

BREAKING NEW GROUND

Computational Science at the
Forefront of Discovery

COALITION FOR ACADEMIC SCIENTIFIC COMPUTATION • CASC

2019



ABOUT THE COVER

IMAGE COPYRIGHT CHRISTELLE WAUTHIER

CONTENTS

Tackling Brain Injury	3
When Machines Learn	5
From Computation to Clinic	8
Engineering the Future	11
Fighting Cyberattacks	14
Discovery on Land, Sea and Air	16
Phenomenal Physics	20
Inspiring Students	24
CASC Membership	26



Hawaii's Kilauea volcano captured the world's attention when it began vigorously erupting in May 2018, spewing geysers of fiery rock and sending streams of lava racing down residential streets. The event also captured the attention of scientists, who used advanced instruments and high performance computation to understand what was happening beneath the volcano's tumultuous surface.

This image, produced by Penn State volcanologist Christelle Wauthier, is based on a satellite-mounted observation technique called interferometric synthetic-aperture radar (InSAR). InSAR lets researchers examine changes in ground surface elevation, even tiny shifts of just a few centimeters. The colors indicate how dramatically the ground around Kilauea shifted as magma moved beneath the surface, creating cracks, lava flows and earthquakes around a magma intrusion in the East rift zone in June 2007.

Wauthier uses seismic and InSAR data to infer how magma channels are organized beneath Kilauea and other volcanoes. Because this requires a lot of number-crunching, she relies on supercomputers at Penn State's Institute for CyberScience to combine large data sets and model the inner workings of geologic formations that often span many square miles.

Wauthier and other researchers use these techniques to better understand how and why volcanoes erupt, with an ultimate goal of developing better warning systems to save lives and money. In addition to volcanoes, the work could also yield insights relevant to understanding and detecting earthquakes.

ABOUT CASC

The Coalition for Academic Scientific Computation is an educational nonprofit 501(c)(3) organization with 87 member institutions representing many of the nation's most forward-thinking universities and computing centers. CASC is dedicated to advocating for the use of the most advanced computing technology to accelerate scientific discovery for national competitiveness, global security, and economic success, as well as develop a diverse and well-prepared 21st century workforce. In addition, CASC collaborates with the United Kingdom High Performance Computing Special Interest Group (HPC-SIG) to advance the use of scientific computing across all disciplines

Executive Committee

Sharon Broude Geva, University of Michigan, Chair
Neil Bright, Georgia Institute of Technology, Vice Chair
Craig Stewart, Indiana University, Secretary
Scott Yockel, Harvard University, Treasurer
Rajendra Bose, Columbia University, Past Chair
Lisa Arafune, Director

Communications Committee

Andrew Bell, University of Virginia
Vivian Benton, Pittsburgh Supercomputing Center
Marisa Brazil, Purdue University
Melyssa Fratkin (Chair), Texas Advanced Computing Center
Tom Furlani, University at Buffalo, SUNY
Kristin Lepping, Rutgers University
Dan Meisler, University of Michigan
Paul Redfern, Cornell University
Robert Schoon, Indiana University
Jan Zverina, University of California, San Diego
Lisa Arafune, Director

TACKLING BRAIN INJURY

Concussions and other forms of traumatic brain injury are a major concern in high-impact sports such as football. Although safer football helmets are being developed, there is still much more to learn about how brain injuries occur during sports, and the best way to protect the brain during impact.



IMAGES COPYRIGHT
MATTHEW B. PANZER,
UNIVERSITY OF VIRGINIA

Matthew B. Panzer and his team from the University of Virginia are using computational modeling to build detailed models of football helmets used in the NFL, to better understand how head impacts lead to brain injury. His research could not only help make helmets safer but might also be used to better understand other sources of head injury such as car wrecks.

As part of the NFL's Engineering Roadmap, Panzer's research team, along with teams from the University of Waterloo, KTH Royal Institute of Technology and the Wake Forest University School of Medicine, are building computational models of football helmet designs in current use. Pictured is the University of Virginia's model of the Vicis Zero1 helmet, which uses a highly engineered structure to reduce impact forces.

"We are making these models publicly available so that universities and manufacturers can use them to drive innovation in sports equipment," said Panzer. "The extremely detailed models could be used, for example, to run simulations that test whether a new foam material increases helmet safety without requiring the lengthy process of building an entire helmet model or prototype from scratch."

The project required modeling the helmets on multiple levels. The researchers first tested individual helmet materials such as the foam, the helmet shell and the facemask. They then performed testing of larger-scale components such as the pads and energy-absorbing structures. Finally, all the components of the helmet were assembled and assessed with dummy models under various impact conditions.

"These helmets each feature different energy absorbing systems, such as rubber columns that buckle, or air bladder systems that absorb impact," said Panzer. "Capturing the mechanics of how these complex mitigation systems respond to various types of hard hits occurring at fast speeds and then incorporating that information into the computational models was very challenging."

Validating the models

After using the experimental information to build the computational models, the researchers put the models through a comprehensive validation and evaluation process that compared the models' performance with how real helmets perform in tests conducted with dummies. This final step ensured that the computer models will provide accurate results even when used under simulated conditions that aren't the same as those under which the helmets were tested.

Panzer plans to use the models to identify helmet modifications that could reduce the likelihood of injury. He also wants to incorporate the models into classes to allow students who may not have a comprehensive background in computer modeling to use the model helmets for design work.

Injury at the neuron level

In another project, Panzer is using advanced computational modeling to better understand how a head impact leads to brain injury. "When football helmets are evaluated for safety, the head motion of a dummy is recorded and related to risk of injury," said Panzer. "However, that head motion doesn't directly cause the brain injury. The motion causes deformation in the brain and the brain's neurons, and it is the deformation that actually causes the brain injury."

The researchers are developing computer models that can be used to better understand the relationship between head motion, brain deformation and injury. The detailed brain models that will result from his work could be used together with the helmet models to better understand how to modify the helmets to reduce tissue deformation in the brain.

As a first step, the researchers developed a model that reconstructs axon pathways—the connections between neurons in the brain—based on brain scans from hundreds of medical records. The image shows the reconstructed axon tracks with the different colors indicating the orientation of the axons.

"By embedding those axon tracks within a 3D brain model, we can now look at deformations along the axons that result from a given head impact," said Panzer. "This allows us to more closely study the strains in very specific regions of tissue as related to how that tissue becomes injured. Eventually we want to be able to predict whether a given impact will cause injury."

