Success Story

"We are using the COSMOS supercomputer to develop techniques to analyze data from the world's most ambitious CMB experiment, the Planck Surveyor satellite, due for launch in 2007."

—Stephen Hawking, Professor, Cambridge University

The Challenge

- Create a scientific networking grid that enables U.K. cosmologists at distant locations to share data and computing resources in their study of the origins of the universe
- Establish a central computing and visualization facility and enable remote grid participants to use it to collaborate in real time by interacting visually with large cosmological data sets

The Solution

- A 64-processor SGI® Origin® 3800 supercomputing system with 32GB of shared memory, connected through a SAN switch to a 1.6TB SGI® TP9100 RAID
- A dedicated 8-CPU SGI® Onyx® family visualization facility
- A 64-processor Beowulf cluster with 1 GHz Pentium® III processors
- SGI® Origin® 300 servers at eight or more U.K. institutions, enabled to collaborate over a Visual Area Network (VAN) using OpenGL Vizserver and broadband facilities

The Result

- CosmoGrid has established its proof of concept through cross-Atlantic demonstrations at scientific meetings, enabling researchers to collaborate in real time by jointly manipulating visualizations of cosmological data
- UK cosmologists at widely separate locations are now able to share compute and visualization resources through the central Cambridge facility

Stephen Hawking, SGI, and CosmoGrid: The Collaborative Search for Beginnings

The most basic question in science is, How did it all begin? The search for the origins of the universe has led cosmologists to the Big Bang theory, powerfully reinforced by the discovery of uniform cosmic microwave background (CMB) radiation throughout the universe. Today cosmology research focuses on reconstructing the critical fractions of a second that followed the Big Bang. U.K. researchers, led by principal investigator Professor Stephen Hawking of Cambridge University, have combined forces to attack this compute- and visualization-intensive problem with the U.K. Cosmology Grid (CosmoGrid).
The heartbeat of CosmoGrid is an SGI Origin 3800 supercomputer.

CosmoGrid: Linking Hands to Answer the Universe's Biggest Question

Widely accepted theory states that the Big Bang created an expanding universe, much as an aerial fireworks explosion creates a globular pattern of light. But to produce the universe we inhabit, there must have been subtle variations, and they would be reflected in the CMB—the cool, snapshot relic of an unimaginably hot and brilliant event. These variations showed up, as predicted by theory, when the COBE satellite was launched in 1989. "The COBE satellite images were the first to show that on top of the smooth CMB, you began to see these small ripples," said Dr. Paul Shellard, director of the COSMOS project. "These primordial ripples, or perturbations, evolved into the structures of the universe—galaxies, stars, etc."

Cosmologically, these ripples are of intense interest. Just what happened during the first few ticks of the cosmic clock? To test competing theories, the COSMOS project aims at nothing less than modeling the history of the universe from those primordial picoseconds over 10 billion or 20 billion years to our present day.

Like the Human Genome Project, the COSMOS project is of staggering proportions, but not beyond comprehension. If many hands make light work, many minds—and many computing cycles—can simplify this enormous challenge. The purpose of CosmoGrid, a collaborative computational and visualization network, is to accelerate COSMOS by making more data, computing power, and visualization power accessible to U.K. cosmologists.

CosmoGrid, funded by the U.K. Higher Education Funding Council, is a joint initiative by SGI, Platform Computing, and the U.K. National Cosmology Consortium with additional support from the U.K e-Science program. The focal point of CosmoGrid is the COSMOS supercomputer, a 64-processor Origin 3800 system in the Department of Applied Mathematics and Theoretical Physics at Cambridge University. This central system is linked to SGI® systems at the universities of Sussex, Manchester, Oxford, and Portsmouth and at the Imperial College in London. Platform Computing, the leader in distributed computing software, has customized its grid computing software solutions, including Platform LSF, Multicluster, and Platform Globus™, to integrate CosmoGrid's distributed IT infrastructure and enhance its efficiency.

Proof of Concept: Global Grid Collaboration

Cosmologists at collaborating U.K. universities can now use CosmoGrid access to store and process data either on their local facilities or, if more appropriate, on the hierarchical data storage management system of the Origin 3800 server at Cambridge, which includes an SGI TP9100 Fibre Channel RAID. They can also collaborate visually with scientists at other CosmoGrid institutions. This project is exciting for SGI because it exploits all the major strengths of SGI technology: collaborative visualization, management of complex data, and distributed high-performance computing.

