



2018

**U.S. Computational
and Data Science:
Driving Economic Growth
and Discovery**



Coalition for Academic Scientific Computation

About The Cover

Pictured is a computer simulation of the process used to extract water from wood pulp during paper manufacturing. Researchers created the simulation in an effort to find new ways to cut the energy and water required to make paper. If manufacturers can reduce energy use by 20 percent, it would save the industry up to \$400 million annually. See page 3 for more on how computation is being used to improve American manufacturing. Image courtesy of David Trebotich, Lawrence Berkeley National Laboratory.

About CASC

The Coalition for Academic Scientific Computation is an educational nonprofit 501(c)(3) organization with 84 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, and to support economic and workforce development in high performance computing-related fields.

CASC Executive Committee

Rajendra Bose (Chair), Columbia University
 Sharon Broude Geva (Vice Chair), University of Michigan
 Neil Bright (Secretary), Georgia Institute of Technology
 Andrew Sherman (Treasurer), Yale University
 Curtis W. Hillegas (Past Chair), Princeton University
 Lisa Arafune (Director)

Communications Committee

Dustin Atkins, Clemson University
 Andrew Bell, University of Virginia
 Vivian Benton, Pittsburgh Supercomputing Center
 Marisa Brazil, Purdue University
 Sarah Engel, Indiana University
 Melyssa Fratkin, Texas Advanced Computing Center
 Tom Furlani, University at Buffalo, SUNY
 Karen Green, Renaissance Computing Institute
 Dan Meisler (Chair), University of Michigan
 Paul Redfern, Cornell University
 Jan Zverina, University of California, San Diego
 Lisa Arafune, CASC Director

Writing: Anne Frances Johnson and Kathleen Pierce, Creative Science Writing
 Design: Bullhaus - Durham, NC (Heide Randall, lead designer)

Contents

MANUFACTURING A STRONG ECONOMY 3

MAPPING RESOURCES AND RISKS 6

HEALTH MATTERS 9

TRAINING FOR TOMORROW 13

UNVEILING THE UNIVERSE 15


CASC MEMBERSHIP 18



Manufacturing a Strong Economy

Although the role of computers may be more obvious in the tech industry, computation is just as essential in other key drivers of the American economy, such as manufacturing, medicine and energy. Many industries rely on computation and computer simulations to solve problems, streamline operations and spark innovation. Staying at the forefront of computation fuels growth and helps strengthen the economy now and for years to come.

These images model the microscale dynamics of water as it moves through felt pressers used to extract water from wood pulp. Improving the pressing process could yield large efficiency gains in paper manufacturing.



A Public-Private Partnership to Propel Papermaking

Since people have been making paper for more than 2,000 years, you might think we would have the process perfected by now. Yet scientists at the Department of Energy's (DOE) national laboratories are applying high-tech advanced computational science to an emblem of low-tech—paper—in an effort to make the papermaking process significantly more efficient.

Among manufacturers, the papermaking industry is the country's third-largest energy user behind petroleum refining and chemical production. In a unique collaboration

between government labs and a consortium of paper manufacturers, researchers are using high performance computing to help reduce the amount of energy and water needed to make paper.

The research is led by scientists at Lawrence Livermore National Laboratory and Lawrence Berkeley National Laboratory, who used data supplied by paper manufacturers to design computer models of the pressing and drying process used to remove water from wood pulp, a

Continued

crucial (and energy-hungry) step in papermaking. The project is designed to lead to practical improvements that can be implemented on-the-ground in manufacturing facilities in the near term.

“We successfully used pore-scale simulation of flow in the paper press from microscale image data to develop a more accurate engineering scale model of the wet press process,” explains David Trebotich, a computational scientist at Lawrence Berkeley National Laboratory who works on the project. “This project is a very good example of high performance computation for manufacturing. Pore-scale simulation from microstructure image data requires large-scale computing. And the results obtained from these simulations have been successfully integrated in an engineering scale model.”

The team’s goal is to save the industry money by making the pulp and paper-drying process more energy efficient. If the paper’s dryness were increased by 10-15 percent after the pressing stage,

researchers estimate that the energy used in the subsequent drying step could be reduced by up to 20 percent. That, in turn, could save as much as \$400 million annually across the U.S. papermaking industry.

The research is one of many projects fueled by a DOE program designed to help businesses take advantage of the nation’s supercomputers and expertise. DOE hosts five of the top twelve most powerful computers in the world. The High Performance Computing for Manufacturing (HPC4Mfg) program gives industry partners access to these state-of-the-art computational resources for projects geared toward boosting domestic manufacturing and improving U.S. competitiveness in the global marketplace.

“Industrial partners do not always have, or are not always able to allocate, resources for fundamental research as these efforts typically occur on longer time scales than industry can accommodate,” Trebotich says. “DOE

national laboratories in particular are able to engage in this level of research. The HPC4Mfg program leveraged existing computational efforts and facilities at labs in Berkeley and Livermore to develop a usable modeling tool for an industrial manufacturing purpose.”

HPC4Mfg has helped address a wide range of challenges since the program’s launch in 2015, with a broad focus on boosting productivity, improving energy efficiency, and developing new energy technologies. In addition to the collaboration with America’s paper manufacturers, the program has supported advances in automotive and aerospace manufacturing, metals and materials, semiconductors and electronics, transportation and shipping, lighting, energy generation and other industries.

“Giving industrial partners access to DOE computational resources ultimately helps fulfill the DOE mission of delivering energy solutions to the world,” Trebotich says. •

Gearing Up for Smarter Testing


When a manufacturer needs to know if a part from a car or airplane is healthy or defective, getting an answer typically requires a battery of expensive, time-consuming physical tests. Computer modeling offers a better way. Many companies are developing sophisticated software to automate such inspections for faster and more reliable results. But a major hitch is that running these models requires a substantial amount of computing power—more than most smaller businesses can provide.

A project to link business needs with university resources aims to close that gap. The Ohio-based company Advanced Numerical

Solutions was awarded a NASA grant to determine the feasibility of moving its models to a cluster computing system to take better advantage of existing computational infrastructure.

To test pilot the approach, the company partnered with The Ohio State University to run analyses on the Ohio Supercomputer Center’s Oakley Cluster.