Although the axon tracks are currently based on an average of several hundred people who received brain scans, the researchers are working at modeling the axons of individual brains. This will allow them to examine how the shape and size of the brain may affect a person's vulnerability to head injury. ■



IMAGE COPYRIGHT PITTSBURGH SUPERCOMPUTING CENTER

WHEN MACHINES LEARN

Machine learning and artificial intelligence (AI) algorithms are opening up vast new capabilities in research and business.

Hold on...a Computer Can Play Texas Hold'Em?

Computer scientists have long pitted algorithms against human opponents in games such as chess or trivia competitions. The field of AI reached a significant new milestone with the crushing defeat of top poker professionals by an artificial intelligence program called Libratus. The program, developed by Carnegie Mellon University (CMU) professor Tuomas Sandholm and PhD student Noam Brown, amassed more than \$1.8 million in chips and beat each of its human opponents over the course of a 20-day poker competition.

It's not only an impressive poker achievement, but a remarkable one for AI. No-Limit Texas Hold'Em involves an astronomical number of decision points. Players must make decisions based on what is in their hand but also make guesses about what opponents are holding and whether they are bluffing.

Libratus tackles the challenge with a three-pronged approach. First, it breaks the vast universe of possible decision points into a more manageable set of problems. It then creates a “blueprint” strategy for the game that includes a higher level of granularity for the early rounds and a rougher outline of the strategy for later ones. Finally, it improves the blueprint strategy over the course of the competition. During the contest Libratus drew its computing power from the Bridges computer at the Pittsburgh Supercomputing Center, a joint program of CMU and the University of Pittsburgh.

The demonstration’s impacts go far beyond gaming. Advances in the ability for AI to make quick decisions based on imperfect information—including purposeful deception—could be useful in real-world applications such as cybersecurity, finance, strategic pricing and military applications. ■

A High-Tech Window to the Past

The Meserve-Kunhardt Collection at Yale University’s Beinecke Rare Book and Manuscript Library offers a rich visual documentation of U.S. history in the earliest days of photography, including material related to Abraham Lincoln and his Civil War contemporaries. But the collection contains tens of thousands of images. How best to explore the historical treasures hidden within?

Thanks to an innovative computational solution, historians and other researchers now have an easy way to peruse the entire collection based on visual similarities between the images it contains. Yale Digital Humanities Laboratory researchers Douglas Duhaime and Peter Leonard used a neural network algorithm to process each image, find similarities among them, and display them in an interactive visualization. The program takes advantage of advanced graphics processing capabilities and multi-dimensional image analysis technology to assess each photograph’s similarity to others in the collection, based on 2,048 different “ways of seeing”—a feat that would be far too complex for human observers to perform.

Users can explore the full collection, zoom in on groupings of related images and select individual items of interest. For example, one “neighborhood” of images contains women in sweeping 19th-century gowns, while another showcases boxers in their distinctive poses. Researchers can even quickly find images in the collection that feature specific components such as buttons or swords.

Techniques for efficiently visualizing thousands of images can be useful for a variety of applications from art to biology. The ability to visually explore the outputs of computer vision models could also potentially help reveal and address social biases that are often hidden within such models. ■



IMAGE COPYRIGHT DOUGLAS DUHAIME AND PETER LEONARD, YALE UNIVERSITY

Learning How to Make Better Materials

Martensite is an extremely hard crystalline structure often formed in the manufacture of steel. The martensite content of steel is crucial to its hardness; too much martensite and the steel becomes brittle, too little and it becomes soft.

To better understand the formation and behavior of martensite, University of Michigan post-doctoral researcher Koki Sagiyama and professor Krishna Garikipati constructed intricate 3D models of its microstructures using computational resources from the Extreme Science and Engineering Discovery Environment (XSEDE) and the National Energy Research Scientific Computing Center. In this image, each color denotes a variant within the martensite microstructure.

Recently, Garikipati has begun combining his numerical modeling acumen with AI to accelerate the design and discovery of new materials. A project conducted under a \$2.4 million grant from the Silicon Valley-based Toyota Research Institute, with additional resources from University of Michigan, focuses on creating better materials for batteries, such as those in electric cars. ■




IMAGE COPYRIGHT
KOKI SAGIYAMA,
UNIVERSITY OF
MICHIGAN





IMAGE COPYRIGHT
KHAN M.
IFTEKHARUDDIN,
VISION LAB,
OLD DOMINION
UNIVERSITY


Boosting Doctors' Ability to Tell Tumor from Healthy Brain Tissues

Gliomas, the most common brain tumor in adults, tend to infiltrate healthy brain tissue rather than growing in a single mass. This makes them extremely difficult to remove surgically, and also complicates diagnosis, monitoring and prognosis.

Researchers led by Khan Iftexharuddin at Old Dominion University's (ODU) Vision Lab, along with clinical collaborators at Children's Hospital of Philadelphia, the University of California San Diego and the University of Iowa, want to give doctors a new weapon in the fight against gliomas. With funding from the National Institute of Biomedical Imaging and Bioengineering and the help of ODU's high performance computing resources, the team developed robust Brain Tumor Segmentation methods designed to help doctors see where tumor tissue ends and healthy brain tissue begins.

The researchers expose their computer program to large collections of brain magnetic resonance imaging (MRI) scans, training the system to find clues in the images that correlate with various health outcomes. After this training, when the program is exposed to new brain images it uses sophisticated computational methods to see patterns and determine the tumor's shape. Complementing the work of clinicians, the program then predicts the patient's prognosis. This image shows how the program processes MRI images (left) to automatically generate an accurate volume segment of the brain tumor (right).

The ODU Segmentation system ranked #1 in an international competition of computer-assisted medical image assessment and patient survival prediction technologies. The researchers hope it can eventually act as a virtual assistant to help doctors diagnose gliomas and determine the best treatment approach and dose. ■



FROM COMPUTATION TO CLINIC

Sophisticated computational approaches accelerate discovery in basic biology and medicine.

Making Safer Heart Medicines

Though effective in most cases, today's heart medications can cause serious side effects. Many of these medicines act on M2 muscarinic acetylcholine receptors (M2 mAChRs), proteins found in the membranes of heart cells, to decrease heart rate and reduce heart contractions. But the genetic sequence of the receptor's primary binding site is found in a variety of places throughout the body, and side effects can occur when heart medicines act on these receptors outside of the heart.

Using a unique computational approach to sample proteins in millisecond time intervals, a research team from the University of California at San Diego (UCSD) and Monash University in Australia identified promising drug leads that may selectively combat heart disease, from arrhythmias to cardiac failure, with a lower risk of unwanted side effects.

