

AN EXCITING NEW ERA!

2020 brings change—and an exciting new era—to PSC. Beginning his role this fall is PSC's new director, William Kramer, who comes to us after serving as principal investigator of the National Center for Supercomputing Applications' Blue Waters, one of the most powerful supercomputers in the world. Also, PSC will deploy its newest supercomputer, Bridges-2 (see below), which will build on the successes of the current Bridges system.



Carnegie Mellon University

University of Pittsburgh

NSF Funds Bridges-2 Supercomputer at PSC

\$10-Million System Will Expand National Capacity for Big-Data Science, Artificial Intelligence and Research by Nontraditional Supercomputing Users

A \$10-million grant from the National Science Foundation (NSF) is funding a new supercomputer at PSC. In partnership with Hewlett Packard Enterprise (HPE), PSC will build Bridges-2, a high-performance-computing system designed to provide researchers in Pennsylvania and the nation with massive computational capacity and the flexibility to adapt to the rapidly evolving field of data- and computation-intensive research. Bridges-2 will be available at no cost for research and education, and at cost-recovery rates for other purposes.

"Unlocking the power of data will accelerate discovery to advance science, improve our quality of life and enhance national competitiveness," says Nick Nystrom, PSC's chief scientist and principal investigator (PI) for Bridges-2. "We designed Bridges-2 to drive discoveries that will come from the rapid evolution of research, which increasingly needs new, scalable ways for combining large, complex data with high-performance simulation and modeling."

Bridges-2 will accelerate discovery to benefit science, society, and the nation. Its unique architecture will catalyze breakthroughs in critically important areas such as understanding the brain, developing new materials for sustainable energy production and quantum computing, assembling genomes of crop species to improve agricultural efficiency, exploring the universe via multimessenger astrophysics and enabling technologies for smart cities. Designed to be particularly accessible to researchers who have not before used supercomputers, Bridges-2 will also democratize high-performance computing for new communities of scientists who never before needed computers, let alone supercomputers.

"We look forward to having Bridges-2 as the next step in providing computational infrastructure for the open research

community in the United States," says William Kramer, director of PSC. He noted that PSC's unique position as a joint center of Carnegie Mellon University and the University of Pittsburgh, as well as its ongoing relationship with the University of Pittsburgh Medical Center, will also allow the new computer to be of particular value to those communities.

Building on PSC's experience with its very successful Bridges system, Bridges-2 will take the next step in pioneering converged, scalable high-performance computing (HPC), artificial intelligence (AI) and data. Designed to power and scale applications identified through close collaboration with the national research community, Bridges-2 will integrate cutting-edge processors, accelerators, large memory, an all-flash storage array and exceptional data-handling capabilities to let researchers meet challenges that otherwise would be out of reach. By enabling AI to be combined with simulation and modeling and through its focus on ease of use and researcher productivity, Bridges-2 will drive a new era of research breakthroughs.

"Enabling the execution of science, engineering and non-traditional workflows at scale while leveraging and further developing artificial intelligence is vital to keeping the United States at the forefront of scientific discovery now and into the future," says Paola Buitrago, Director of Artificial Intelligence & Big Data at PSC and co-PI of Bridges. "The Bridges-2 system is the way to realize this and more. I look forward to all the knowledge, discoveries and progress this new system will produce."

Bridges-2 will feature new technology such as Intel's 10nm Ice Lake processor, along with other Intel CPUs. It will also offer NVIDIA's graphics processing units (GPUs) for AI research and research questions that rely on analyzing graphic information rather than "traditional" calculus-based supercomputing. Like Bridges-1, the new machine will also offer nodes with different RAM, including 256 GB (gigabytes), 512 GB, and 4 TB (4,000 GB). For comparison, a high-end personal computer may have 64 GB of RAM. These nodes will enable work requiring the computer to analyze vast amounts of data.

Bridges-2 will be deployed in the summer of 2020.

PITTSBURGH SUPERCOMPUTING CENTER provides university, government and industrial researchers with access to several of the most powerful systems for high performance computing, communications and data storage and handling available to scientists and engineers nationwide for unclassified research. PSC advances the state of the art in high performance computing, communications and data analytics and offers a flexible environment for solving the largest and most challenging problems in computational science.