The result was a valuable proof-of-concept illustrating the benefits simulations can bring to industry at many scales. Advanced Numerical Solutions can use the outcomes to determine where to focus the company’s technology investments for the greatest return. Pictured here is a bull gear (part of a gear box), as analyzed by the company’s powerful modeling software. •



A Laser Focus on Light-Based Innovations

Modern lasers capable of emitting extremely concentrated light beams are an essential tool used in hundreds of industrial processes, consumer products, medical treatments and research applications. They've found their way into the operating room and onto the battlefield; they are essential in precision manufacturing facilities and in the barcode scanners at your local supermarket. Despite their many uses, we still lack a fundamental understanding of how lasers transform objects at the atomic level.

A team led by Leonid Zhigilei of the University of Virginia performs simulations to answer this essential question.



Using the powerful Rivanna computer, the team models the behavior of millions of atoms simultaneously, providing an extreme close-up of the complex interactions that occur when materials interact with laser light.

Focusing first on metals, the research could lead to improvements in the use of lasers for hardening materials and manufacturing metal products. Ultimately, the team plans to expand the simulations to investigate lasers' effects on a variety of other materials, including nanomaterials and biological tissue. •

Envisioning a New Energy Future With Fusion

Ever wonder how stars generate the tremendous amount of energy it takes to shine across galaxies for billions of years? The answer is a powerful reaction known as fusion. In one of today's most ambitious energy research projects, scientists from 35 countries are working to build the world's largest thermonuclear fusion plant, International Thermonuclear Experimental Reactor, known as ITER ("the way"). The experimental facility, now under construction in southern France, could pave the way for an energy revolution in the coming decades.


ITER will generate energy using a tokamak, a doughnut-shaped, magnetized vacuum chamber designed to confine plasma, a state of matter that tends to expand. Confined plasma creates a fusion reaction, generating extreme heat in the walls of the device. This heat can then be harnessed to produce electricity. ITER is designed to hold nearly 10 times more plasma than the current largest experimental fusion reactor and will provide an important proof-of-concept for commercializing fusion technology.

Plasma expert Wendell Horton of the University of Texas at Austin is

contributing to the effort by modeling the dynamics that occur within the tokamak. Shown here is a simulation of the electromagnetic fields (visualized as colorful repeating patterns) that are generated in the lining of the magnetic device as the plasma inside creates fusion energy. The visualization, created by Greg Foss of the Texas Advanced Computing Center, shows a tokamak's doughnut-shaped vessel collapsed into a rectangular solid. The white

contours at the front show magnetic surfaces confining the plasma in the core (left) and exhaust chamber (right) plasmas.

These models provide a test bed for validating the tokamak's design parameters with the aim of creating a safe, efficient and vastly more powerful power plant. •



Mapping Resources and Risks

High performance computing is invaluable for improving our understanding of—and ability to predict—large-scale natural phenomena. As sensors become more powerful and prevalent, researchers are complementing fine-grained real-world measurements with computer-based simulations to decipher the behavior of land, water and air. This research has gained increased urgency as expanding populations expose more people and infrastructure to hazards such as earthquakes, hurricanes and landslides.


Big Data for Big Maps

Nevada's Black Rock Desert Wilderness is just what you would imagine based on its name—a vast, bone-dry span of desert scrub stretching between dark-hued peaks. However, as the snowpack melts each spring, the Quinn River emerges across this arid landscape. The river flows through deposits from ancient Lake Lahonton, a giant, high-elevation lake that reached its greatest depth during the last Ice Age, about 14,000 years ago. The unique region is a potential analog to the surface of Mars, where the echoes of ancient rivers are seen in the planet's Aeolis Dorsa region.

This image of the Quinn River was generated with LIDAR, a light-based remote sensing technology used to examine the Earth's surface. The LIDAR data was collected by the National Science Foundation (NSF)-supported National Center for Airborne Laser Mapping and used by University of Virginia doctoral researcher Yo Matsubara. In the map, the unmistakable swoops and swirls of water become the most prevalent feature of the now-dry landscape, potentially offering insights into how similar features on Mars may have developed.

Matsubara's project is just one of the many applications of LIDAR-based high-resolution topography data made available via OpenTopography, a platform for the distribution and processing of these massive datasets. Supported by the NSF, OpenTopography is based at the San Diego Supercomputer Center (SDSC) and is a collaboration between the University of California, San Diego, Arizona State University and UNAVCO, a non-profit research consortium.

"By leveraging the high performance computing resources available at SDSC, we can integrate resource-intensive algorithms to enable more sophisticated data processing and analysis," says Viswanath Nandigam, Principal Investigator for the OpenTopography facility. "OpenTopography acts as a gateway that makes NSF-funded XSEDE high performance computing



resources more accessible to the scientific community by significantly reducing complexity and barriers to their use.” The Extreme Science and Engineering Development Environment (XSEDE) is an NSF-funded virtual organization that integrates and coordinates the sharing of advanced digital services—including supercomputers and high-end visualization and data analysis resources—with researchers nationally to support science.

With more than a trillion LIDAR-based measurements of the Earth’s surface available from OpenTopography, providing efficient access to these data for the research community is a Big Data challenge. OpenTopography

utilizes large-scale data management and high performance computing to provide efficient web-based access and processing tools for large, high-resolution topographic datasets.

By providing ways to share and use large Earth science data sets, OpenTopography expands access to rich, incredibly detailed models of topography and bathymetry. “The power of OpenTopography is that it provides simple access to what are otherwise quite challenging datasets to process and analyze,” says Christopher Crosby who manages day-to-day operations at OpenTopography. “This ease of access to high-resolution topography leads to considerable data reuse, innovative new