In this image the M2 mAChR is shown as orange ribbons. The yellow spheres represent drugs that bind to areas of the receptor similar to those found elsewhere in the body. Purple spheres, on the other hand, bind to regions of the receptor that are much more specific, reducing the likelihood of adverse side effects.

The research team, led by Yinglong Miao and Andrew McCammon at UCSD and Celine Valant and Arthur Christopoulos from Monash University, modeled these various binding sites using supercomputers at the San Diego Supercomputer Center and the Texas Advanced Computing Center. The team performed an unprecedented survey of protein structures and tracked their myriad shapes and conformations in order to zero in on promising targets for tomorrow's heart medicines. ■

IMAGE COPYRIGHT
YINGLONG MIAO,
HOWARD HUGHES
MEDICAL INSTITUTE
AT UC SAN DIEGO




IMAGE COPYRIGHT BRIAN KELLER, THE OHIO STATE UNIVERSITY WEXNER MEDICAL CENTER -
IMAGE IS A COMPOSITE OF IMAGES PUT TOGETHER IN ITS FINAL FORMAT BY BRIAN KELLER

Molecular Dynamics with a Hollywood Flair

To understand how proteins and other cellular components behave and interact, computational biologists such as Jacob Durrant at the University of Pittsburgh often use molecular dynamics simulations. However, these simulations produce massive amounts of detailed data, making it hard to use high-quality 3D rendering techniques such as those common in the film and video game industries.

Seeking the best of both worlds, Durrant's lab developed a program called Pyrite to create leaner versions of these simulations that are more amenable to photorealistic, Hollywood-style 3D visualization. Pyrite works as a plugin to Blender, an open-source 3D modeling program commonly used by scientists. It simplifies simulations by considering the positions of only select atoms and updating those positions less frequently. This limited information is used to estimate the locations of the missing atoms, so less data needs to be stored in memory.

This Pyrite visualization shows a protein complex covered in bushy green hair. In the video version, the hair waves as the protein moves according to a scientifically rigorous simulation. Proteins with

green hair don't exist, but the simulation demonstrates Pyrite's ability to marry simulation and photorealistic computer imagery.

Creating even these simplified Pyrite animations requires considerable computing power. Durrant collaborates with Pitt's Center for Research Computing to speed the rendering process. Seeing Pyrite's potential as a tool for teaching as well as research, Durrant also works with lecturers at Pitt to develop and scale virtual reality tools for teaching structural biology and microbiology. ■

IMAGE COPYRIGHT JACOB
DURRANT, UNIVERSITY OF
PITTSBURGH, DEPARTMENT
OF BIOLOGICAL
SCIENCES





Sequencing to Improve Lung Transplant Outcomes

Why do some lung transplants succeed while others fail? Researchers think microbes such as viruses, bacteria and fungi may play a big role.

Brian Keller, a pulmonologist at The Ohio State University's Wexner Medical Center, uses metagenomics, or the study of microbes as a collective community, to evaluate the relationships between microbes, environmental factors and clinical outcomes in lung disease. Keller and his colleagues created a lung transplant specimen repository along with a comprehensive clinical database that provides researchers and doctors with incredibly detailed information about the microbial communities associated with each specimen. The repository contains 110 patient specimens, and collection is ongoing.

The researchers use genetic sequencing to determine the “virome” of each specimen—that is, the genetic content of all of the viruses present. The researchers analyze this virome data in the context of information about other microbes and the lung tissue itself to gain a deeper understanding of the lung microenvironment. The data is combined with environmental and clinical factors to determine how a lung's microbial community might make it healthier or more prone to disease.

The work will help researchers better understand how lung disease develops and how it affects the functioning of the lungs, with the ultimate goal of improving the outcomes for lung transplant patients. This image is a visual representation of Keller's work, showing the metagenomic sequencing of the DNA of the lung virome (left), leading to an understanding of the various types of viruses and other microbes that cause lung disease. ■

3D Imaging for a Better Picture of Autism

In Samuel Wang's laboratory at the Princeton Neuroscience Institute, researchers examine mouse brains to understand how the cerebellum interacts with the rest of the brain. The team is particularly interested in the cerebellum's role in non-motor functions and its relationship to cognitive development, with a focus on autism.

These images are part of a project led by MD/PhD candidate Tom Pisano that looks at the regional connectivity of the cerebellum to the rest of the brain. Pisano maps connections in the cerebellum using transsynaptic viral tracers—viruses used to trace neural pathways—and whole brain clearing, a specimen preparation process that makes it possible to clearly view high-resolution 3D tissue samples.

The image at left shows a mouse brain cleared and imaged using light-sheet microscopy, a process that uses fluorescence to create high-resolution 3D images. To capture more information in a single image, the researchers downsize the volumetric image, project it in 3D, and map it with color to represent depth.

The image at right uses the same brain clearing technique and shows the protein zebrin (Aldolase C, green areas), which is found in Purkinje cells, the principle cells of the cerebellum. Zebrin gets its name for its striping pattern and the visualization is designed to help researchers understand the brain-to-brain variation of these stripes. By understanding the cerebellum and its role in cognitive functions such as attention and language development, researchers hope to shed light on the role of the cerebellum in autism. ■

IMAGES COPYRIGHT
PRINCETON UNIVERSITY



ENGINEERING THE FUTURE

Computation fuels engineering innovations to improve safety, efficiency and sustainability.




IMAGE COPYRIGHT BRONISLAW GEPNER,
UNIVERSITY OF VIRGINIA

Staying Safe in a Crash

More than 1 million people worldwide are killed in automobile crashes each year. Although new sensing and processing technologies are helping drivers avoid crashes, and are automating some aspects of driving, these advances also create crash scenarios that didn't exist before.

Bronislaw Gepner and colleagues from the University of Virginia are using computer modeling and simulation to understand how new technology and advances in vehicle safety devices affect car safety. The researchers are also working to improve methods used for crash testing, and identify the best ways to protect vulnerable occupants with body sizes outside the ranges for which cars are optimized, such as children, small adults and people who are obese.

The image shows a computational simulation of a National Highway Traffic Safety Administration research crash test in which a deformable barrier moving at 56 mph impacts a passenger vehicle. By including details from the full vehicle, impact barrier, seatbelts, airbags and crash test dummies, the simulation provides important information about how the occupants and vehicle respond to the impact, the likelihood of injuries, and how the vehicle occupants interact with seatbelts and airbags.

