World's Largest Coherent Shared Memory System Used for Multitude of Research Projects

Flexibility, Power and Ease of Use Enable Amazing Insights Quickly

Key Facts

Organization:
Pittsburgh Supercomputing Center

Locations:
Pittsburgh, PA

Application:
Higher Education & Research

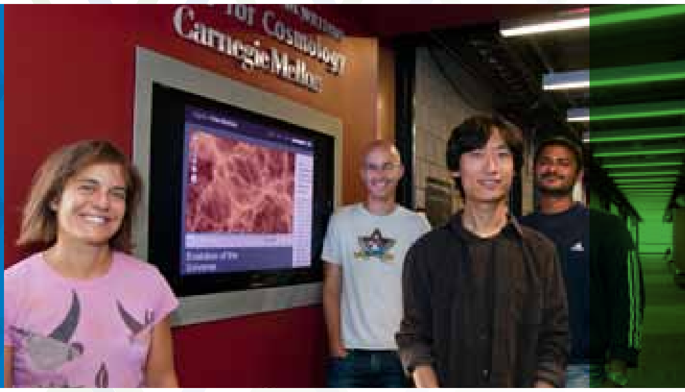

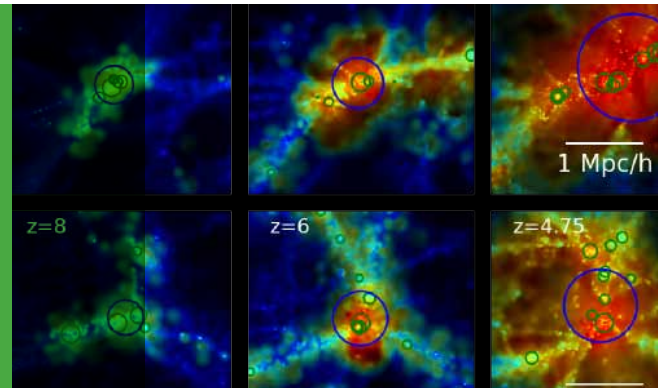

Aided by a $2.8 million award from the National Science Foundation (NSF), and offered as a unique computational resource through the Extreme Science and Engineering Discovery Environment (XSEDE) program, the installation of "Blacklight" is enabling fantastic discoveries by researchers. This system, which the Pittsburgh Supercomputing Center (PSC) acquired in July 2010, features the SGI scalable, shared memory computing platform and associated disks. The SGI UV 1000 system's extremely large, coherent shared memory opens new computational capability for U.S. scientists and engineers.

Hardware and Software Combination Makes Running Jobs Easy

Blacklight is an SGI UV 1000 shared memory system comprising 256 blades. Each blade holds two Intel® Xeon® X7560 eight-core processors, for a total of 4,096 cores across the entire machine. Each core has a clock rate of 2.27GHz, supports two hardware threads and can perform nine Gflops. Thus, the total floating point capability of the machine is 37 Tflops. The 16 cores on each blade share 128GB of local memory, giving each core 8GB of memory with a total machine capacity of 32TB, though the machine is also partitioned such that it can also look like two 16TB machines. Blacklight runs on SUSE® Linux® Enterprise Server. Enhancements to the operating system that were necessary to enable a machine like Blacklight have been made available to the community to benefit from SGI and PSC's shared experience.

The Intel® C, C++ and Fortran compilers and the GNU Fortran, C and C++ compilers are installed on the system, as are the facilities to enable researchers to run threaded, MPI and hybrid threaded programs. OpenMP programs are commonly run on the system. UPC and Java are also available on the machine. Various performance tools, debuggers and libraries are also available.

Blacklight Memory Advantage Program[1]

"Shared memory" means that a system's memory can be directly accessed from all of its processors, as opposed to distributed memory, where each node's memory is directly accessed only by the processors on that node. Because all processors can share and access a single copy of data, a shared memory system is generally easy to program and use.

To help researchers take advantage of Blacklight, PSC provides a Memory Advantage Program to develop applications that can effectively use Blacklight's shared memory capabilities. These include rapid expression of algorithms, such as graph-theoretical software, for which distributed memory often presents obstacles, and interactive analysis of large data sets, which often can be loaded in their entirety into Blacklight's shared memory. For such projects, a PSC consultant can provide advice on debugging, performance-analysis and optimizations.

[1] http://psc.edu/publicinfo/news/2011/110911_BLProgress.php

Researchers from All Over Taking Advantage of the System

“Because of the extraordinary memory size and relative ease of programming made possible by this system’s shared memory structure, scientists and engineers are able to solve problems that were heretofore intractable,” says PSC scientific directors Michael Levine and Ralph Roskies in a joint statement. “For many research communities, including data analysis and many areas of computer science, it has opened the door to the use of high-performance computation, thereby expanding the abilities of scientists to ask questions and gain insight.”

As of early 2012, a total of 1,316 users are using Blacklight for 373 projects in machine learning, natural language processing, software security, analysis of molecular biology simulations, extreme-scale performance engineering, chemistry, fluid dynamics, the early universe, condensed matter, seismic analysis, nanomaterials, astrophysics, geophysics, climate modeling and genomics. “Researchers are extremely excited about what they are achieving with 16 terabytes of coherent shared memory, or even accessing all 32 terabytes as can be accomplished with a bit more programming,” says Nick Nystrom, PSC’s director of strategic applications. “Also revolutionary is that Blacklight offers unparalleled ease of use for rapidly testing new ideas and for expressing complex algorithms, dramatically increasing users’ productivity.”