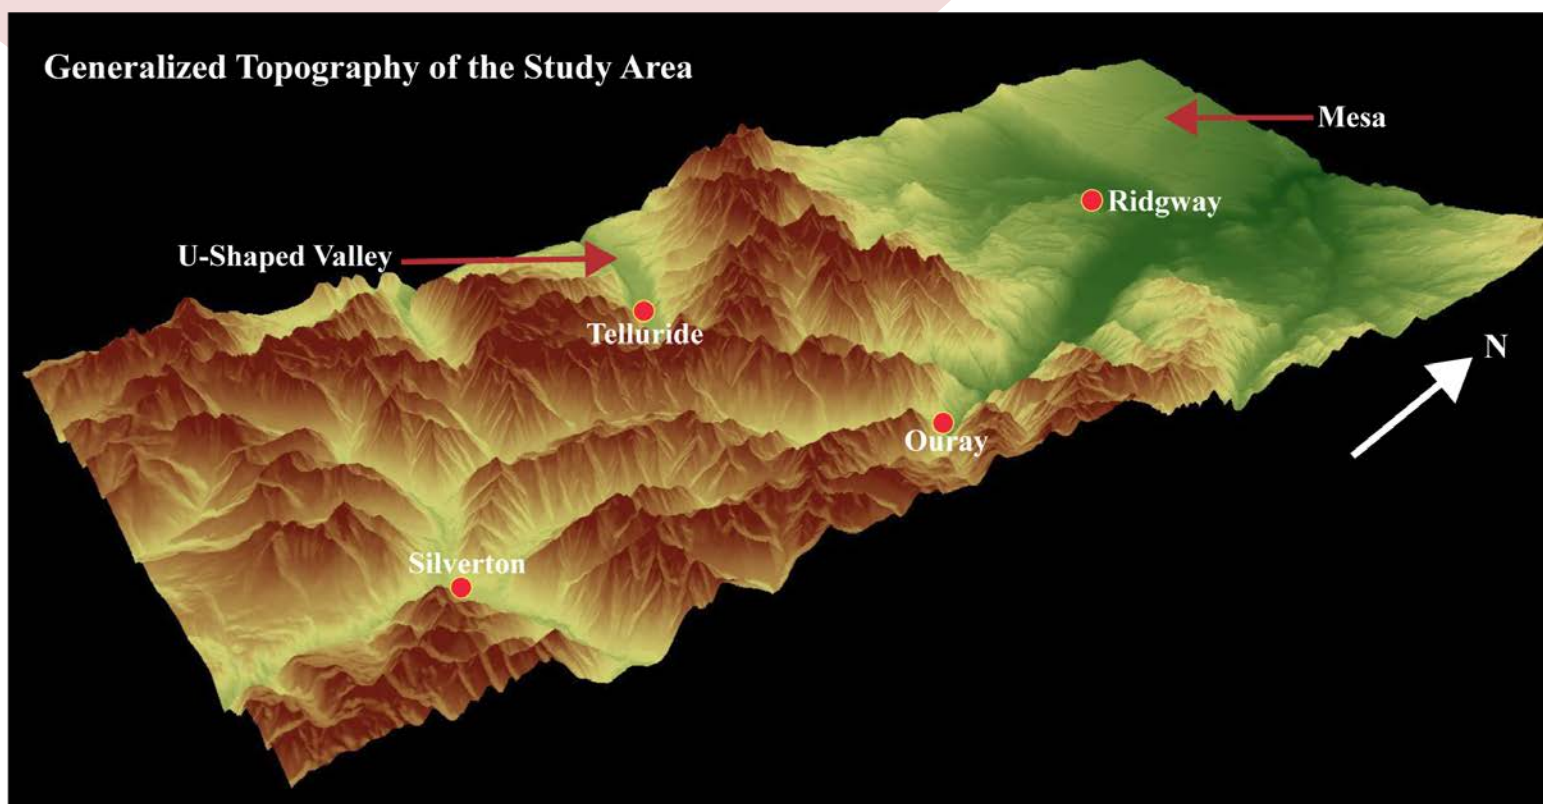
applications for data collected initially for other research topics and greater return on investment for funders who invest in the initial data collection.”

These maps support a wide range of applications, from forestry and ecological conservation to disaster resilience. For example, researchers can study the area around Kaikoura, New Zealand as it looked before and after the 2016 earthquake, track the changes occurring in the McMurdo Dry Valleys of Antarctica or analyze floodplains in Teton National Park. OpenTopography has enabled nearly 300 peer-reviewed publications including academic works in Earth science, ecology, hydrology, geospatial science, engineering and computer science. ●

Managing Terrain Both Tenacious and Tenuous

Visitors to Colorado’s San Juan mountain range enjoy a breathtaking variety of terrain, from rugged bare rock to gentle, forested hills and valleys. But while the landscape is beautiful to look at, it is also highly vulnerable to mass movements such as landslides, avalanches and rockfalls. In recent years, development and climate-related changes have compounded this natural fragility. With a rapidly growing resident population and increased tourism, more lives and infrastructure are now in the path of danger.

To help the region’s decision makers better understand and manage risks, researcher Kaytan Kelkar of Texas A&M University used the Immersive Visualization Center at Texas A&M to create minutely detailed digital elevation models of the area, shown here. By visualizing the area’s complex geomorphology at a fine scale, these sophisticated models help reveal areas prone to landslides or other dangers. Decision makers can use that information to predict hazards, promote best land use practices and minimize risk. ●




When the Earth Shakes

For anyone living on the U.S. West Coast, earthquakes are a fact of life. While it remains impossible to predict exactly when an earthquake will strike, studying quakes over the years has dramatically improved our ability to predict the likelihood and

expected impacts of future earthquakes, as well as helped to inform building codes, disaster plans and recovery efforts.

With a new software system, known as EDGE (Extreme-Scale Discontinuous Galerkin Environment), researchers can study earthquake mechanisms in a virtual laboratory. By enabling the fastest seismic simulations to date, EDGE allows seismologists to perform more detailed computer simulations than ever thanks to efficient utilization of the latest and largest supercomputers. The system was developed by SDSC researcher Alex Breuer and colleagues as part of a public-private partnership between Intel Corporation and SDSC at the University of California San Diego.


This image, created using EDGE, shows a simulated seismic wave in the San Jacinto fault zone between Anza and Borrego Springs in California. Colors denote the amplitude of the particle velocity, with warmer colors corresponding to higher amplitudes. ●



Keeping a Steady Eye on Unsteady Weather


Unsteady vortices are natural weather phenomena that occur in tropical storms near the equator and cold weather patterns like the polar vortex. The complex dynamics of these vortices make it extremely difficult to precisely predict the size and track of a storm.

Researchers Robert Krasny and Ling Xu of the University of Michigan (UM) are studying the behavior of unsteady vortices in an effort to improve weather forecasts and potentially save lives. Using a high performance computing cluster at UM's Advanced Research Computing facility, the team generated this visualization employing leading-edge numerical methods to carry out accurate and efficient simulations of unsteady vortices. The colors in the image show that over time, vortices remain compact and tighter in the center, but more dispersed and unpredictable at the edges. The study suggests advancements in numerical algorithms and high performance computing resources can enhance our understanding and ability to predict unsteady vortex dynamics. ●



Health Matters

Computation has long been an essential tool for understanding the inner workings of our bodies. The ability to collect fine-scale data on the genes, structures and interactions that drive living things has increased dramatically in recent years. But harnessing all that data to generate meaningful advances is a significant challenge, leading to a growing focus on finding smart ways to transform big data into new medical tools to fight pernicious problems such as cancer and infectious disease.