The results of this research will be used to inform regulators, researchers and manufacturers of challenges that may arise from the adoption of new car technology, and how this technology could affect the types of crashes that occur and the safety of vehicle occupants. ■





IMAGE
COPYRIGHT
ROHINI BALA
CHANDRAN,
UNIVERSITY OF
MICHIGAN

Using Sunlight to Turn Water and Carbon Dioxide into Fuel

Cleaner and more efficient alternatives to fossil fuels could help slow climate change and reduce air pollution. One such alternative being studied is solar thermochemical fuels, which use the high-quality heat from concentrated sunlight to turn water and carbon dioxide into hydrogen and carbon monoxide that can be used to produce gasoline and other liquid fuels.

IMAGE COPYRIGHT SEYED
HOSSEIN JAMALI



Identifying New Materials to Remove Salt from Water

Today, desalinated water can cost up to ten times more than groundwater. Finding an economical way to remove salt from ocean water could help solve water shortages in many parts of the world.

Li-Chiang Lin and his team from The Ohio State University, along with other collaborators, are trying to identify materials that could reduce the cost of desalinating water. Using molecular simulations performed at the Ohio Supercomputer Center, they identified promising ultrathin-film membranes, such as aluminosilicate nanotubes, two-dimensional hydrocarbon polymers, and zeolite nanosheets (pictured above).

State-of-the-art molecular simulations have allowed the researchers to probe and understand exactly how water molecules travel through a membrane, as well as how salt ions interact with the membrane structure. Zeolite nanosheets, for example, can limit the number of salt ions passing through via well-defined pore structures while easily transporting water because of their ultra-short diffusion paths.

The researchers have conducted a large-scale investigation to explore a number of possible zeolite candidates with diverse structural features to find the ones that offer outstanding salt-trapping efficiency while allowing high levels of water flow. Through this research, they were also able to establish the detailed relationship between structural features and separation performance, which will facilitate the future development of zeolite nanosheet membranes and potentially lead to more efficient technologies for producing desalinated water. ■

Solar fuels look promising if they can be produced on a commercial scale at costs low enough to compete with traditional fuels.

Maximizing the efficiency of solar-to-fuel conversion requires understanding not only the chemistry involved but also the complex heat and mass transfer that takes place. Rohini Bala Chandran and colleagues from the University of Minnesota created a 3D model of a prototype solar reactor that uses cerium dioxide (ceria) as an oxygen transfer agent, to split water and carbon dioxide into hydrogen and carbon monoxide in a nearly isothermal process. Pictured are the temperature distributions in the reactor predicted by the 3D model.

The simulations revealed the rates at which gas to solid reactions in ceria occur in the bench-top reactor at a temperature of 1773 Kelvin (about 2700 degrees Fahrenheit). The analysis also showed the interdependencies among spatial and temporal variations in temperature, chemical concentration and reaction rates, and provided insights into how optical and thermophysical properties, combined with design and operating conditions, affect reactor performance.

Now at the University of Michigan, Bala Chandran and her students are using computational modeling to explore how to use solar thermochemical methods to produce hydrogen from methane, as well as how to produce higher value carbon products such as carbon black, graphene and aromatics. The computational modeling was performed on the high performance computing resources of the Advanced Research Computing at the University of Michigan. ■



IMAGE COPYRIGHT SEAN C. GARRICK,
UNIVERSITY OF MINNESOTA

Boosting Engine Efficiency

Engines powering rockets, aircrafts, and cars all work by breaking liquid fuel into small droplets that are heated, evaporated and burned. Fuels that form the smallest droplets perform best because they heat up rapidly, evaporate sooner and enable efficient, clean combustion. To inform the design of next-generation engines that are clean and efficient, Sean Garrick's team from the University of Minnesota developed new numerical simulation tools that can capture droplet breakup dynamics in their entirety.

Correctly predicting the droplet temperatures and subsequent combustion is challenging because droplets can break up in a variety of ways. The researchers use a unique set of mathematical techniques to efficiently and accurately describe details of the liquid and simulate the complicated dynamics as the liquid flows and breaks up into smaller pieces.

The top image shows a computational fluid dynamics simulation of liquid fuel breakup and heating. The cooled fuel breaks up as hot air flows from left to right. Some small droplets form directly from the large, dome-like structure, while others are produced from the breakup of long threads of fuel. As the droplets break away, they are carried downstream and heated by the flowing air. The small droplets are hot while the larger droplets remain cooler. The bottom images show the time progression for a different droplet breakup mechanism that was simulated with the same numerical methods. The simulations were performed at the Minnesota Supercomputing Institute. ■

FIGHTING CYBERATTACKS

As cybercriminals step up their game, researchers are putting big data and computational power behind cybersecurity research, prevention and response.

IMAGE COPYRIGHT EMILY STERNEMAN, INDIANA UNIVERSITY

“White Hat” Hackers Aim to Beef Up Cloud Security

Cloud computing lets organizations save money on computing hardware, maintenance and IT support, but it also can increase the risk of security breaches. A multi-institutional cybersecurity research team is developing innovative ways to detect cyberattacks on cloud infrastructure in the early stages.

As white hat hackers, or ethical computer security researchers, the team simulates attacks in the cloud using servers in research labs at the University of Arkansas at Pine Bluff and at North Carolina A&T State

University. While they are “hacking” the system, the researchers simultaneously run intrusion detection and prevention applications, helping them understand how real hackers may evade these protections and attempt to hide in everyday network traffic.

Because simulating cyberattacks on campus infrastructure is frowned upon, the team, led by Jessie Walker of the University of Arkansas at Pine Bluff, turned to Chameleon, a computing system hosted by the University of Chicago and the Texas Advanced Computing Center. The researchers used Chameleon to set up virtual machines, then used the virtual machines to carry out their simulated attacks. Walker hopes the work will give organizations in academia, government and industry new ideas on how to use the cloud without exposing themselves to hackers. ■



IMAGE COPYRIGHT EMILY STERNEMAN, INDIANA UNIVERSITY

Research Shows Internet Denial of Service Affects Millions

The first large-scale analysis of internet denial-of-service (DoS) attacks provides eye-opening statistics. Between March 2015 and February 2017, hackers launched more than 20 million DoS attacks targeting about 2.2 million internet addresses. That's about 30,000 attacks per day.

The attacks studied included direct attacks, in which traffic is sent directly to the target by the attackers; and reflection attacks, in which third-party servers are hijacked to carry out the attacks. Even though those are just two of several ways DoS attacks can be carried out, the numbers are a thousand times bigger than previous studies have shown—a staggering result, says lead researcher Alberto Dainotti of the University of California at San Diego and the San Diego Supercomputer Center.

