

The National Resource for Biomedical Supercomputing

National Leadership in High-Performance

Computing for Biomedical Research

Established in 1987, PSC's National Resource for Biomedical Supercomputing (NRBSC) was the first external biomedical supercomputing program funded by the National Institutes of Health (NIH). Along with core research at the interface of supercomputing and the life sciences, NRBSC scientists develop collaborations with biomedical researchers around the country, fostering exchange among experts in computational science and biomedicine and providing computational resources, outreach and training. In October 2006, NRBSC received \$8.5 million from NIH's National Center for Research Resources (NCRR) to renew its work for five years.

"Over the past decade, computing has become essential to almost all aspects of biomedicine," says PSC's Joel Stiles, director of NRBSC, a medical doctor who also holds a doctorate in physiology. "Here at the NRBSC, we're developing and distributing computational tools in simulation,

visualization, and education that are helping to transform our understanding of life and disease."

MORE INFORMATION:
<http://www.nrbsc.org>

COMPUTATIONAL SERVICE & TRAINING

Since NRBSC's inception, PSC and NRBSC together have provided access to computing resources for more than 1,400 biomedical research projects involving more than 4,200 researchers at 274 research institutions in 46 states and two territories. Among these are several projects featured in this booklet (pp. 18 & 22).

NRBSC training activities reach hundreds of scientists each year. More than 3,500 researchers have participated in NRBSC workshops in such areas as spatially realistic cell modeling, volumetric



← The NRBSC team: (seated, l to r) Boris Kaminsky, Pallavi Ishwad, Jenda Domaracki, (standing) Art Wetzel, Jack Chang, Markus Dittrich, Troy Wymore, Christal Banks, Alex Ropelewski, Joel Stiles, Adam Kraut, Greg Hood. Not pictured: Hugh Nicholas and Jacob Czech

NRBSC BIOMEDICAL COLLABORATIONS

Albert Einstein College of Medicine
Carnegie Mellon University
Cornell University
Duke University
Harvard University

Howard University
Janelia Farm, Howard Hughes Medical Institute
Marine Biological Laboratory, Woods Hole
Morgan State University
North Carolina Central University


Rockefeller University
The Salk Institute
The Scripps Research Institute
University of California at Davis
University of California at San Diego
University of Kansas

University of North Carolina, Chapel Hill
University of Pittsburgh
University of Pittsburgh School of Medicine
University of Puerto Rico, Medical Sciences Campus

RESEARCH

NRBSC research focuses on three areas of biomedicine that span many scales of space and time: spatially realistic cell modeling, large-scale volumetric visualization and analysis, and computational structural biology.

SPATIALLY REALISTIC CELL MODELING centers on realistic → 3-D simulations of movements and reactions of molecules within and between cells, to better understand physiological function and disease. *MCell*, *DReAMM* and *PSC_DX* software is developed at the NRBSC and used to model and visualize events such as this image, which represents neurotransmitter release in one dendritic spine.



data visualization and analysis, protein and DNA structure, genome sequence analysis and biological fluid dynamics."


NRBSC participates in a range of undergraduate and graduate training programs.

THESE INCLUDE:

- A joint Carnegie Mellon and University of Pittsburgh Ph.D. program in computational biology (www.compbio.cmu.edu)
- The Ray and Stephanie Lane Center for Computational Biology at Carnegie Mellon (lane.compbio.cmu.edu)
- The University of Pittsburgh Department of Computational Biology (www.cccb.pitt.edu)
- The undergraduate Bioengineering & Bioinformatics Summer Institute (www.cccb.pitt.edu/bbsil), sponsored by NRBSC, Carnegie Mellon, the University of Pittsburgh, and Duquesne University, and funded jointly by NSF and NIH.


K-20 SCIENCE OUTREACH

The NRBSC and PSC have developed innovative Computational Modules In Science Teaching (CMIST) for high school and undergraduate biology, chemistry, physics, computer science and math. CMIST modules bring critical concepts to life in novel ways, using realistic models and simulations with visually appealing, scientifically accurate animations (see p. 46). NRBSC distributes the modules online and on DVDs. They include lecture slides, animations, lesson plans aligned to national and state standards, worksheets and answer keys. This year NRBSC developed a second CMIST module on atomic and molecular movements, bridging enormous space and time scales that are important to biological systems. A third module on enzyme structure and function is also currently under development.