The Power of Prediction

Proteins are the workhorses of living cells, carrying out the essential functions that keep life ticking. When it comes to proteins, shape matters—a lot. A protein's shape is what determines how it interacts with other cell components and anything else it may encounter, such as invading viruses or bacteria. As a result, a great deal of modern biology focuses on deciphering the structures of proteins and how their shape affects their function.

One particularly promising area of research is protein design—the art of architecting and manufacturing proteins to accomplish difficult tasks, for example degrading harmful chemicals in the soil or repairing injured cells. However,

protein structures can be exceedingly complex, and their seemingly infinite variety makes it challenging to puzzle out the key facets that make a protein work.

It is against this backdrop that Daisuke Kihara and his team at Purdue University developed sophisticated software to predict protein structures and behavior. The software includes a protein docking program that can model protein complex structures and a virtual drug screening program that can identify potential drugs for a target protein.

Continued

“Solving protein structures by experiment is time-consuming and expensive,” explains Kihara. “Accurate computational methods for predicting protein structures and complex structures can provide protein structure information quickly and with much less cost.”

The software tool has proven extremely effective in solving real-world problems, earning the team recognition in three prestigious research competitions. First, the software was able to accurately predict the 3D structure of a protein (previous page, left image), given only its amino acid sequence, in a contest sponsored by the U.S. National Institute of General Medical Sciences. Second, the program was used to predict how two or more proteins bind to create protein complexes (previous page, right image) in an international competition sponsored by the U.S. National Institutes of Health (NIH)

and the European Union. Finally, the team earned national recognition in a challenge to predict protein function in a contest sponsored by the NSF.


Kihara hopes that these demonstrations will inspire many more projects to use protein predictions to advance discovery and find new applications. For example, protein structure predictions can also dramatically reduce the timeline for developing new medicines. ●

Resisting Antibacterial Resistance

One of today’s most alarming public health problems is the emergence of so-called “superbugs”—infectious bacteria that are resistant to current antibacterial drugs. To combat this problem, scientists are designing antibodies that can go after new targets within the invading bacteria.

Researcher James Gumbart at the Georgia Institute of Technology created this simulation, which reveals how an antibody can pass through the outer layer of the bacterial membrane and bind to a specific protein at the cell surface. The image shows an antibody (at top) targeting a protein (blue, green, and red) that is buried within the outer layer (gold and gray) of *Salmonella*.

By helping researchers tinker with different variables and understand precisely what mechanisms are required for antibodies to recognize and bind to targets, simulations like these are crucial tools in the quest to create more effective weapons in the war against superbugs. ●





Cancer Mappers

“Cancer metastasis” may be one of the most dreaded phrases ever uttered and also one of the most studied biological phenomena in human history. But scientists still do not have a complete understanding of how this

complex, multistep process occurs. Researcher Abdul N. Malmi-Kakkada and colleagues with the Thirumalai Group at the University of Texas at Austin theorized that studying the physical signatures of cancer cell

movements during metastasis could help us understand, and one day counteract, this deadly process. He used the modeling resources at the Texas Advanced Computing Center (TACC) to create highly detailed maps of tumor cell proliferation.



A model of tumor cell proliferation, created by Anne Bowen of TACC, shows the different velocities of individual cells within a tumor. Black arrows indicate the direction of movement of cells in the full (left) and bisected (right) view. While there are cell-to-cell variations, there are also some clear patterns. Malmi-Kakkada and collaborators discovered, for example, that while the inner cells are basically static and move in random directions, cells on the outside move much faster in a directed manner thereby playing an important role in propelling tumor growth and metastasis. ●


Getting Brainy with BRAIN-I

High-resolution brain imaging is a powerful tool to help scientists understand the brain and its ailments. But today’s astonishingly detailed imaging technologies also create an overwhelming amount of data, making it difficult for researchers to share their data and replicate studies.

Researcher Jason Stein at the University of North Carolina at Chapel Hill (UNC) ran into this very problem when he started studying high-resolution images of 3D mouse brain slices, 1 micron thick, with the goal of understanding how genome variations affect brain development and potentially cause neuropsychiatric illnesses.

Mike Conway and Ashok Krishnamurthy of UNC’s Renaissance Computing Institute (RENCI) stepped in with the solution. Through a collaboration known as BRAIN-I, Stein’s lab is using RENCI’s cross-disciplinary cyberinfrastructure to host the images and tools required to properly analyze large stores of brain imaging data, transforming it from a confusing jumble to a rich source of impactful discoveries. By shifting

data from an impediment to an opportunity, well-designed cyberinfrastructure allows scientists to devote their time to asking—and answering—big questions about the brain. ●