Dainotti and colleagues in The Netherlands and Germany conducted their study using two complementary raw data sources: the UCSD Network Telescope, which captures evidence of attacks involving spoofed addresses; and AmpPot distributed denial-of-

service honeypots, which provide simulated services to lure attackers and capture reflection and amplification of DoS attacks.

More research will help security experts better understand the problem and how to combat it to keep the internet as secure as possible. In the future, Dainotti's team hopes to assess the impact of attacks to understand when and why they actually succeed in taking down networks. They're also studying political attacks similar to those experienced in Egypt and Libya in 2011, when government-ordered internet blackouts were used to control anti-government protests. ■



IMAGE COPYRIGHT EMILY STERNEMAN, INDIANA UNIVERSITY




IMAGE COPYRIGHT
SCOTT PEARSE

DISCOVERY ON LAND, SEA AND AIR

Big data gives scientists new tools to understand our dynamic planet.

Simulating Winds to Save Millions in Energy Costs

Wind is a fast-growing source of clean, renewable energy. But because wind speeds vary over time and terrain, wind turbines require constant load balancing, which is the process of storing excess energy in strong winds for use when the wind is weaker. In addition, understanding wind flows and wind wakes—the slowing of wind speeds caused by turbulence as winds move over large, stationary objects—promises to make wind energy cheaper and more efficient.

Pedro Jimenez of the National Center for Atmospheric Research uses a model known as the Weather Research Forecast to predict how changing winds may impact wind turbines. This image, produced using the model, shows a wind wake off the leeward side of Mt. Hood east of Portland, Ore. Although Oregon has a wind generating power capacity of more than 3 Gigawatts (enough to power more than 2 million homes for a

year), wind wakes are a common problem in the state's mountainous areas.

As shown in dark blue, the wind slows as it passes over the mountain, creating a wind wake. Farther from the mountain, the wake dissipates and wind speeds pick up again. As the wake of slow wind passes through wind generators, it can effectively turn the entire turbine system off and back on over the course of hours.

By improving power plant operators' ability to predict when wind wakes will happen and balance energy loads appropriately, models like the Weather Research Forecast can potentially help save hundreds of millions of dollars in energy costs. ■


A 3D Map of Earth's Tremors and Shakes

The ability to visualize where earthquakes happen and how bad they are helps researchers pinpoint areas with the most seismic activity and understand where humans and infrastructure are most at risk. It can also inform building and development strategies in earthquake prone areas.

The problem is that earthquakes generally occur along fault lines, which makes visualizing their quantity and magnitude over time challenging. With so many events concentrated in the same geographic areas, data depicting earthquakes that are older or smaller can become obscured by more recent or bigger ones.

Eliot Feibush of Princeton University sought to overcome this problem with the 3D visualization tool ParaView, which he uses to show the more than 234,000 earthquakes in the U.S. Geological Survey earthquake database, dating back to 1900. The 1,300 strongest earthquakes (magnitudes 7.0 to 9.5) are shown as triangles, while quakes in the 6.0 to 7.0 range are depicted as color-coded vertical lines. Narrower, translucent lines show smaller earthquakes (magnitude 5.0 to 6.0). The longer the line, the heavier the concentration of earthquakes recorded in a given spot.

To create the visualization, Feibush grouped the raw data into “bins” containing data with similar values. He then integrated that information with terrain imagery from NASA's Visible Earth Blue Marble collection and ocean depth data from the National Oceanic and Atmospheric Administration's Geophysical Fluid Dynamics Laboratory. The result is a single, global visualization that helps researchers, planners and residents in earthquake prone areas mine more than a century of data and see how their risk compares. ■



Looking Below a River's Surface to Create Better Engineering Practices

It's a real drag—and potentially a waste of infrastructure investments—when engineers dig or dredge a river only to have it fill in with sediment months later. The phenomenon, known as the Bulle Effect, occurs at points in a river where an existing channel splits into two and the amount of sediment entering the two channels differs from the amount of water flowing into each channel. Because the water at a diversion point takes on a corkscrew-like flow, the Bulle Effect causes sediment in the lower part of the water column to move into the new lateral channel, even when more of the water continues to flow along the main channel.

This visualization of research conducted by Som Dutta, Paul Fischer and Marcelo H. Garcia of the University of Illinois at Urbana-Champaign illustrates the difference in sediment flow at different depths of a river channel at diversions. Most of the sediment traveling near the bottom of the channel moves into the lateral channel of the river, even though the water flowing into the channel is proportionally smaller.

The detailed visualization was generated using the Blue Waters supercomputer at the University of Illinois' National Center for Supercomputing Applications. The work adds to the fundamental understanding of an important natural phenomenon, and provides insights that could help in the design of future engineered channels used for irrigation or navigation, and to restore deltas that have lost land because of rising sea levels. ■



IMAGE COPYRIGHT JENNA PRESTON AND ELIOTT FOUST, NCAR

IMAGE COPYRIGHT MATT REHME, NCAR



Putting Climate Science in the Palm of Your Hand

Many climate visualizations are designed by and for atmospheric researchers. The National Center for Atmospheric Research (NCAR) aims to break that mold by putting sophisticated climate visualizations within the reach of anyone with enough curiosity to use them.

For example, the image at left shows an augmented reality application that showcases Earth and geoscience visualizations by overlaying 3D virtual objects onto the physical world, as viewed through a head-mounted device or even an iPad. The app, created by Arizona State University doctoral student Nihanth Cherukuru during a summer internship at NCAR, takes advantage of the camera on a personal device. When the camera is pointed at an image from a visualization—in this case a computer model of monthly surface temperature anomalies—the visualization pops up onto a 3D globe that can be spun around with a finger. By making such visualizations accessible in an augmented reality environment, NCAR hopes to make science more engaging and interactive, particularly for younger audiences.

The image at right shows another climate-related visualization that could potentially be viewed in the augmented reality environment. The visualization was created using the Community Earth System Model (CESM) and shows a typical present-day weather pattern. The colorful swirls represent surface wind speeds in the “roaring 40s,” the region of the southern hemisphere between 40 and 50 degrees latitude where strong winds are unhindered by land masses. As winds increase, the colors change from blue to yellow to red. An NCAR research team led by Susan Bates uses the CESM to simulate future climate scenarios and to study weather phenomena and climate on a global scale. Studying climate and weather with augmented reality is one way to make climate data sets less abstract and more understandable for students and the general public. ■




IMAGE COPYRIGHT SOM DUTTA, UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN





IMAGE COPYRIGHT
TEXAS ADVANCED
COMPUTING CENTER,
UNIVERSITY OF TEXAS
AT AUSTIN

Transporting Heat in a Cold Ocean

Although water that is just a couple degrees above freezing may not seem warm, a little bit of warmth means a lot from the standpoint of Arctic sea ice and the Greenland ice sheet. A team at the University of Texas at Austin, led by An Nguyen and Patrick Heimbach, uses computer simulations to study heat transport into the Arctic and how it relates to mass loss of the Greenland ice sheet and the decline in the summer Arctic sea ice cover.