VOLUMETRIC →

VISUALIZATION using the NRBSC's PSC_VB software enables multiple users to share, view and analyze extremely large datasets and time series obtained from light and electron microscopes, CAT and MRI scanners, etc. This image from an MRI dataset, viewed with PSC_VB's volume browser, shows detailed structure of the left ventricle of a mouse heart with an error range of less than 3-percent. Mice hearts are frequently used in cardiac research, and this high degree of accuracy facilitates extended duration studies of cardiac development and recovery after injury.



NRBSC STRUCTURAL BIOLOGY focuses on computational tools used to determine the structure of proteins from their amino acid sequence and development of quantum-mechanical simulation methods for biomolecules such as enzymes. This image shows the 3-D structure of r-hpcdh, an enzyme that catalyzes a key reaction — a coupled proton/hydride transfer. PSC-developed software enables researchers to simulate enzyme reactions, to reproduce experimental reaction rates and gain new insight into enzyme function, which facilitates design of new therapeutic drugs.

Networking the Future

One of the leading resources in the world for network know-how

PSC's Advanced Networking group is one of the leading resources in the world for knowledge about networking. Through 3ROX (Three Rivers Optical Exchange), a high-speed network hub, they provide high-performance networking for research and education. Their research on network performance and analysis — in previous projects such as Web100 and the NPAD diagnostic server — has created valuable tools for improving network performance nationally.

MORE INFORMATION:
<http://www.psc.edu/networking/>



← Tim Devlin, program director of instructional media services, Allegheny Intermediate Unit, discusses educational impact of PSC's new PoP in Allegheny County.