Pittsburgh Supercomputing Center is a joint effort of Carnegie Mellon University and the University of Pittsburgh. It was established in 1986 and is supported by several federal agencies, the Commonwealth of Pennsylvania and private industry.

Major Computing Resources

Bridges converges HPC, AI and Big Data to empower new research communities, bring desktop convenience to advanced computing, expand remote access and help researchers to work more intuitively. **Bridges-AI** is the premier resource for scalable AI in NSF-supported cyberinfrastructure, enabling deep learning training on the largest data and with the most sophisticated models. Available at no cost for open research, Bridges-AI consists of an NVIDIA DGX-2 Enterprise AI Research System with 16 Volta GPUs and nine HPE Apollo 6500 servers, each with 8 Volta GPUs (88 Volta GPUs total), and is fully integrated with Bridges to support AI-enabled simulation and complete data science workflows.

Anton 2 – a special-purpose supercomputer for biomedical simulation designed and constructed by D. E. Shaw (DESRES). A successor to Anton, Anton 2 is a 128-node system, made available to PSC by DESRES without cost for non-commercial research use by U.S. universities and other not-for-profit institutions. It is hosted by PSC with support from the National Institute of General Medical Sciences.

PSC gratefully acknowledges significant support from the following:

The National Science Foundation
The National Institutes of Health
The U. S. Department of Defense
The U. S. Department of Energy

The Commonwealth of Pennsylvania
D. E. Shaw Research
University of Pittsburgh

**ARTIFICIAL
INTELLIGENCE
ADVANCES EARLY
PREDICTION OF BREAST
CANCER DEVELOPMENT**

**PITTSBURGH SUPERCOMPUTING CENTER
SCIENCE HIGHLIGHTS**

FALL 2019

WWW.PSC.EDU/SCIENCEHIGHLIGHTS

PEACE OF MIND THROUGH AI

FEATURE

"DEEP LEARNING" ARTIFICIAL INTELLIGENCE ON PSC SUPERCOMPUTERS PROMISES FEWER FALSE ALARMS AND EARLY PREDICTION OF BREAST CANCER DEVELOPMENT

[PSC.EDU/BREASTCANCER](https://psc.edu/breastcancer)

ABOVE: Using AIs to identify false recalls by classifying the three categories (negative, false recalls, and malignancy) of digital mammogram images in breast cancer screening.

Why It's Important

Despite a lot of progress in improving survival and quality of life for women with breast cancer, the disease remains a major threat to women's health. It's the most common cancer in women and is either the first or second most common cause of cancer death for women in the largest racial and ethnic groups, accounting for 41,000 deaths in 2016 alone.

Screening mammography is an important tool for getting early warning, when the disease is easiest to treat. But it's not perfect. For women whose scans show no signs of breast cancer, doctors wonder whether that scan may contain information they could use to predict future risk. More than 10 percent of women who get mammograms are "recalled" for further testing. But nearly 90 percent of the time it's a false alarm. That's something like 3 million women in the U.S. who go through the stress of unnecessary recall each year.


Expert radiologists can tell a lot from a modern digital mammogram. But Shandong Wu and his colleagues at the University of Pittsburgh Medical Center (UPMC) wondered if artificial intelligence (AI) could detect subtle hints in mammograms that the human eye can't see. They tested their "deep learning" AIs on digital mammograms from UPMC patients whose status was already known, running the programs on PSC's Bridges-AI.

How PSC Helped

The task for the UPMC scientists' deep-learning software was a big one. Each digital mammogram is more than 2,000 by 3,000 pixels large—that's dozens of megabytes of data for each image. And to do their study, they needed to "train" and then test their AIs on thousands of images. Their AIs are also pretrained on large datasets with tens of thousands of images. The size of the data for AI modeling was enormous, making the computations slow on the computers available to the researchers at their own laboratory.