This simulation shows how warm water of subtropical origin is carried by the Norwegian Atlantic boundary current through the Nordic and Barents Seas and into the Arctic Ocean. The visualization looks north from the Nordic Seas toward the Eastern Arctic. It shows the Atlantic water boundary current transporting heat toward the Eastern Arctic and eddies mixing heat from the steep slopes into the Arctic Ocean interior.

Reds, yellows and greens depict the warmest waters (1 to 2 degrees Celsius) while light and dark blues represent cooler waters (-1 to -2 degrees Celsius). The area depicted is about 5,200 by 3,800 kilometers with depths ranging from 0 to 800 meters. Images made from this data reveal eddies of length scales of less than 15 kilometer radius. The simulation's level of detail required more than 1,000 hours of high performance computer time to produce. In the future, stronger winds could increase turbulent vertical mixing in the water, erode the in-between cold buffering water layer, and potentially enable this Atlantic water heat source to reach the surface and further melt sea ice. ■



PHENOMENAL PHYSICS

Computing power helps scientists explore physical phenomena from the miniscule to the massive at home and in outer space.

Mapping Weather on Mars

Mars is plagued by sudden and sometimes extreme dust storms such as the one that completely engulfed the planet in the summer of 2018. This extraterrestrial extreme weather forced the team operating NASA's Opportunity rover on Mars to shift the rover to minimal operations to save power. Dust storms not only block sun from reaching the solar panels that power rovers, but can also be detrimental for future vehicles scheduled to land on the planet.

A research team led by Steven Greybush of Penn State has created hourly weather maps depicting winds, temperatures and pressures on the red planet over the past decade. This information will help scientists better understand how planet-wide dust storms evolve from local storms, and provide the ability to track the traveling

weather systems (pictured here) that give rise to those storms. The maps are created by the Ensemble Mars Atmosphere Reanalysis System (EMARS) tool, which combines temperature and dust information obtained from orbiting spacecraft with numerical weather prediction computer simulations.

EMARS could help NASA keep its spacecraft safe and help with planning future missions. It can also be used by scientists to explore the predictability of dust storms on Mars and even to study weather on Earth, which has regional traveling weather systems that can resemble those on Mars. ■



IMAGE COPYRIGHT MICHAEL
CRONIN AND CLAUDE-
ANDRÉ FAUCHER-GIGUÈRE,
NORTHWESTERN UNIVERSITY

Using Simulation to Understand the Origin of the Universe


Scientists aren't sure exactly how galaxies—which contain hundreds of billions of stars held together by gravity—formed after the Big Bang 14 billion years ago. Did small particles group together to gradually form stars and star clusters that turned into galaxies, or did the universe start out as enormous clumps that later separated into galaxies?

To find out how the universe developed the structure it has today, researchers led by Claude-André Faucher-Giguère of Northwestern University are using simulations to study galaxy formation. Each image, or snapshot, in a galaxy simulation can contain millions of particles, making it difficult for scientists to explore and make discoveries using these complex data sets. To

help solve this problem, Faucher-Giguère and his team developed a 3D app called FIRExplorer that allows the user to visualize and fly through simulation snapshots.

With the FIRExplorer app, simulation data can be loaded directly without any preprocessing. Researchers around the world can then view very large simulations of galaxy formations in real time on any device. The FIRExplorer snapshot from a cosmological galaxy simulation pictured here consists of approximately 21 million particles representing stars, gas and dark matter in a young galaxy.

In addition to helping researchers understand our universe's structure, FIRExplorer can be used to study components of the universe such as dark matter and dark energy as well as how stars and black holes form. ■



IMAGES COPYRIGHT YOUSRA NAHAS
AND SERGEI PROKHORENKO,
UNIVERSITY OF ARKANSAS

Large-Scale Simulation of Thin Film Physics

Lead zirconate titanate, a manmade ceramic material commonly known as PZT, exhibits several unusual properties. Not only does it change shape when an electric field is applied, it also becomes polarized when exposed to an electric current, and remains polarized when the electric current is removed. These unique properties make the material useful for forming ultrasound images, and for sensors that measure changes in pressure, acceleration, temperature, strain, or force, by converting the changes to an electrical charge.

Yousra Nahas, Sergei Prokhorenko and Laurent Bellaiche of University of Arkansas are using complex computational and simulation methods within the DARPA Topological Excitations in Electronics program to study the physics of PZT thin films and their reaction to various external parameter changes.

PZT has a dipolar structure, meaning that it contains unit cells with opposite electric charges, or polarity, that are close together. Using a large-scale Hamiltonian simulation developed by Bellaiche's group at the University of Arkansas, the researchers studied how the film's unit cells form different dipolar patterns in response to temperature and external field changes.

In the image on the right, the simulation revealed a labyrinth-like pattern when the film temperature was decreased to 10 Kelvin. On the left, a quasi-hexagonal lattice structure of bubbles is formed when the film temperature was dropped to 10 Kelvin under an electric field. These dipolar patterns correspond to ferroelectric domains and have promising technological properties. Grey indicates dipoles pointing upward and burgundy indicates those pointing downward.

Probing self-organization in ferroelectrics such as PZT could reveal universal behaviors among disparate systems and be used in bottom-top design of novel devices including new kinds of electrical circuits and components of artificial neural networks. ■




IMAGE COPYRIGHT CHRISTOPHER N. SHINGLEDECKER,
UNIVERSITY OF VIRGINIA

Space Chemistry Could Reveal How Life Began

Researchers are looking to space to understand how life originated. They are particularly interested in studying the chemical evolution of interstellar objects because it could provide clues as to how the organic molecules necessary for life originally formed.

The interstellar medium is found in the space between star systems. It contains gases, dust, cosmic rays and radiation. Christopher N. Shingledecker and colleagues from the University of Virginia are investigating how cosmic rays—which are predominantly made of high energy protons—can cause physical and chemical changes when they collide with ice covering interstellar dust grains. The researchers developed a new model called the Chemistry of Ionizing Radiation in Solids (CIRIS) that is providing the first microscopic scale simulations of these types of collisions.

