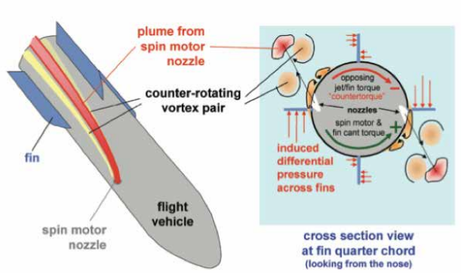

Illustration of jet/fin interaction in which vortices generated by the jet in cross-flow interaction induce pressures on the fins to create a counter torque opposing the torque generated by the spin rockets [source: Beresh EtAl, AIIA Journal, 45(8):1827-1840, 2007].

NEW PLATFORM PROVIDES INNOVATION

STEVE MONK, PI | *smonk@sandia.gov*

Sky Bridge is the primary unclassified scientific computing platform at Sandia, serving all Mission Areas within the Laboratories. At the formal ribbon cutting in December 2014, Laboratories Director Paul Hommert emphasized the importance of the Institutional Computing Program in providing a critical core Laboratories capability. The deployment of Sky Bridge more than doubled the computing cycles available in Sandia's restricted computing environment, while incorporating the same underlying architecture as the existing ASC/TLCC2 compute clusters, thus allowing many existing computer codes to run on the system without any major modifications. Sky Bridge immediately addressed the shortfall in projected computing cycles for Sandia's programs of energy and environment, nuclear weapons Life Extension Programs. From a technological perspective, Sky Bridge is contributing to the readiness of the data center facility and operations personnel for a transition to liquid cooling in future HPC systems and other high-density systems. By comparing the performance and energy costs of Sky Bridge to similar systems (e.g., Dark Bridge) that are air-cooled rather than liquid-cooled, it will also provide valuable data for making return on investment decisions for future purchases. Installation of Sky Bridge enables decommissioning of the Red Sky supercomputer, reduce energy and support costs, and make room for future HPC systems already planned through the ASC and Institutional Computing programs. Sky Bridge's use of liquid cooling (cold plate to the systems processors) and more efficient power distribution (480 V direct to the rack) saved over $700,000 in construction costs while reducing annual operating costs by $120,000.

