BLUE WATERS PROJECT

QUARTERLY NEWSLETTER

Volume 1
Issue 3
September 2010

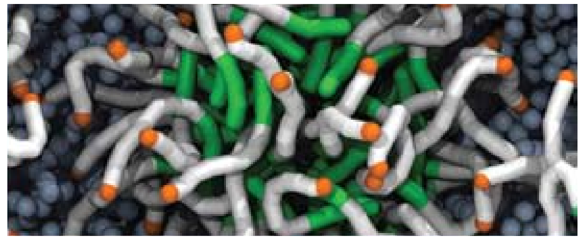

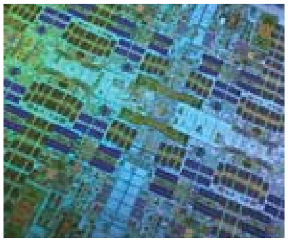

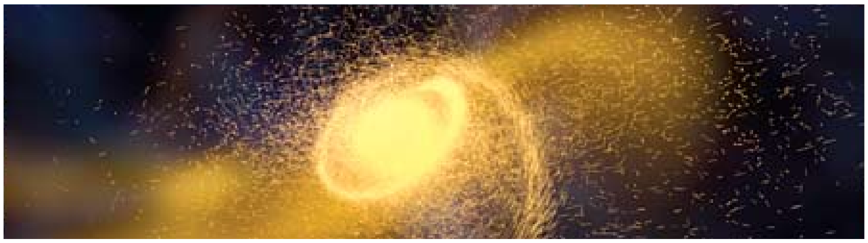

Blue Waters Interconnect

BLUE WATERS will employ new IBM interconnect technology that combines concepts from IBM's Federation supercomputing switches and InfiniBand. The high-bandwidth, low-latency interconnect will scale to hundreds of thousands of cores.

The key component is the hub/switch chip. The four POWER7 chips in a compute node Quad-chip module (QCM), are connected to a hub/switch chip, which serves as an interconnect gateway as well as a switch that routes traffic between other hub chips. The system therefore requires no external network switches or routers, providing considerable savings in switching components, cabling, and power.

The hub chip provides 192 GB/s to the directly connected POWER7 QCM; 336 GB/s to seven other nodes in the same drawer on copper connections; 240 GB/s to 24 nodes in the same supernode (composed of four drawers) on optical connections; 320 GB/s to other supernodes on optical connections; and 40 GB/s for general I/O, for a total of 1,128 GB/s peak bandwidth per hub chip.

The Blue Waters interconnect topology is a fully connected two-tier network. In the first tier, every node has a single hub/switch that is directly connected to the other 31 hub/switches in the same supernode. In the second tier, every supernode has a direct connection to every other supernode. These inter-supernode connections terminate at hub/switches, but a given hub/switch is directly connected to only a fraction of the other supernodes. Messages traveling from a hub/switch in one supernode to another supernode in many cases must be routed through another hub/switch in both the sending and receiving supernodes. Multiple routes are possible with direct and indirect routing schemes. The minimum direct-routed message latency in the system is approximately 1 microsecond.

Indirect routes add an intermediate supernode to the path described above, which (depending on the running program's communication pattern) can be useful for avoiding contention for the link that directly connects the two communicating supernodes. A key advantage of Blue Waters interconnect is that the number of hops that must be taken by a message passed from any supernode to any other supernode in the system is at most three for direct routes and five for indirect routes, independent of the number or location of the processors that are communicating.

The hub chip participates in the cache-coherence protocol within the node, enabling it to extract the data to be sent over the interconnect from either the sending processor's caches or main memory and to place data from the interconnect into either the L3 cache or the main memory of the receiving processor, potentially reducing memory access latency significantly.

The hub chip provides specialized hardware called the Collectives Acceleration Unit (CAU) to significantly reduce the network latency associated with frequently used collective operations, such as barriers and global sums, offloading this work from the processors. The CAU can perform multiple independent collective operations occurring simultaneously between various groups of processors.

The hub includes hardware support for global address space operations, active messaging, and atomic updates to remote memory locations. These features help improve the performance of one-sided communication protocols as well as PGAS languages, such as X10, UPC, and Co-Array Fortran.

Researchers from IBM and NCSA's Torsten Hoefler collaborated on a paper and presentation for the Hot Interconnects conference in August. The paper is published in the *Proceedings of 18th Symposium on High-Performance Interconnects*.

For more information, see:
www.ncsa.illinois.edu/BlueWaters/system.html

Blue Waters is a joint effort of the University of Illinois at Urbana-Champaign, its National Center for Supercomputing Applications, IBM, and the Great Lakes Consortium for Petascale Computation. It is supported by the National Science Foundation, the University of Illinois at Urbana-Champaign, and the state of Illinois.