3ROX BRANCHES OUT: PSC AND FIBERTECH OPEN NEW POP

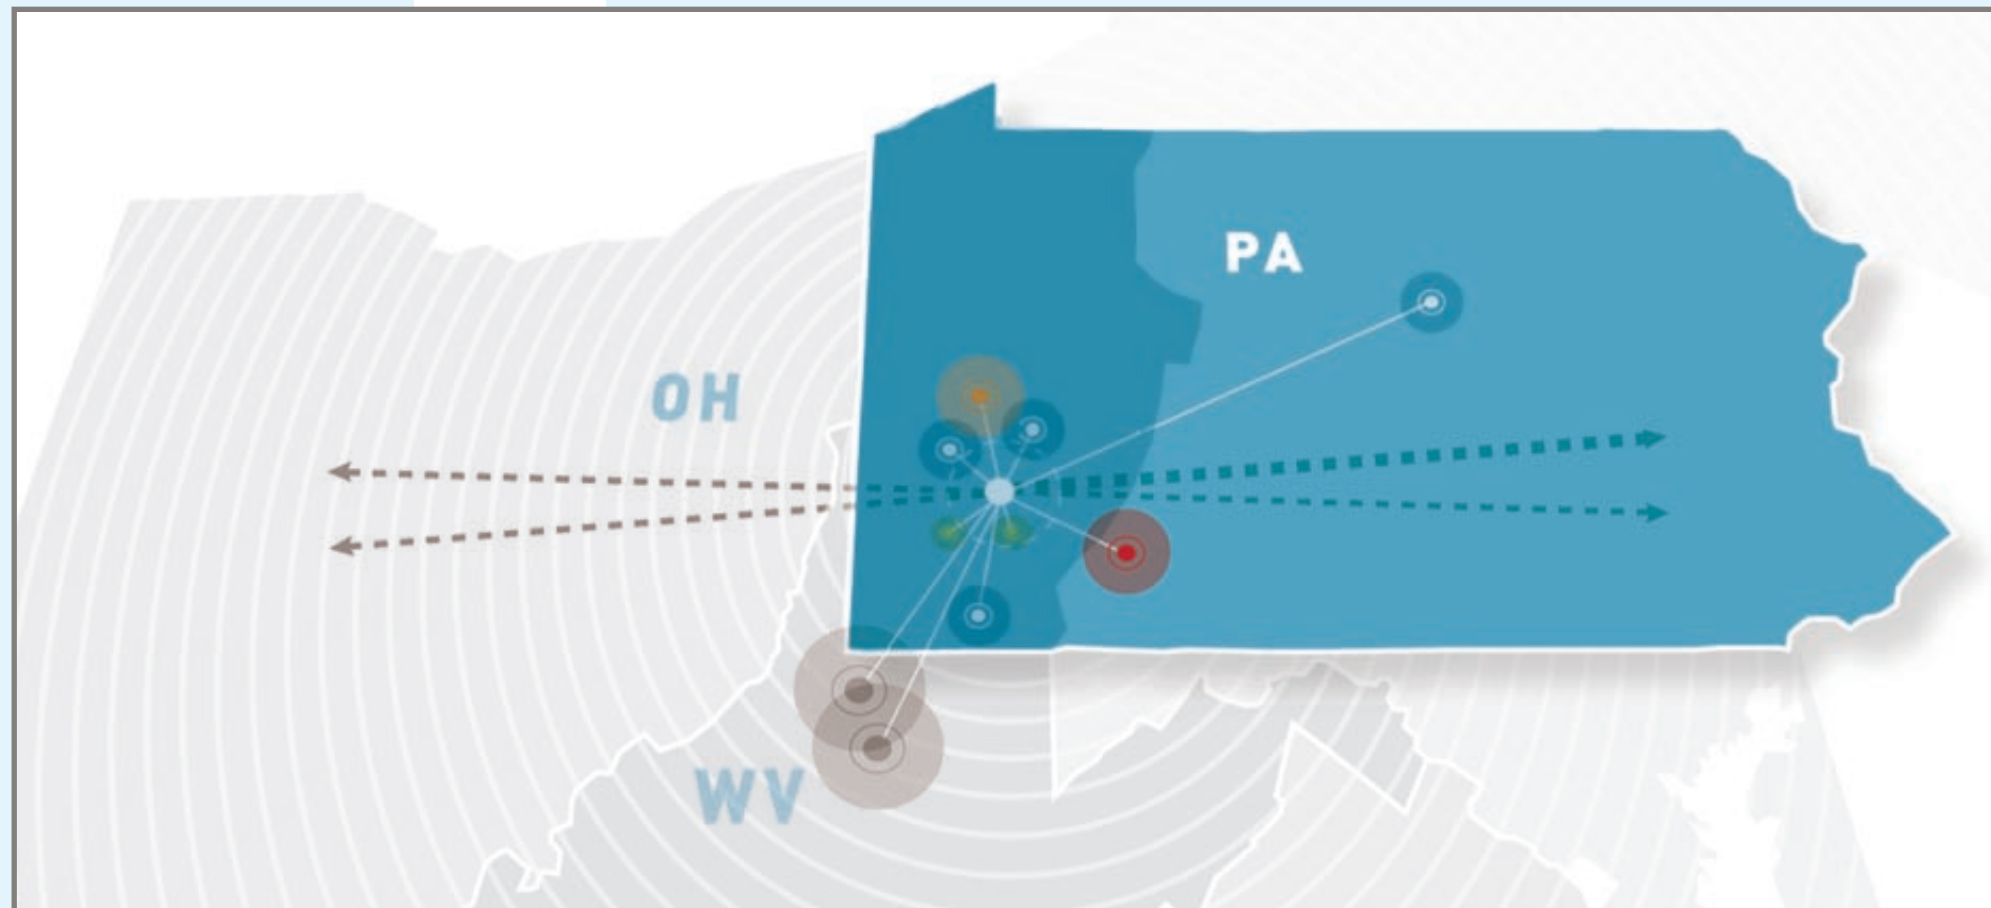
Through 3ROX, PSC connects universities and public schools in Pennsylvania and West Virginia to high-performance networks, such as Internet2, which links leading U.S. universities, corporations, government research agencies, and not-for-profit networking organizations.

This year 3ROX extended its connectivity in southwest Pennsylvania by opening a second PSC PoP (point of presence) in Fibertech's new collocation facility in Pittsburgh's Allegheny Center Mall. At an April kickoff event, PSC Director of Networking Wendy Huntoon noted in brief remarks that the new PoP is PSC's first off-campus location, making connectivity through 3ROX more cost-effective in providing underlying support for economic growth in the southwest Pennsylvania region.

Tim Devlin, program director of instructional media services, Allegheny Intermediate Unit, described the networking program of the Allegheny Intermediate

Unit, 42 suburban school districts surrounding the city of Pittsburgh. Optical fiber will link all 200 school buildings of all 42 districts and, through the Allegheny Center Mall (ACM) PoP, connect back to 3ROX, providing access to state-of-the-art educational resources.

FiberTech Networks operates a Metropolitan Wide Area Network in Pittsburgh that uses fiber optic-based transport services. Available to local businesses, the Allegheny Center Mall facility will enable customers to meet their current and future connectivity requirements with unlimited scalability.



