

Pitt-CMU's Powerful Supercomputer 'Bridges' Expands Its Reach

By Adam Reger

Issue Date: September 19, 2016

Researchers in science, engineering, and other big data fields have been using supercomputers to model and visualize complex processes, and to mine big data and optimize its sharing on the world's fastest computers, including those housed at the Pittsburgh Supercomputing Center (PSC), a joint venture of Pitt and Carnegie Mellon.

"The PSC is one of five National Science Foundation-funded supercomputing centers in the nation, and the only one in the Northeast," said Ralph Roskies, the center's scientific director and a professor of physics in Pitt's Department of Physics and Astronomy.

Now those vast computing capabilities are available to a far broader range of researchers through a new system designed by a PSC team in response to a National Science Foundation competition calling for a supercomputer that would engage "nontraditional users."

The result is Bridges, a supercomputer that possesses 17 petabytes of data storage capacity and the potential to move and access data at extremely high speeds. One petabyte, equivalent to one million gigabytes, could hold about one-half of all the content contained in the nation's academic research libraries. That capacity puts Bridges in elite company.

What really makes Bridges special, said Roskies, is its accessibility and ease of use, especially for scholars who may never have considered using a supercomputer for their work. And those are exactly the users that the designers of Bridges had in mind, said Nick Nystrom, the principal investigator of the Bridges project.

The PSC has long supported nationwide efforts in, for example, physics, chemistry, molecular dynamics, engineering, and climate simulation. But what about scholars in academic fields not traditionally associated with supercomputers?



From left, Nick Nystrom, Pittsburgh Supercomputing Center's senior research director, and Ralph Roskies, PSC's scientific director and a Pitt physics professor (Photo by Emily O'Donnell)

OTHER STORIES FROM THIS ISSUE SEPTEMBER 19, 2016

[Year of Diversity: Getting Comfortable with Differences](#)

[First Person: What's in Your Culture Box?](#)

[Swanson School Innovates with "Make Research Work" Private-Sector Partnerships](#)

[Laura Zullo: Managing Pitt's Carbon Footprint](#)

[Pitt-CMU's Powerful Supercomputer 'Bridges' Expands Its Reach](#)

[Lynn Emanuel Honored for Outstanding Poetry](#)

[Call for Nominations: Chancellor's Awards](#)

[Katz MBA Student Leads Drive for Global War on Terror Memorial](#)

[Ready to Learn Mentor Program Makes Progress at Pittsburgh Millions UPrep](#)

[Funding Given for Innovation in Teaching of Natural Sciences](#)

[Allegheny Observatory Open House Set for Oct. 7](#)

[Awards & More](#)

[Happenings](#)

ON THE FREEDOM ROAD



Follow a group of Pitt students on the Returning to the Roots of Civil Rights bus tour, a nine-day, 2,300-mile journey crisscrossing five states.

[Introduction](#)

[Day 1: The Awakening](#)

Roskies cited the example of "distant reading," an emerging research method in the digital humanities that is made possible by supercomputers. Bridges is capable of scanning hundreds of thousands of digital documents to identify literary and cultural trends, perform authorship attribution, and link to other fields.

"No one can read a hundred thousand books," said Roskies, "but a system like this can."

Another example is the Center for Causal Discovery (CCD), which is working to understand the root causes of cancer, lung disease, and brain dysfunction. CCD is a National Institutes of Health Big Data to Knowledge Center for Excellence, and it is led by researchers at Pitt, CMU, PSC, and Yale University. Bridges is uniquely capable of supporting the Center for Causal Discovery's need to integrate high-performance computing with big data through its large-memory nodes and its dedicated nodes for databases and web portals. (A node is a discrete computational element in the supercomputer.)

And the new supercomputer is ideal for problem-solving in the public health arena, too. PSC's Public Health Applications Group used Bridges' massive computing power to design a program that "virtually" explored the public's potential acceptance of a new nasal mist flu vaccine versus the traditional, injected, option. By drawing on vast existing collections of data from large metropolitan areas and using high-performance computing to cull and analyze those data, Bridges was able to suggest that, compared to current vaccination programs that offer just one type of vaccine, giving individuals a choice of options would result in higher vaccination rates.

Bridges' flexible design can accommodate many different researchers and projects at once. Currently, it is allocating space to 299 research groups and 1,932 individual users.

Another key element of Bridges' appeal to nontraditional supercomputer users is its prominent use of popular "gateways," which enable users to connect to the system directly through convenient web interfaces. This allows users to get quickly and easily to work, directly from their offices without needing to learn new programs.

"Gateways allow users who are experts in their fields to use Bridges without becoming programmers or even knowing that they're using a supercomputer," said Nystrom, who also is the PSC's senior director of research and a CMU research physicist. Bridges' gateways transparently deliver supercomputing as a service—similar to how a single Internet search is actually accomplished via the vast computational power of internet search engines.

Even for researchers working in fields where high-powered computing is frequently used, Bridges' flexible architecture opens new avenues of inquiry. Users can import pre-existing successful software into the supercomputer without rewriting it for Bridges' system. Bridges also offers popular high-productivity languages such as R, Python, MATLAB, and Java, which are not available on traditional supercomputers.

PSC staff are available to help users tailor Bridges to fit their needs, assisting researchers in designing experiments and fine-tuning software applications to use the

Day 2: [Deep Impressions](#)
 Day 3: [Music, Montgomery, and More](#)
 Day 4: [Looking Back, Looking Forward](#)
 Day 5: [Learning to Remember](#)
 Day 6: [The Mountaintop](#)
 Day 7: [Slavery and Beyond](#)
 Day 8: [Lessons to Bring Home](#)
 Day 9: [Final Lessons](#)



The Bridges supercomputer comprises several refrigerator-sized cabinets, each containing multiple processors.

system's unique capabilities to their fullest. Staff also conduct frequent workshops to help users exploit the system's full capabilities.

Those capabilities extend well beyond the system's pure speed and power. Intel's Omni-Path Architecture (OPA) technology fully interconnects Bridges' 908 distinct nodes and the system's shared file system. This interconnectivity enables data to flow freely between all system components. "Bridges' OPA network enables the highest possible performance for the kinds of applications that researchers are running," said Nystrom. Other high-bandwidth networking enables efficient data transfers to and from other sites over the Internet.

There is no fee to use Bridges for U.S.-based research groups that are engaged in open research, subject to peer review of their proposals for access. Nystrom added that access is also available, for a fee, to industry engaged in appropriate, computer- or data-intensive research.

PSC's offices, home to approximately 70 full-time staff, are located on Craig Street in Oakland. Bridges, and other PSC supercomputers like the molecular simulation system Anton and the public-health oriented Olympus, are housed in PSC's secure data center in Monroeville.

The project's Phase 2 Technical Upgrade, which will enhance Bridges' processing capability, is expected to be completed this fall.



Pitt Chronicle
Office of University Communications
200 South Craig Street
Pittsburgh, PA 15260

ucomm@pitt.edu
412-624-1033

Copyright 2018 | [Send feedback about this site](#)