Blacklight Goes to Work[2]

The world’s largest shared memory system, Blacklight has rapidly proven itself as a productive tool in research across a range of fields, and has opened new capability for U.S. scientists and engineers. Blacklight became available for NSF allocations in October of 2010. “As we expected it would, Blacklight has rapidly become a force across a wide and interesting spectrum of fields.” said Levine and Roskies.

For astrophysicists Tiziana Di Matteo and Rupert Croft, Blacklight has revolutionized discovery from large-scale simulations of how the cosmos evolves. The ability to hold an entire snapshot of their MassiveBlack simulation — between three and four terabytes of data — in memory at one time was instrumental in their ability to reveal “cold gas flows” as a phenomenon that accounts for supermassive black holes in the early universe, resolving what had been a puzzle in the Cold Dark Matter model of the universe.[3]

Also processed on Blacklight is a simulation entitled “Flux Ropes,” by Homa Karimabadi of the University of California San Diego and collaborators. “Blacklight’s shared memory architecture,” says Karimabadi, “is critical for analysis of these massive datasets.” In this large geophysics project, a team of physicists used Blacklight to produce scientific visualizations that made it possible to see a fundamental phenomenon of space weather called magnetic reconnection, which can disrupt satellites, spacecraft and power grids on Earth. The researchers used additional XSEDE resources for very large simulations that characterize how turbulence within sheets of electrons generates structures — called “flux ropes” — that play a large role in magnetic reconnection. “One run can generate more than 200 terabytes,” says Karimabadi. “Blacklight’s shared memory architecture is critical for analysis of these massive datasets.”[4]

In genomics, Blacklight has helped to open a potential bottleneck in processing of next-generation sequencing data. In one project, involving billions of 100-base reads from a sequencer, Blacklight’s shared-memory architecture — along with consulting help from XSEDE’s Extended Collaborative Support Services staff — made it possible to complete a de novo assembly in weeks, progress that had eluded James Vincent of the University of Vermont and colleagues in the Northeast Cyberinfrastructure Consortium for nearly a year in work with other systems.[5]

With limitless quantities of text available on the World Wide Web, Blacklight’s shared memory provides a powerful tool for natural language processing (NLP) — sifting through billions and billions of words in various applications, including automated translators, and innovative predictive modeling. Noah Smith of Carnegie Mellon University produced four studies in diverse areas of NLP within six months of access to Blacklight. “Blacklight has been a very useful resource for us,” says Smith. “We can incorporate deeper ideas about how language works, and we can estimate these more complex models on more data.”[6]

PSC Collaborates with SGI on Shared Memory Enhancements

As a resource of XSEDE, the NSF cyber-infrastructure program, Blacklight serves researchers nationwide across many fields of science and engineering. To optimize Blacklight’s productivity for these researchers, PSC staff, under the lead of J. Ray Scott, Director of Systems & Operations, worked closely with SGI to enhance a number of Blacklight features. These enhancements include:

- Substantially reducing the boot time of the system
- Enabling high input-output file performance for very large datasets without the need to change application software
- Developing new tools to diagnose research application problems — to optimize job layout among processors and to track and isolate hardware problems
- Modifying the portable batch system scheduler so that resources assigned to one shared-memory job don’t interfere with other jobs
- Enabling researchers to use Zest, a high-speed memory snapshot utility, and SLASH2, a wide-area data-replicating file system, to ease the burdens of data management for very large, distributed datasets

[2] http://psc.edu/publicinfo/news/2011/110911_BLProgress.php
[3] http://www.psc.edu/science/2011/supermassive
[4] http://www.psc.edu/science/2011/inprogress/#solarwind
[5] http://www.psc.edu/science/2011/sequencing
[6] http://www.psc.edu/science/2011/language

According to PSC and SGI officials, many of these enhancements will benefit not only XSEDE and the researchers who use Blacklight, but also the research community at large, by enabling improved performance for all large-configuration SGI UV 1000 systems. "We've been at this for 25 years," said PSC scientific directors Roskies and Levine, "and we've become expert at working out the sometimes perplexing problems that go along with bringing a new system through the process of shakedown and final testing to provide researchers with a maximally productive tool."

About XSEDE

The Extreme Science and Engineering Discovery Environment (XSEDE) is the most advanced, powerful, and robust collection of integrated advanced digital resources and services in the world. The five-year, $121-million project is supported by the National Science Foundation, and replaces and expands on the NSF TeraGrid project. XSEDE lowers technological barriers to the access and use of computing resources. Using XSEDE, researchers can establish private, secure environments that have all the resources, services, and collaboration support they need to be productive. The XSEDE partnership includes over a dozen universities and research centers, and is led by the University of Illinois's National Center for Supercomputing Applications (NCSA).

About the Pittsburgh Supercomputing Center

The Pittsburgh Supercomputing Center is a joint effort of Carnegie Mellon University and the University of Pittsburgh together with Westinghouse Electric Company. Established in 1986, PSC is supported by several federal agencies, the Commonwealth of Pennsylvania and private industry, and is a partner in the National Science Foundation XSEDE program.

For more information about what researchers are doing on Blacklight, specifically in the area of Life Sciences, please review the White Paper at sgi.com/pdfs/4347.pdf

Global Sales and Support: sgi.com/global