3ROX MEMBERS

- **Universities**
 Carnegie Mellon University, Pennsylvania State University, University of Pittsburgh, Waynesburg University, West Virginia University
- **K-12 Institutions**
 Allegheny Intermediate Unit (AIU3), Arin Intermediate Unit (IU28), Beaver Valley Intermediate Unit (IU27), Intermediate Unit One, Northwest Tri-County Intermediate Unit (IU5), Riverview Intermediate Unit (IU6), City of Pittsburgh School District (IU2), Woodland Hills School District
- **Business**
 Comcast, Westinghouse Electric Co.
- **Government Laboratory**
 The National Energy Technology Laboratory
- **Other**
 Computer Emergency Response Team

NETWORK CONNECTIONS

- **National Research Networks**
 Internet2 — 1 Gbps, ESnet — 1 Gbps, National LambdaRail PacketNet — 10 Gbps, TeraGrid Extensible Backplane Network — 30 Gbps
- Other Network Connections**
 Southern Crossroads (SOX) — 1Gbps, TransitRail — 1Gbps
- ← **National Commodity Internet Networks**
 Global Crossing — 1 Gbps; Sprint — 1 Gbps
- Pittsburgh Local Exchange Networks**
 Comcast

Note: Gbps: a billion (Giga) bits per second

Networking the Future (continued)



← Matt Mathis, PSC senior network engineering specialist. A paper he co-authored in 1997 won the ACM SIGCOMM Test of Time Award.

ADVANCED NETWORK RESEARCH: PSC NETWORK ENGINEER RECOGNIZED

Three organizations this year recognized PSC senior network engineering specialist, Matt Mathis, for his network research. Mathis has been a network engineer at PSC since 1988, helping to lead such projects as NPAD and Web100 and the related Net100.

In August, the Special Interest Group on Data Communication (SIGCOMM), of the Association for Computing Machinery (ACM), notified Mathis that he will receive their Test of Time Award. ACM SIGCOMM is the leading forum for professional discussion of data communications and computer networks. The award recognizes Mathis's 1997 paper, co-authored with former PSC staff members Jamshid Mahdavi and Jeff Semke and with Teunis Ott (then at Bellcore, now at the New Jersey Institute of Technology). "The macroscopic behavior of the TCP congestion avoidance algorithm," that was a foundation for Internet traffic control standards.


The paper was published in the ACM journal *Computer Communication Review*. The Test of Time Award honors papers from 10 to 12 years ago in *CCR* deemed by a committee to be "an outstanding paper whose contents are still a vibrant and useful contribution today."

In October, Cisco System's Collaborative Research Initiative presented an unrestricted gift of \$65,500 to Carnegie Mellon University to support Mathis's research proposal titled "Rebalancing Internet Congestion Control." The project will explore the possibility of changing how the Internet manages traffic. "Our goal," says Mathis, "is to shift responsibility for allocating network capacity from the end-systems to the network itself, such that the network can support the safe operation of diverse control algorithms."

The Cisco-funded research is related to the earlier paper, explains Mathis: "The macroscopic model paper was key to establishing the 'TCP-friendly' principle that has guided Internet congestion control standards. The Cisco funding intends to move the Internet beyond this principle, to permit new standard network protocols


that are not bound to politely share the network with legacy TCP."

In September, BBN Technologies funded a PSC proposal from Mathis for work on prototyping of the NSF-sponsored GENI suite of network research infrastructure. This work will proceed using a "spiral development" approach in which simultaneous trials will give rapid feedback to guide evolving designs. Multiple competing approaches are funded, of which Mathis's work is one.



↑ PSC's directors and managers (l to r): Katherine Vargo, manager of scientific computing systems; Cheryl Begandy, outreach manager; David Kapcin, director of financial affairs; Elvira Prologo, manager of administration; Joel Stiles, director of NRBS; Bob Stock, PSC associate director; John Kochmar, manager of high-performance computing facilities; Richard Raymond, manager of user support; J. Ray Scott, director of systems & operations; Wendy Huntoon, director of networking; Sergiu Sanielevici, director of scientific applications & user support. **Not pictured:** Nick Nystrom, director of strategic applications; Laura McGinnis, manager of data & information resource services; Janet Brown, manager of networking.

A workshop underway in the PSC Computer Training Center, the David W. Deerfield II Training Center, equipped with 30 "dual-boot" workstations and a projector for overhead display of the instructor's desktop.



PITTSBURGH SUPERCOMPUTING CENTER WORKSHOPS (2007-2008)

Hybrid Programming for Shared-Memory and Clustered SMP Systems
Summer Institute in Bioinformatics (for minority-serving institutions)
Computational Methods for Spatially-Realistic Microphysiological Simulations
Parallel Computing

PROJECTS

IN

SCIENTIFIC

COMPUTING

2008