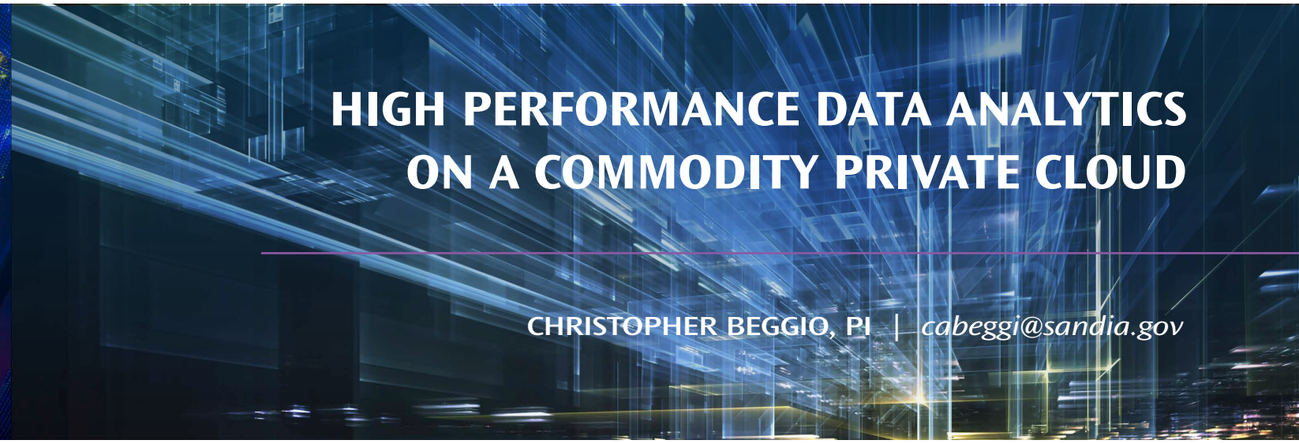

HIGH PERFORMANCE DATA ANALYTICS ON A COMMODITY PRIVATE CLOUD

CHRISTOPHER BEGGIO, PI | *cabeggi@sandia.gov*

Sandia is researching novel methods of analysis for large datasets. The term "Big Data" has been used frequently in recent years to describe a discipline of computer science that attempts to distil an increasing quantity of data into a result or meaningful subset that can be easily visualized and acted upon by human consumers. These data are frequently characterized by one or more of the following: high volume, velocity, variety, and veracity. Sandia scientists and researchers are no strangers to large datasets, which often typify the result of high performance computing workloads. The challenge for Sandia data scientists, however, comes at the arrival of exascale computing, which has become mandated by executive order for national computing centers. Because of power and networking limitations, exascale computing precludes the movement of datasets at exascale from data production to data analysis. Such analysis must be done in-situ, to mitigate the time and power costs of network transport. Currently, many of Sandia's data analytics environments exist as independent entities, often specialized for individual use cases, interconnected to other data analytics, high performance computing, and storage systems.

Sandia is procuring, installing, and researching future platforms for such data analytics. A memory footprint of 256 GB per node, Fourteen Data Rate (FDR) InfiniBand networking, PCIe connected non-volatile memory (NVMe) data storage, and an object data store will characterize the coming institutional data analytics platforms. The combined systems will contain 300 compute nodes. A cloud-inspired software stack resulting from collaboration between the Sandia centers will allow the systems to support a flexible mix of user and system-facing workloads and use cases including code development and testing in R and Python, Apache Hadoop and Spark, distributed in-memory databases, graph analytics, emulytics, system log and predictive analytics, and data visualization for institutional and mission customers. These systems will arrive at the end of FY15 and production is planned by the spring of FY16.

Plato, a pilot system for Sandia's institutional data analytics effort, runs primarily Hadoop and Spark.

MISSION COMPUTING SYSTEMS AND SUPPORT

CHAMA

Usage	*Program*	*TFLOPS*	*Process Hours/yr*
HPC	NW/ASC	392	172,677,120
Nodes	***Cores***	***Memory/Core***	
1,232	19,712	2.0 GB	

Chama, along with Pecos, is a NW/ASC HPC system deployed in 2012 as part of the DOE/NNSA Tri-Labs TLCC2, procurement.

CIELO DEL SUR

Usage	*Program*	*TFLOPS*	*Process Hours/yr*
HPC	NW/ASC	86	77,928,960
Nodes	***Cores***	***Memory/Core***	
556	8,896	2. GB	

Cielo del Sur is a Cray XE6 system that supports users of classified computing as part of the National Security Computing Center operated by Sandia on behalf of DOE.

DARK BRIDGE

Usage	*Program*	*TFLOPS*	*Process Hours/yr*
HPC	NW/ASC	294	129,507,840
Nodes	***Cores***	***Memory/Core***	
924	14,784	4.0 GB	

Dark Bridge is a TLCC2 system that, supports users of classified computing as part of the National Security Computing Center operated by Sandia on behalf of DOE.

DARK SAND

Usage	*Program*	*TFLOPS*	*Process Hours/yr*
HPC	NW/ASC	294	129,507,840
Nodes	***Cores***	***Memory/Core***	
924	14,784	4.0 GB	

Dark Sand is a TLCC2 system that supports users of classified computing as part of the National Security Computing Center operated by Sandia on behalf of DOE.

JEMEZ

Usage	*Program*	*TFLOPS*	*Process Hours/yr*
HPC	Institutional	95	40,366,080
Nodes	***Cores***	***Memory/Core***	
288	4,608	2.0 GB	

Jemez is an Institutional HPC system that was released for general use in FY14. At 96TFlops peak, this system is a major and cost-effective addition to the Institutional resources available to users who are doing classified computing.

MINI SEQUOIA

Usage	*Program*	*TFLOPS*	*Process Hours/yr*
HPC	NW/ASC	107	71,761,920
Nodes	***Cores***	***Memory/Core***	
512	8,192	1.0 GB	

Mini Sequoia is a small version of the Tri-Lab Sequoia system sited at LLNL. The purpose of Mini Sequoia is to provide local code-development and checkout system for the application teams that support Sandia users on Sequoia.