Pictured here is an illustration of how a bombardment of an interstellar dust grain might appear. The red track depicts the cosmic ray and the electrons formed when the cosmic ray collides with molecules in the ice. The researchers generated the track using Rivanna, the University of Virginia's high performance cluster, to run the computationally intensive CIRIS model.




IMAGE COPYRIGHT HECTOR GOMEZ, PURDUE UNIVERSITY

Visualizing Order Within Chaos

Fluid and unstable phenomena such as turbulence are extremely complex, with billions of atoms swirling in seeming chaos. Researchers led by Hector Gomez at Purdue University are working to get a better handle on that complexity by using a large-scale numerical simulation to visualize the solution to the Kuramoto-Sivashinsky equation in 3D (pictured). This important equation is used to study instabilities that occur during chaotic phenomena.

The simulation helped the researchers learn more about the role that geometry and dimensionality play in how fluids transition from order to chaos. The new computational method could be useful for studying various chaotic phenomena including aerodynamics and turbulent fluid flow and inform the design of more efficient aircraft and ships. ■

CIRIS can be used to study a variety of physical and chemical changes that occur in an irradiated solid, which will provide new insight into the chemistry of interstellar objects exposed to cosmic radiation as well as planet-like objects such as large moons or asteroids exposed to solar wind. ■

High Schoolers Create Autonomous Vehicles in Summer Program

Last summer, high school students from Central Texas built and programmed autonomous vehicles as part of the Texas Advanced Computing Center (TACC) CODE@TACC Robotics summer program. CODE@TACC, now in its fourth year, gives students from schools with limited resources exposure to coding and STEM careers.

With support from the KLE Foundation, each student in the robotics summer program received a Raspberry Pi computing system that could be embedded in a small autonomous vehicle. The students learned advanced programming concepts, including how to train a neural network to make decisions based on input from sensors. At end of the program the participants each presented a solution to an autonomous driving problem.

In addition to the robotics program, CODE@TACC also offers summer programs focused on developing computational solutions to societal problems; creating Internet of Things-connected objects; and exploring internet security and cryptography. The programs all include hands-on interactive instruction and career guidance from TACC researchers, panel presentations from undergraduates and professionals, and industry site visits to increase awareness about careers with connections to computing.

According to Dawn Hunter, a senior program coordinator at TACC, the CODE@TACC programs not only expose and train students in technology, but also allow them to explore life's opportunities by providing a safe and encouraging space to ask questions and excel. ■

IMAGE COPYRIGHT TEXAS ADVANCED
COMPUTING CENTER, THE UNIVERSITY
OF TEXAS AT AUSTIN

INSPIRING STUDENTS

Exposing diverse students to computation equips the next generation for careers at the cutting edge.





IMAGE COPYRIGHT XSEDE, NATIONAL CENTER FOR SUPERCOMPUTING APPLICATIONS

Creating a More Diverse and Inclusive STEM Workforce

The U.S. Bureau of Labor Statistics projects that by 2020, one million high-tech jobs will remain unfilled because of a lack of qualified U.S. college graduates. Creating a more diverse and inclusive STEM workforce is vital for filling these future high-tech jobs.

The National Science Foundation's Extreme Science and Engineering Discovery Environment (XSEDE), a virtual organization that integrates and coordinates sharing of advanced digital services, is working to create a more diverse and inclusive STEM workforce through its Broadening Participation program. This program helps underrepresented minorities, women, and faculty from Minority Serving Institutions (MSI) join the XSEDE community so that they can benefit from its resources and services.

The Broadening Participation program supports the research and teaching goals of individual scientists and also collaborates with institutions to promote broad adoption of computational science and data analytics in their research. The program exhibits at conferences, visits campuses and conducts regional workshops to create awareness and increase participation by underrepresented minorities and MSI faculty in the XSEDE community. In addition, the XSEDE Diversity Forum connects outreach and diversity managers at each XSEDE partner to facilitate sharing of best practices, program opportunities and reviewing XSEDE services and programs to ensure that they are inclusive.

Since 2011, over 3,200 women and underrepresented minorities have gained access to an XSEDE computer resource for the first time. In addition, users from Minority Serving Institutions report some of the highest levels of satisfaction across the full spectrum of XSEDE services. ■



IMAGE COPYRIGHT ROBERT GOTWALS, NORTH CAROLINA SCHOOL OF SCIENCE AND MATHEMATICS

High School Chemistry Students Gain Experience with Supercomputers

What do teenagers, fabric dyes, solar energy and supercomputers have in common? They were all recently the focus of a unique educational experience designed to advance energy research while sparking young people's interest in STEM. High school chemistry students at the North Carolina School of Science and Mathematics (NCSSM) explored the suitability of fabric dyes for solar energy applications, using the Bridges supercomputer at the Pittsburgh Supercomputing Center (PSC) of Carnegie Mellon University and the University of Pittsburgh. The students gained important experience that will prepare them for computationally intensive jobs in fields like biology, engineering and the social sciences.

NCSSM computational chemistry teacher and North Carolina State University (NCSU) visiting scholar Robert Gotwals worked with PSC's Marcela Madrid and Elena Jakubikova at NCSU to develop research opportunities for NCSSM residential and online students using NSF funding and access to the Bridges supercomputer. Gotwals and his students used Bridges to explore computational methods to screen a small subset of dyes known as anthroquinones from a library of 90,000 dyes housed at NCSU. The research question was to see if any of the dyes exhibited light absorption and electronic properties that could make them useful components in solar power cells.