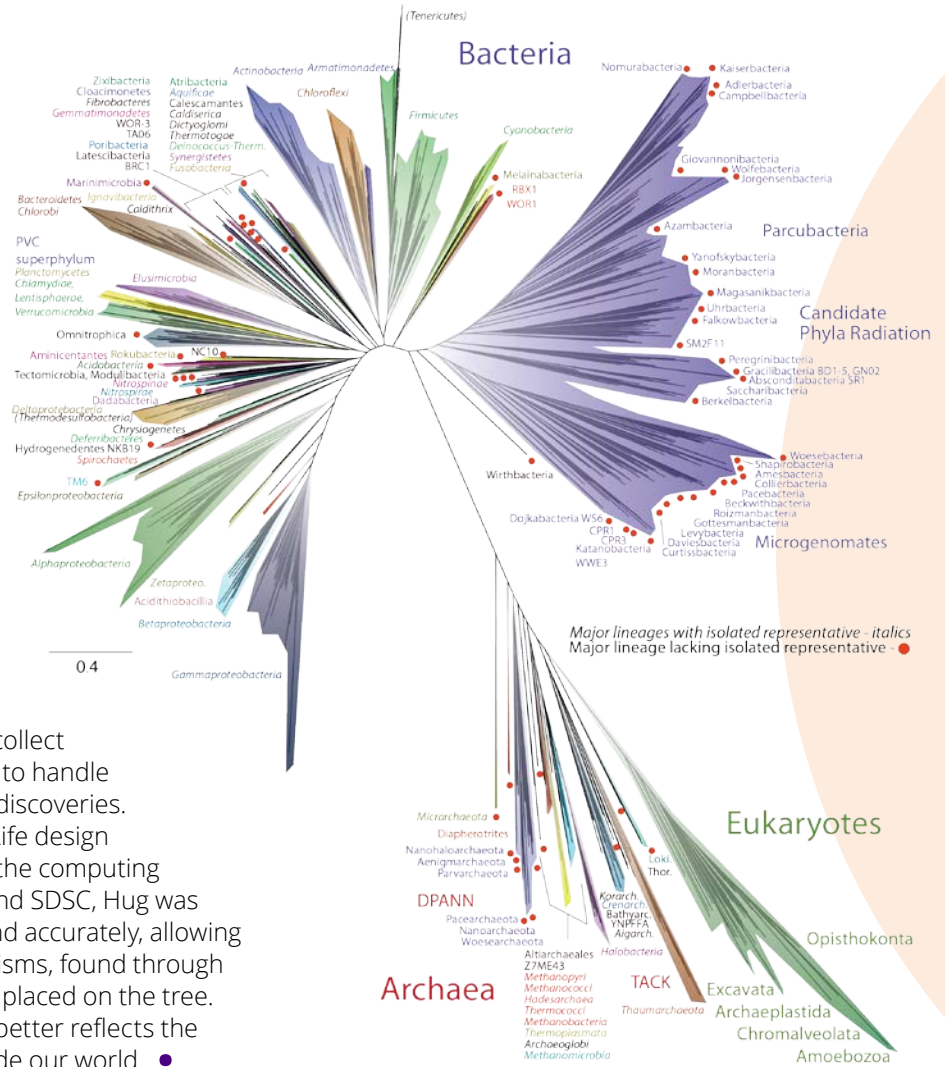


A New Tree of Life


Our planet is teeming with life, yet the organisms we can see with our eyes are only a fraction of the total in existence. Tiny microorganisms are crucial to the health of our bodies and to all ecosystems on Earth.

As we learn more about microbes, we gain an increased appreciation for their diversity and complex history. To deepen our understanding of how microbes have evolved and the diversity they contribute to life on Earth, Laura Hug of the University of Waterloo and Jill Banfield at the University of California, Berkeley developed an updated Tree of Life, which captures the relationships among more than 3,000 organisms.

While it is now possible to determine the DNA sequences of almost any organism, the ability to collect this data often outpaces the technology required to handle such large data sets and mine them for scientific discoveries. Indeed, previous attempts at a microbial Tree of Life design failed from a lack of computing power. Thanks to the computing resources of the NSF's CIPRES Science Gateway and SDSC, Hug was able to infer the characteristics of trees quickly and accurately, allowing her team to discern where newly identified organisms, found through DNA sequencing rather than culturing, should be placed on the tree. The result is a new depiction of life on Earth that better reflects the remarkable—though tiny—organisms that pervade our world. •



With Modern-Day Dowsing Rod, Finding Water in DNA



More than six decades after scientists determined the basic structure of DNA, there is still so much about this remarkable “molecule of life” that we don’t know. With the advent of gene editing technologies, gene therapies and other cutting-edge techniques, scientists have been working with greater urgency to pin down the intricate details of DNA in order to improve our ability to predict and manipulate its dynamics.

Researchers have long known that within the minor groove of a DNA strand is a “spine of hydration,” a line of water molecules that adhere to the strands and provides stability to the double-helix structure. With computational resources from the University of Notre Dame’s Center for Research Computing,

researchers Steven Corcelli and Kristina Davis performed calculations and visualizations to support experiments that successfully measured the vibrational spectrum of the minor groove water molecules. This illustration shows the experimental setup, with the line of water molecules depicted in orange and white.

Now that the spectrum of the spine of hydration can be measured, scientists can study its properties with much more detail and learn how changes within it can alter DNA functioning when foreign bodies, such as biomolecules or drugs, are introduced. For example, water molecules have been thought to prevent drug binding, but it may be possible to find a novel method for binding that bypasses the water or counteracts its effects. •

Training for Tomorrow

As big data grows ever bigger, it also grows ever more integral to almost every industry. For America to stay on top, it is now clear that an increasingly large number of workers need the skills and knowledge to take advantage of the tools and methodologies of data science. Generating—and sustaining—a data-savvy workforce means both educating the next generation and training today’s mid-career professionals in the latest developments and tools.



Distilling Data Science

Whether you’re in manufacturing, healthcare, tech, finance or any number of other fields, data matters. That’s probably why Data Matters™, a short course series from the National Consortium for Data Science, the Renaissance Computing Institute and the Odum Institute at the University of North Carolina at Chapel Hill and North Carolina State University, has grown so popular, so fast. The courses attract students and faculty, as well as working professionals, seeking the timely skills they need to analyze and manage large, heterogeneous data sets.

Recent courses have covered a wide range of topics, including the latest in programming languages, data visualization, machine learning, geospatial analytics, data mining, and data sharing and reuse. Organizers aim for the series to be equally valuable for students just starting out in the field of data science and seasoned professionals who want to bring the power of data to their organization’s processes and decisions. ●

Mining a New National Treasure

Stampede2 is the largest supercomputer ever to be housed at a U.S. university and one of the most powerful systems in the world. When it came online in 2017 at the Texas Advanced Computing Center at the University of Texas at Austin, the new computer opened exciting new doors for researchers in Texas and across the country.

Made possible through a \$30 million NSF grant, the supercomputer is designed to serve thousands of U.S. scientists and engineers to improve the competitiveness of the U.S. workforce and ensure America remains a leader in computational research. In this image, Cornell University Center for Advanced Computing staff lead a one-day training workshop to help researchers take advantage of this potent new resource. ●



Cultivating Young Innovators

Each year, dozens of bright, tech-interested students walk the halls of SDSC at the University of California San Diego. They come for workshops, camps and internships—but most of all, they come for the opportunity to create their own futures. These students know that tech will play a huge role in the economy and their careers in the decades ahead. SDSC offers robust education and experiences for today's students in order to inspire and equip tomorrow's workforce.

One particularly popular program is the Research Experience for High School Students, which pairs students with SDSC mentors to gain experience working on a computation-centric project. In 2017 the program supported research experiences for 66 students selected from more than 200 applicants. Their work ranged from mapping the molecular mechanisms of disease to learning how to monitor and analyze computer networks and software.