MUTRINO

Usage	*Program*	*TFLOPS*	*Process Hours/yr*
HPC	NW/ASC	NA	28,032,000
Nodes	***Cores***	***Memory/Core***	
100	3200	4.0 GB	

Mutrino is a small-scale Cry XC system that is one of two Application Regressions Testbed(ART) systems, supporting the Sandia/LANL ACES partnership Trinity platform located at LANL.

MUZIA

Usage	*Program*	*TFLOPS*	*Process Hours/yr*
HPC	NW/ASC	NA	2,803,200
Nodes	***Cores***	***Memory/Core***	
20	320	2.0 GB	

Muzia is a small-scale Cray XE6 system that was acquired as part of the Sandia/LANL partnership that manages the Cielo platform at LANL.

OPENSTACK

Usage	*Program*	*TFLOPS*	*Process Hours/yr*
Cloud	Institutional	NA	14,016,000
Nodes	***Cores***	***Memory/Core***	
100	1,600	8.0 GB	

OpenStack is an Institutional Cloud system that was acquired at the end of FY13 for use as a research cloud. The system is comprised of 100 Dell r720 nodes having local disk and bound together with a highly configurable Ethernet fabric. An OpenStack cloud environment will support multiple research groups who are experimenting with “Infrastructure as a Service.”

PECOS

Usage	*Program*	*TFLOPS*	*Process Hours/yr*
HPC	NW/ASC	392	172,677,120
Nodes	***Cores***	***Memory/Core***	
1,232	19,712	2.0 GB	

Pecos is a NW/ASC HPC system deployed in 2012 as part of the DOE/NNSA Tri-Labs TLCC2 procurement. Pecos is a primary resource for NW/ASC users.

PLATO

Usage	*Program*	*TFLOPS*	*Process Hours/yr*
Analytics	Institutional	NA	7,148,160
Nodes	***Cores***	***Memory/Core***	
51	816	6.0 GB	

Plato is a Hadoop cluster, which entered production in FY14. The HP-based system runs out-of-the box Cloudera’s CDH enterprise product.

RED MESA

Usage	*Program*	*TFLOPS*	*Process Hours/yr*
HPC	Institutional	180	134,553,600
Nodes	***Cores***	***Memory/Core***	
1,920	15,360	1.5 GB	

Red Mesa is an Institutional HPC system on the collaborative network. Based on the Red Sky platform architecture, Red Mesa is used by EC PMU and other Partners.

RED SKY

Usage	*Program*	*TFLOPS*	*Process Hours/yr*
HPC	Institutional	264	197,835,840
Nodes	***Cores***	***Memory/Core***	
2,823	22,584	1.5 GB	

Deployed in 2010 Red Sky and Red Sky (C) have been the workhorses of Institutional HPC computing. Developed in collaboration with Sun Microsystems, this cluster was the first large-scale HPC system to deliver an Infiniband interconnect based on a Torus network topology.

SKY BRIDGE

Usage	*Program*	*TFLOPS*	*Process Hours/yr*
HPC	Institutional	588	259,015,680
Nodes	***Cores***	***Memory/Core***	
1848	29,568	4.0 GB	

Based on the same hardware as TLCC2, Sky Bridge is the latest Institutional compute cluster. At 588 TF peak, the water-cooled Cray CCS cluster is ranked 123 in the June 2015 TOP500 HPC list. By partnering with Cray and Asetek, Sandia was able to show that a water-cooled HPC cluster can be both reliable and energy efficient. This success has set the stage for future water-cooled HPC clusters at Sandia and elsewhere.

TWINFIN

Usage	*Program*	*TFLOPS*	*Process Hours/yr*
Analytics	Institutional	NA	0
Nodes	***Cores***	***Memory/Core***	
1848	29,568	4.0 GB	

Twinfin is an IBM/Netezza analytics appliance for structured- and semi-structured search. Twinfin came online early in FY13. The system integrates proprietary hardware and software to accelerate structured search integrated with data analytics.