Undergraduate Institute Prepares Students to Integrate Computer Science, Other Research Disciplines

INTEGRATED LEARNING is the driving principle behind the Undergraduate Petascale Institute, which is training a new generation of scientists. They'll be computational scientists who—from the very beginning of their educations and careers—make computational simulation an integral part of their research and understand how those simulations can drive discovery.

The institute—a collaboration between NCSA and Shodor—is part of the Blue Waters project. It kicked off this summer, with more than 20 students giving up their Memorial Day weekend, then spending the next two weeks at NCSA in an intensive computer science workshop. At the workshop, they learned about the architecture of the Blue Waters supercomputer and the programming languages they will need to use the computer. They also found the time for some Ultimate Frisbee.

These skills—well, maybe not the Frisbee—will be applied more broadly throughout the students' careers.

"All the stars interact at the same time. All the electrons in an atom interact at the same time," says Shodor's Bob Panoff. "If we're going to get better science from these machines, especially a machine like Blue Waters, which is going to have hundreds of thousands of processing cores, we're going to have to be able to map the parallelism in nature more closely to what the machine is capable of computing."

During the institute, students learned from experts in the field. And from each other.

The program didn't stop with the summer workshop. Each student is paired with a faculty mentor at their home institution. Together, the institute students and mentors are working on real-world research projects that require computer simulation in fields like chemistry, biology, and astrophysics.

Watch students describe the experience at the institute:
www.ncsa.illinois.edu/News/Video/

When Tornadoes Strike

In many ways, tornadoes remain a mystery. While their powerful winds can leave broad swaths of devastation—damaged and destroyed homes and businesses, injuries and fatalities—twisters are a challenge for atmospheric scientists to study. It's difficult (not to mention hazardous) to be in the exact right place at the precisely the right time to see one in action, and models that capture all of the critical small-scale details within a dynamic storm system are computationally demanding.

That's where the sustained-petaflop Blue Waters supercomputer comes into the picture. The University of Illinois' Bob Wilhelmson is the principal investigator for a Petascale Computing Resource Allocation (PRAC) project. The National Science Foundation's PRAC program enables scientists (in this case Wilhelmson and collaborators at Illinois, the National Center for Atmospheric Research, the National Oceanic and Atmospheric Administration, the University of North Dakota, and Central Michigan University) to work closely with the Blue Waters team, preparing their codes to take full advantage of its several hundred thousand cores.

One target for Wilhelmson's work is tornadogenesis—the process by which a tornado forms. Forecasters can identify conditions in which tornadoes are likely to form, but it's not yet possible to predict exactly when and where tornadoes will form, how strong they will be, or what path they will follow. To address these questions, the PRAC team is working with CM1, a sophisticated model of atmospheric dynamics and thermodynamics, in order to simulate the interaction between the birth of a tornado from a supercell in the ultra-high resolution needed to adequately represent important small-scale features that influence the evolution of the tornado. More detailed simulations using Blue Waters may provide valuable insights into the environmental and particular storm conditions that support tornado development and longevity.

Better understanding of what conditions lead to the most dangerous tornadoes can be used to improve the lower-resolution models used in routine forecasting, leading to more precise warnings about the most dangerous twisters.

Getting breakthrough results on Blue Waters won't just be a matter of using existing storm-tornado modeling codes on a more powerful computer.

The PRAC team also is working to improve the CM1 model, originally developed by team member George Bryan at the National Center for Atmospheric Research, and the microphysics using the expertise of Matt Gilmore at the University of North Dakota.

Team members:
George Bryan, National Center for Atmospheric Research
Matthew Gilmore, University of North Dakota
Brian Jewett, University of Illinois at Urbana-Champaign
Joe Klemp, National Center for Atmospheric Research
John Michalakes, National Center for Atmospheric Research
Leigh Orf, Central Michigan University
Glen Romine, National Center for Atmospheric Research
Bill Skamarock, National Center for Atmospheric Research
Lou Wicker, National Oceanic and Atmospheric Administration
Robert Wilhelmson, (principal investigator) University of Illinois at Urbana-Champaign
Paul Woodward, University of Minnesota

More than 1,000 Visit New Supercomputing Facility

More than 1,000 people visited the University of Illinois' new National Petascale Computing Facility June 17, touring the state-of-the-art building and learning more about the supercomputers it will house. The building will be home to Blue Waters, a supercomputer capable of performing 10 quadrillion calculations every second. Scientists will use Blue Waters to better understand a wide range of phenomena, from the formation of tornadoes to the complex workings of cells in our body. Although the supercomputer is still being built and won't be installed until next year, IBM provided a sample Power7 node so that visitors could see what the new technology will look like.

GLCPC Outlines Goals for Coming Year

"WHILE THE GLCPC'S MISSION IS NOT LIMITED TO BLUE WATERS, NCSA has committed time on the petascale machine to GLCPC members, so we should use it as effectively as possible," says new GLCPC Board President Maxine Brown of the University of Illinois at Chicago. "There are very few research teams today who can tackle major problems that consist of complementary yet interdependent subsystems and who can utilize the power of petascale. For GLCPC, the whole can truly be greater than the sum of the parts, assuming we have the right mix of synergy and energy focused on issues of regional and national priority."