Over its eight-year history, the program has inspired many students who have since gone on to pursue study at many eminent academic institutions. ●




Unveiling the Universe


When it comes to big data (or anything else, for that matter), you can't get much bigger than the universe. That's what makes high performance computation essential to almost every area of astronomy and astrophysics. This research can lead to findings with many practical applications here on Earth, such as improvements in GPS technology, in addition to revealing the wonders of space.

A Gravitational Dance

Plumbing the infinite depths of black holes can enhance our understanding of the universe at its most basic level. Sourav Chatterjee and Aaron Geller of Northwestern University are particularly interested in solving how two black holes can come so close to each other that they eventually merge: Were they close to each other "by birth," or did they somehow "find each other" billions of years later? The question arose recently when the merging of two black holes was detected by the NSF-funded Laser Interferometer Gravitational-Wave Observatory (LIGO).

To address the question, Geller developed this visualization of the latter scenario, described as a "gravitational dance" involving three black holes and one normal star. The visualization shows how the enormous gravity coming from the black holes bends the starlight. As a result of the gravitational interactions depicted here, two black holes would come so close to each other that they would merge, and the other black hole and the star would be shot outward into space at a high velocity. In addition to advancing astronomy, studying black holes can also lead to insights about gravity, galaxy formation and general relativity. ●





Grouping Galaxies

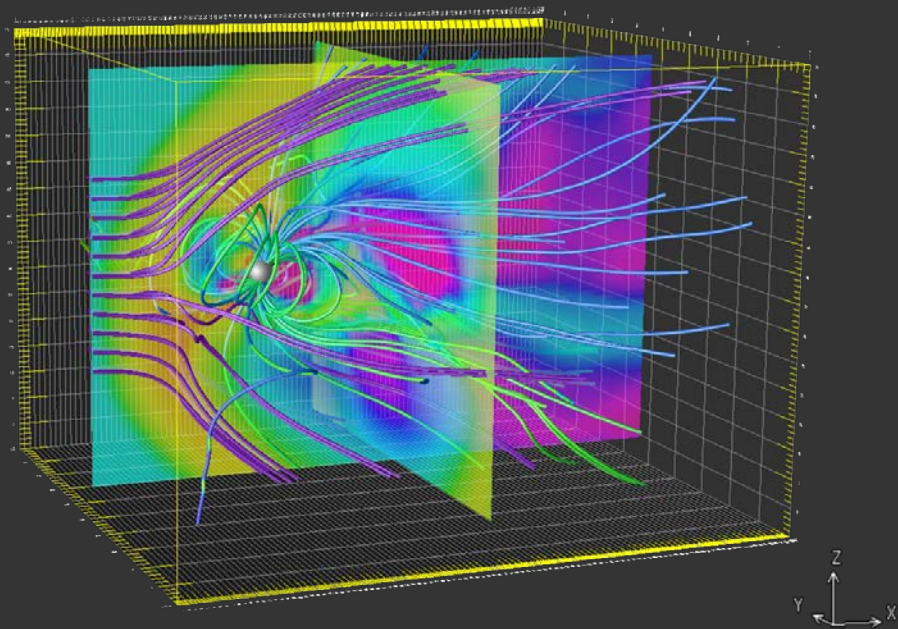
Galaxy clusters—swirling masses of hundreds or thousands of galaxies—are the largest bodies in our universe. By tracking their size and number and better understanding how they grow, scientists hope to learn more about dark matter and dark energy, which together make up roughly 95% of the universe. Yale University researchers developed this galaxy cluster simulation to decipher how gas behaves when clusters form. The multi-colored lines represent complex turbulent gas flows, which gain speed as new galaxies merge into an existing cluster. The varying colors represent temperatures, which are higher in the core, where the gravitational pull is more powerful.

Researchers led by Daisuke Nagai ran the experiments, Luis Fernando created the visualization, and Kaylea Nelson and Erwin Lau ran the computations. ●

Understanding Uranus

Uranus, the seventh planet from our Sun, holds many mysteries. How was it formed? Why does it tilt asymmetrically? What's under all that dense cloud cover? Studying these questions can enhance our understanding not just of Uranus, but help answer more fundamental questions about how our solar system was formed, how it continues to evolve and whether any of our neighboring planets could harbor life.

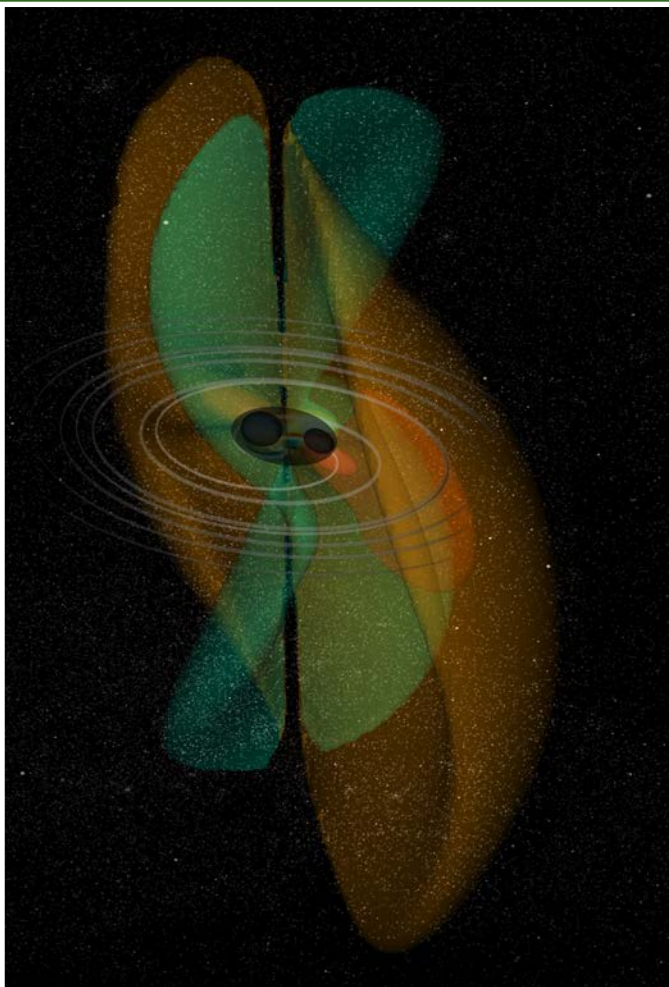

Researchers Xin Cao and Carol Paty of the Georgia Institute of Technology created this simulation of how Uranus' magnetosphere interacts with the local solar wind and seasonal variations. The purple lines represent the flow of solar wind, which diverts around the magnetosphere. The blue and green lines are magnetic field lines with a component pointing either into (blue) or away from (green) the solar wind. ●



In Search of Planet Nine

While Pluto was reclassified as a dwarf planet in 2006, there may yet be another “Planet Nine” hiding somewhere in the far reaches of our solar system, waiting for discovery. To expand on a theory of this mystery planet’s existence, Yale University researchers Sarah Millholland and Gregory Laughlin set out to calculate where the planet would be expected to orbit—if it indeed exists.