By using a real supercomputer, the students in the course gained essential skills such as how to manage file systems when resources are limited. Findings from the project will be published in a scientific journal, giving the students exposure to writing a scientific paper and the peer-review process. NCSSM has the largest computational sciences high school program in the United States, with computational courses in biology, chemistry, physics, drug-design chemistry, nanoscience, digital humanities and research. ■

CASC MEMBERSHIP

**Arizona State University
Research Computing**
Tempe, Arizona

Boston University
Boston, Massachusetts

**Brown University
Center for Computation and
Visualization**
Providence, Rhode Island

**Carnegie Mellon University &
University of Pittsburgh
Pittsburgh Supercomputing
Center**
Pittsburgh, Pennsylvania

**Case Western Reserve
University
Core Facility Advanced
Research Computing**
Cleveland, Ohio

**City University of New York
High Performance
Computing Center**
Staten Island, New York

**Clemson University
Computing and Information
Technology (CCIT)**
Clemson, South Carolina

Columbia University
New York, New York

**Cornell University
Center for Advanced
Computing**
Ithaca, New York

**Georgetown University
Research Technologies**
Washington, District of Columbia

**Georgia Institute of
Technology
Partnership for an Advanced
Computing Environment
(PACE)**
Atlanta, Georgia

Harvard University
Cambridge, Massachusetts

**Icahn School of Medicine at
Mount Sinai**
New York, New York

**Indiana University
Pervasive Technology
Institute**
Bloomington, Indiana

Johns Hopkins University
Baltimore, Maryland

**Lawrence Berkeley National
Laboratory**
Berkeley, California

**Louisiana State University
Center for Computation &
Technology (CCT)**
Baton Rouge, Louisiana

**Michigan State University
High Performance
Computing Center**
East Lansing, Michigan

**Michigan Technological
University**
Houghton, Michigan

**Mississippi State University
High Performance
Computing Collaboratory
(HPC2)**
Starkville, Mississippi

**Montana State University
Information Technology
Center**
Bozeman, Montana

**National Center for
Atmospheric Research
(NCAR)**
Boulder, Colorado

New York Genome Center
New York, New York

New York University
New York, New York

**North Dakota State
University
Center for Computationally
Assisted Science &
Technology**
Fargo, North Dakota

Northwestern University
Evanston, Illinois

NYU Langone Medical Center
New York, New York

**Oak Ridge National
Laboratory (ORNL)
Center for Computational
Sciences**
Oak Ridge, Tennessee

**Oklahoma State University
High Performance
Computing Center**
Stillwater, Oklahoma

Old Dominion University
Norfolk, Virginia

**Pennsylvania State
University
Institute for CyberScience**
University Park, Pennsylvania

Princeton University
Princeton, New Jersey

Purdue University
West Lafayette, Indiana

**Rensselaer Polytechnic
Institute**
Troy, New York

**Rice University
Ken Kennedy Institute**
Houston, Texas

Rutgers University
Piscataway, New Jersey

**Stanford University
Stanford Research
Computing Center**
Stanford, California

**Stony Brook University
Research Technologies**
Stony Brook, New York

**Texas A&M University
Institute for Scientific
Computation**
College Station, Texas

**Texas Tech University
High Performance
Computing Center**
Lubbock, Texas

**The George Washington
University**
Washington, District of Columbia

**The Ohio State University
Ohio Supercomputer Center
(OSCC)**
Columbus, Ohio

**The University of Alabama at
Birmingham
Research Computing
Services**
Birmingham, Alabama

**The University of Texas at
Austin
Texas Advanced Computing
Center**
Austin, Texas

**University at Buffalo, State
University of New York
Center for Computational
Research**
Buffalo New York

**University of Alaska
Fairbanks
Research Computing
Systems**
Fairbanks, Alaska

**University of Arizona
Research Technologies**
Tucson, Arizona

**University of Arkansas
High Performance
Computing Center**
Fayetteville, Arkansas

**University of California,
Berkeley
Berkeley Research
Computing**
Berkeley, California

**University of California,
Irvine
Research
Cyberinfrastructure Center**
Irvine, California

**University of California, Los
Angeles
Institute for Digital Research
and Education**
Los Angeles, California

**University of California, San
Diego
San Diego Supercomputer
Center (SDSC)**
San Diego, California

**University of Chicago
& Argonne National
Laboratory
Research Computing Center**
Chicago, Illinois

**University of Colorado
Boulder
Research Computing**
Boulder, Colorado

**University of Connecticut
Booth Engineering Center
for Advanced Technology
(BECAT)**
Storrs, Connecticut

University of Florida
Gainesville, Florida

**University of Georgia
Georgia Advanced
Computing Resource Center
(GACRC)**
Athens, Georgia

**University of Illinois at
Urbana-Champaign
National Center for
Supercomputing
Applications (NCSA)**
Champaign, Illinois

University of Iowa
Iowa City, Iowa

**University of Kentucky
Center for Computational
Sciences**
Lexington, Kentucky

University of Louisville
Louisville, Kentucky

**University of Maryland
Division of Information
Technology**
College Park, Maryland

University of Massachusetts
Shrewsbury, Massachusetts

**University of Michigan
Advanced Research
Computing (ARC)**
Ann Arbor, Michigan

**University of Minnesota
Minnesota Supercomputing
Institute for Advanced
Computational Research**
Minneapolis, Minnesota

**University of Nebraska
Holland Computing Center**
Omaha, Nebraska

**University of Nevada, Las
Vegas
National Supercomputing
Institute (NSI)**
Las Vegas, Nevada

**University of Nevada, Reno
Research Computing**
Reno, NV

**University of New
Hampshire
Research Computing Center**
Durham, New Hampshire

**University of New Mexico
Center for Advanced
Research Computing**
Albuquerque, New Mexico

**University of North Carolina
at Chapel Hill
Renaissance Computing
Institute (RENCI)**
Chapel Hill, North Carolina

**University of North Carolina
at Chapel Hill**
Chapel Hill, North Carolina

**University of Notre Dame
Center for Research
Computing**
Notre Dame, Indiana

**University of Oklahoma
OU Supercomputing Center
for Education and Research**
Norman, Oklahoma

**University of Oregon
Research Advanced
Computing Services (RACS)**
Eugene, Oregon

**University of Pittsburgh
Center for Research
Computing**
Pittsburgh, Pennsylvania

University of Rhode Island
Kingston, Rhode Island

**University of South Florida
Research Computing**
Tampa, Florida

**University of Southern
California
Information Sciences
Institute**
Marina del Rey, California

**University of Tennessee
National Institute for
Computational Sciences
(NICS)**
Knoxville, Tennessee

**University of Utah
Center for High Performance
Computing**
Salt Lake City, Utah

**University of Virginia
Advanced Research
Computing Services (ARCS)**
Charlottesville, Virginia

**University of Wyoming
Advanced Research
Computing Center (ARCC)**
Laramie, Wyoming

**Vanderbilt University
Advanced Computing Center
for Research and Education**
Nashville, Tennessee

**Virginia Tech
Advanced Research
Computing**
Blacksburg, Virginia

West Virginia University
Morgantown, West Virginia

**Yale University
Yale Center for Research
Computing (YCRC)**
New Haven, Connecticut

**Coalition for Academic
Scientific Computation**

Lisa Arafune, Director

626C Admiral Drive, Suite 530
Annapolis, Maryland 21401
(202) 930-CASC (2272)
casc.org
info@casc.org