Working with Roberto Gomez and other staff at PSC, the team used the graphics processing unit (GPU) nodes of Bridges-AI to train and run their AIs. Deep-learning, which works by building up layers of different kinds of information and then pruning connections between the layers that don't produce the desired result, tends to work best on GPUs. The new NVIDIA "Volta" GPUs in Bridges-AI contain accelerators, called "tensor cores," specifically designed for deep learning. Bridges-AI's GPU nodes combine eight to 16 GPUs each for up to 512 gigabytes of extremely fast GPU memory. The large memory available to PSC's GPU nodes was central to the success of the AIs, bringing the computation time down from weeks to hours. The NVIDIA DGX-2 node, deployed in Bridges-AI as a first for open research, and its massive memory were particularly useful.

When an AI is designed to produce a binary result—yes or no, positive or negative—scientists often report that result as a graph of true positives versus false positives. The larger the "area under the curve," or AUC, the better the AI's accuracy.



One of the predicted new low cobalt structures of $\text{Li Nix Mn y Co } 1-x-y \text{ O}_2$ with a ratio of nickel to manganese to cobalt of 18:5:1. The nickel is shown in grey, the manganese in magenta, and the cobalt in blue. The lithium layer is shown in green and oxygen is shown in red. Gregory Houchins.

HIGHLIGHTS

HE WHO HESITATES - [PSC.EDU/VIRUS](https://psc.edu/virus)

Viruses such as influenza and HIV take a heavy toll, both in human life and in dollars. But these viruses are shape-shifters, changing their outer proteins via mutation so human immunity and antiviral drugs have a hard time keeping up. A team from Rice University and Baylor College of Medicine has used the Anton 2 supercomputer at PSC to simulate the role that a portion of a critical protein called "hemagglutinin (HA)" plays in the flu virus merging with host cells and injecting its genes. They found that the protein's "stem domain," HA2, pauses in the process—a critical step that may offer a target for future therapy that can outsmart the viral transitions. The work offers a completely new approach to therapy since it focuses on the critical injection mechanism rather than the virus's outward structures.

BUILDING BETTER BATTERIES - [PSC.EDU/BETTERBATTERIES](https://psc.edu/betterbatteries)

The move toward cleaner, cheaper energy would be much easier if we had more powerful, safer battery technologies. Carnegie Mellon University scientists are using PSC's Bridges supercomputer to simulate new battery component materials that are inherently safer and more powerful than currently possible.

LIFE ISN'T CHESS—TIMES SIX - [PSC.EDU/PLURIBUS](https://psc.edu/pluribus)

Artificial intelligence (AI) research took a great leap forward when a Carnegie Mellon University computer program overcame the world's best professional players in a series of six-player poker games. Running on PSC's Bridges supercomputer, the Pluribus AI was the first to surpass humanity's best at a multiplayer, "incomplete information" game. Experimenting with such games offers more useful lessons for real-world problems such as security, business negotiations and cancer therapy than "complete information" games such as chess or GO.

GOING WITH THE FLOW - [PSC.EDU/HEATTRANSFER](https://psc.edu/heattransfer)

Smaller electronic components offer us more power in our pockets. But thinner and thinner components pose engineering problems. Anisotropic materials—those with properties that vary in direction—hold promise for being unusually versatile. Still, their properties, particularly as they become "two dimensional" or ultra-thin, are not well understood. Using a combination of calculations on PSC's Bridges and lab experiments, a UCLA-led group showed that an atomistic model can explain and predict the transfer of heat between aluminum and black phosphorous, a highly anisotropic material with possible applications in future devices.

The HA2 protein pauses half-way through unfolding (left), allowing its fusion peptide "arms" (blue) to reach between the viral membrane and the host cell membrane.



NEWS IN BRIEF

- AIDR 2019 Conference Brings AI Researchers to Pittsburgh
- PSC-Affiliated Cybersecurity Center of Excellence Receives \$12.5-Million Renewal Grant
- Pittsburgh Girls Code
- BEST Summer Institute Trains High School Teachers in Computational Biology
- Computer Simulation: Antibody-Resistant Infection Registry Can Reduce Prevalence
- PSC Team Wins Best Student Paper Award
- Bridges Powers AI Poker Demonstration
- PSC Interns Explore Data, Biology, Computing Architecture