The GLCPC is poised to play a major role in accelerating the development of the research teams, tools, and system software to enable multi-site collaborations for complex problem solving, she says, with these goals:

- ***The GLCPC should foster inter-institutional collaborations*** by creating "networks of scientists" to tackle outstanding problems of regional, national will or global importance. The GLCPC must pursue a long-term strategy of organizing "bottom-up" workshops, along specific themes, for the purpose of bringing together people from the various GLCPC institutions that have complementary expertise and/or resources to meet, brainstorm solutions to complex problems that single investigators cannot do alone, to create plans, to apply for time on Blue Waters, and to write grants to fund these activities.
- ***GLCPC will actively create "champions"*** or evangelists, at its member institutions who proactively seek out interested researchers and educators at their home institutions and keep them informed about GLCPC and petascale developments. At some point, these champions should create formal mechanisms to communicate with their local constituents and to promote the talents and resources of those local researchers to the greater GLCPC community.
- ***The GLCPC leadership will keep its members informed*** if these institution "champions" are to succeed. The GLCPC must promote its activities in a variety of media communications in a clear, concise and timely manner. And, GLCPC must continually create opportunities for inter-institutional collaborations and publicize potential funding opportunities.
- ***GLCPC will also be involved in prototyping a tiered collaboration and visualization infrastructure of the future,*** initially, with more powerful compute/storage/network/visualization for GLCPC/PRAC participants and compatible lower-cost entry systems for others.

More information about GLCPC can be found at: www.greatlakesconsortium.org.

Brown Elected President of GLCPC Board

Great Lakes Consortium for Petascale Computation named Maxine Brown their board president at a meeting in May. Brown is associate director of the Electronic Visualization Lab at the University of Illinois at Chicago. Stan Ahalt, director of the University of North Carolina's Renaissance Computing Institute, was voted president-elect. William Punch, director of Michigan State's High Performance Computing Center, was elected secretary/treasurer. Immediate past president John Ziebarth, senior vice president and chief operating officer of the Krell Insititute, will continue to serve on the executive committee, as will Srinivas Aluru of Iowa State University, Steve Gordon of the Ohio Supercomputer Center, Tom Jones of the University of Minnesota, Joe Paris of Northwestern University, and Padma Raghavan of The Pennsylvania State University.

Events of Interest

- **USA Science & Engineering Festival**
  October 10-24, Washington, D.C., details at:
  www.usasciencefestival.org/
- **Workshop on Petascale Computing and Personalized Medicine**
  October 17-19, Urbana, Illinois, details at:
  www.ncsa.illinois.edu/Conferences/PCPM/index.html
- **Fall PRAC Workshop**
  October 18-21, Urbana, Illinois
- **SC10**
  November 13-19, New Orleans, Louisiana, details at:
  http://sc10.supercomputing.org
- **Blue Waters User Group Meeting**
  November 18, 1-5 pm,
  at SC10 Meeting, New Orleans, Louisiana
- **Webinar: Using Eclipse**
  December 10, 10-11 am
  Save the date, details to come

About the Blue Waters Project

Blue Waters is expected to be the most powerful supercomputer in the world for open scientific research when it comes online. It will be the first system of its kind to sustain 1 petaflop performance on a range of science and engineering applications. Scientists will create breakthroughs in nearly all fields of science using Blue Waters.

www.ncsa.illinois.edu/BlueWaters

Blue Waters Project Team

Project Director
Thom Dunning | tdunning@ncsa.illinois.edu

Deputy Project Director
Bill Kramer | wkramer@ncsa.illinois.edu

Senior Project Manager
Cristina Beldica | cbeldica@ncsa.illinois.edu

Chief Software Architect
Marc Snir | snir@illinois.edu

Chief Applications Architect
Bill Gropp | wgropp@illinois.edu

Chief Hardware Architect
Wen-mei Hwu | w-hwu@illinois.edu

IBM Systems Technical Program Manager
Bill Kramer | wkramer@ncsa.illinois.edu

Applications Technical Program Manager
Bob Fiedler | rfiedler@ncsa.illinois.edu

Software Technical Program Manager
Brett Bode | bbode@ncsa.illinois.edu

Supporting Hardware Systems Technical Program Manager
Mike Showerman | mshow@ncsa.illinois.edu

Education, Outreach & Training Technical Program Manager
Scott Lathrop | scott@shodor.org

Industrial Engagement
Merle Giles | mgiles@ncsa.illinois.edu

Building & Facilities Technical Program Manager
John Melchi | jmelchi@ncsa.illinois.edu

Director, Virtual School of Computational Science and Engineering
Sharon Glotzer | sglotzer@umich.edu

President, Great Lakes Consortium for Petascale Computation
Maxine Brown | maxine@uiuc.edu