The team used more than 250,000 CPU hours at the Yale Center for Research and Computing to compute complex calculations based on a few basic assumptions about the potential planet. This image shows the results: the team’s best guess at Planet Nine’s orbit, its interactions with nearby Kuiper Belt objects, and how well each simulation matches known observations. The researchers released a manipulatable 3-D version of the model in hopes that other scientific teams can use the data to build on the theory and perhaps someday find a previously undiscovered planet. ●



Ripples in Spacetime

When researchers at the NSF-funded LIGO observatory finally observed gravitational waves 100 years after Einstein’s predicted “ripples in the spacetime fabric of the universe,” they relied on a combination of experimental observations and numerical simulations to know it was the real deal. One of the simulations involved was created by Deirdre Shoemaker, Pablo Laguna and colleagues at the Georgia Institute of Technology. With the XSEDE computers at the Texas Advanced Computing Center, the team ran hundreds of simulations using Einstein’s equations to predict the behavior of a ripple of energy produced when two black holes collide. Black hole collisions are the most energetic events in our universe, emitting tremendous bursts of gravitational radiation. The image shows a simulation of one such collision. ●

This research is just one example of the contributions many CASC member institutions made to the observation of gravitational waves, a remarkable scientific achievement that was recognized with a Nobel Prize in Physics in 2017.

CASC Membership

Arizona State University
ASU 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 (PSC)
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

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 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

Princeton University
Princeton, New Jersey

Purdue University
West Lafayette, Indiana

Rensselaer Polytechnic Institute
Troy, New York

Rice University
Ken Kennedy Institute for
Information Technology (K2I)
Houston, Texas

Rutgers University
Piscataway, New Jersey

Stanford University
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 (OSC)
Columbus, Ohio

The Pennsylvania State University
University Park, Pennsylvania

The University of Texas at Austin
Texas Advanced Computing Center (TACC)
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 Computing
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
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 Houston
Center for Advanced Computing
and Data Systems
Houston, Texas

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
Minneapolis, Minnesota

University of Nebraska
Holland Computing Center
Lincoln, Nebraska

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

University of Nevada, Reno
Research Computing
Reno, Nevada

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 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
Los Angeles, 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