UNO

Usage	*Program*	*TFLOPS*	*Process Hours/yr*
HPC	NW/ASC	58.0608	23,546,880
Nodes	***Cores***	***Memory/Core***	
168	2,688	4.0GB	

Uno is the first high-throughput cluster deployed at Sandia and features three capabilities: 168 traditional compute nodes for a single node jobs, 25 compute+GPU nodes and 8 many core nodes. Based on a Dell compute node, Uno is designed to provide high-throughput and fast turnaround for single-node jobs. Uno provides a variety of heterogeneous nodes (small and large memory, processors and accelerators) with its interconnect and file systems tuned for singe-node activities.

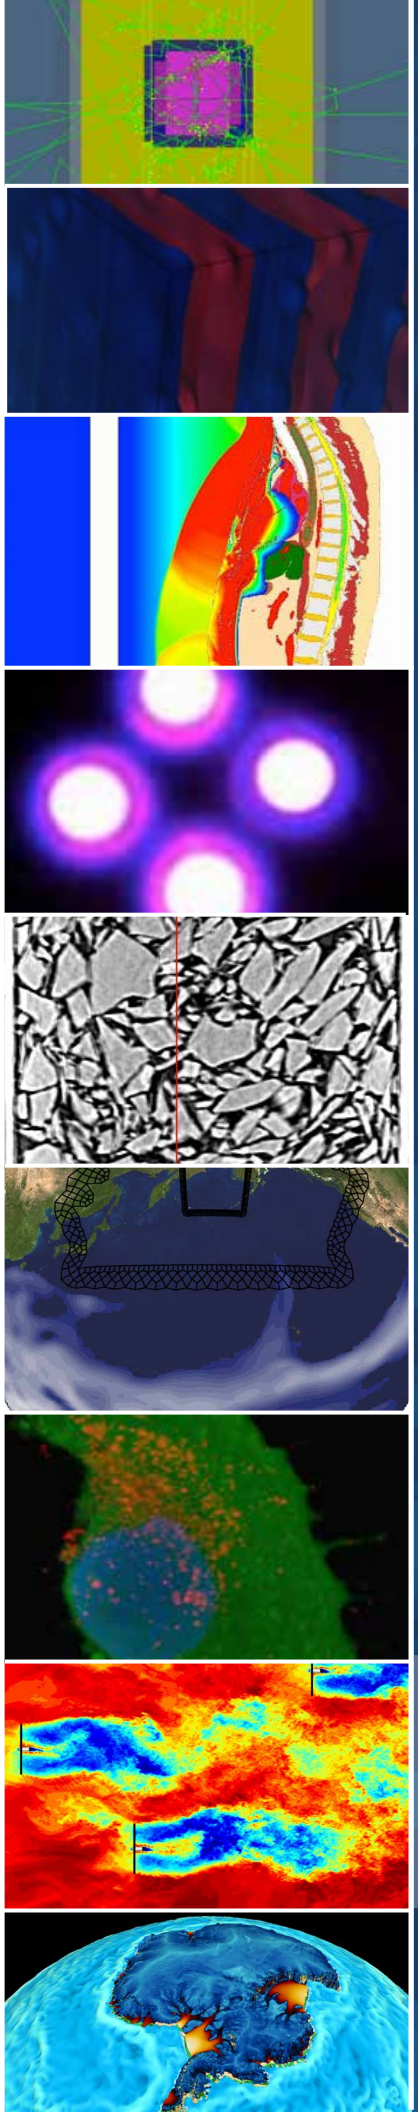

On the Cover: Sky Bridge is the primary unclassified scientific computing platform at Sandia, serving all Mission Areas within the Laboratories. At the formal ribbon cutting in December 2014, Laboratories Director Paul Hommert emphasized the importance of the Institutional Computing Program in providing a critical core Laboratories capability. The deployment of Sky Bridge more than doubled the computing cycles available in Sandia's restricted computing environment, while incorporating the same underlying architecture as the existing ASC/TLCC2 compute clusters, thus allowing many existing computer codes to run on the system without any major modifications.

Sandia National Laboratories

U.S. DEPARTMENT OF ENERGY

National Nuclear Security Administration

Sandia National Laboratories is a multiprogram laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000. SAND2015-8505 R