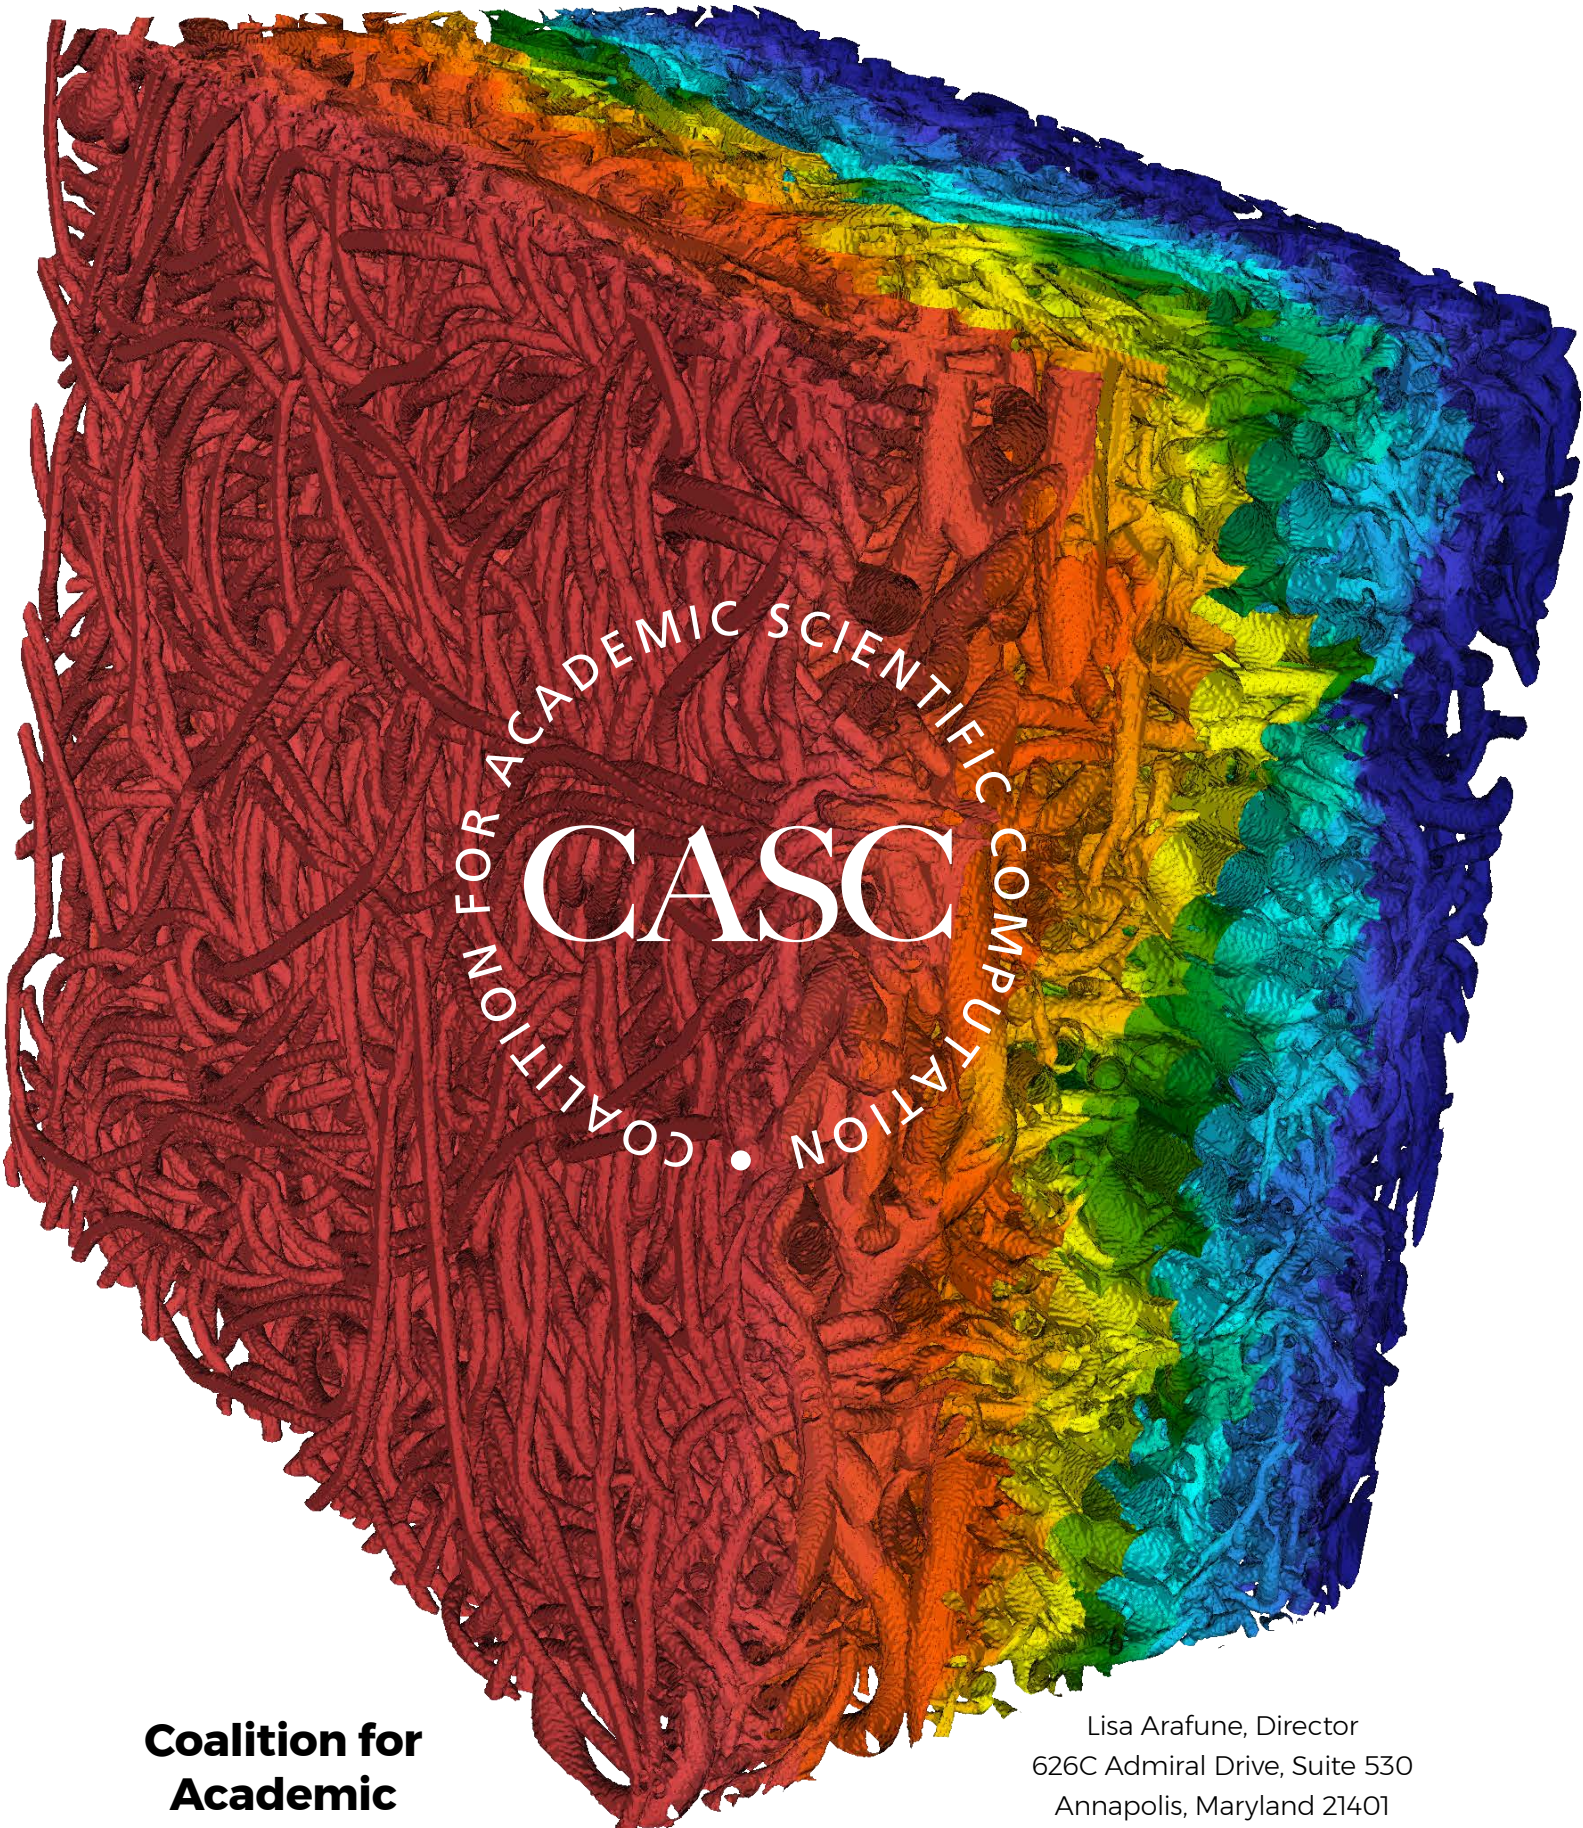
West Virginia University
Morgantown, West Virginia

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



The Cheyenne supercomputer, which came online in 2017, is operated by the National Center for Atmospheric Research (NCAR). Housed in the NCAR-Wyoming Supercomputing Center, Cheyenne is one of the world's most powerful and energy-efficient supercomputers, with a theoretical performance of more than 34 million calculations per second for every watt of power consumed.

Image copyright University Corporation for Atmospheric Research (UCAR), licensed under a Creative Commons Attribution-NonCommercial 4.0 International (CC BY-NC 4.0) License, via OpenSky.



COALITION FOR ACADEMIC SCIENTIFIC COMPUTATION • CASC

**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