Evolution of COSMOS/CosmoGrid Technology

The original COSMOS platform, installed in 1997, was a 32-processor SGI® Origin® 2000 server. In the subsequent years a number of upgrades enhanced the system by adding a graphics pipe, which converted the server to an SGI Onyx family system, and increasing the number of processors to 64. But this was only the beginning. Under a new five-year agreement with Professor Hawking, SGI has recently installed a new 64-processor Origin 3800 server with 32GB of shared memory. This is complemented by:

- A dedicated 8-processor Silicon Graphics® Onyx2® visualization facility
- An SGI® Data Migration Facility on an SGI Origin 300 system and Tape Management Facility to connect to a 4TB StorageTek® Tape Library
- Connecting through a SAN switch to a 1.6TB SGI TP9100 RAID storage array
- Installing the SGI® CXFS™ clustered filesystem
- Installing a 64-processor Beowulf cluster with 1 GHz Pentium III processors
- Installing OpenGL Vizserver
- Installing SGI Origin 300 servers in at least eight other U.K. universities

By February 2003, SGI is planning to complete a remarkable series of upgrades:

- Replacing the 64-processor Origin 3800 server with a 128-processor Origin 3800 with faster processors
- Increasing the size of the Beowulf cluster through SGI Professional Services integration
- Upgrading the Onyx family system

The consortium has already performed proof-of-concept CosmoGrid demonstrations using collaborative visualization technology. During Supercomputing 2001, members of the consortium, working at the Cambridge facility, demonstrated and discussed their work on galaxy formation and the physics of the early universe with conference attendees through daily real-time fly-throughs of cosmological data. The Cambridge facilities, using NCSA's Virtual Director, were linked with the National Center for Supercomputing Applications (NCSA) in Illinois, the Astrophysics Department at the University of Hawaii, and the SGI booth at the conference in Denver, sharing large data sets among SGI Onyx family systems using the grid-enabled navigation and collaboration tools. A similar demonstration linking Cambridge, NCSA, and the SGI booth was conducted July 23–25 at SIGGRAPH 2002.

Ahead: Detailed Data and More Computing Cycles

One major source of frustration for COSMOS investigators—low-resolution CMB data—is slowly going away and will disappear in about five years. COBE was followed by NASA's MAP satellite (June 2001), with 100 times the resolution of COBE. The announcement of the first MAP data is expected by the end of 2002. In the meantime, cosmologists are preparing simulations based on what they believe MAP's data will reveal. All this compute-intensive work is in preparation for a major event that will make CMB data available at twice the resolution of MAP, with a corresponding need for many more computing cycles.

"We are using the COSMOS supercomputer to develop techniques to analyze data from the world's most ambitious CMB experiment, the Planck Surveyor satellite, due for launch in 2007," says Professor Hawking. "This analysis is very computationally challenging, but it will give us unprecedented new information about the origin and present state of our universe.

"Without the computing grid provided by SGI hardware and Platform software, we would be unable to do the calculations required or to collaborate effectively between our remote sites. CosmoGrid will bring together and greatly enhance the U.K.'s leading research in computational cosmology."

Meeting the Challenges of Remote Visualization

CosmoGrid faces challenges that any grid-based, shared-resources program will face: preserving a secure collaboration environment, administering remote job submissions, allocating computing resources, and achieving distributed load balancing, to name a few. The Platform Computing grid software suite will take care of all these management needs for

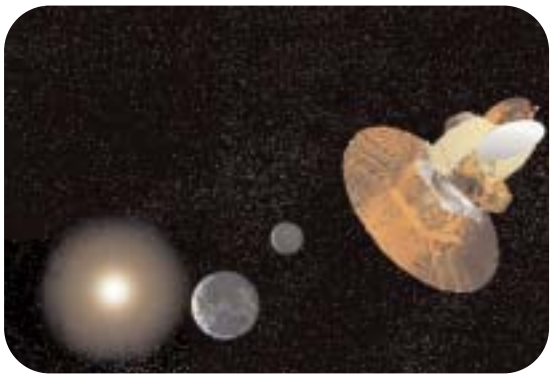

the consortium, making the system transparent to remote users.

Data for visual collaboration sessions will be processed and rendered on the SGI systems at Cambridge, then served to remote clients. The remote clients will use SGI workstations to access visual data, manipulating it as though the application on the server was running on the local workstation.

CosmoGrid is an outstanding example of the basic benefit of grid computing—enabling a community of users to access resources transparently over a network. It is also an example of the way SGI and Platform Computing have collaborated to provide the infrastructure for major scientific grid computing projects around the world.

"At SGI, it is our mission to provide the computing power that underlies all significant scientific and creative progress in the digital world of the 21st century," says SGI CEO and Chairman Bob Bishop. "We are deeply supportive of Professor Hawking's work to model the universe on SGI machines and consider this endeavor as one of the greatest challenges our equipment will ever face. Teaming up with Platform Computing has been an additional bonus."

An alliance between SGI and Platform, announced in July 2002, will enable the two companies to develop innovative solutions that integrate data, compute, and collaborative visualization resources. This includes Visual Area Networking, a concept introduced by SGI, which allows scientists and engineers at any location to interact remotely with high-performance interactive visualizations, by themselves or collaboratively, using any client device.

"CosmoGrid will demonstrate the power of these technologies—grid computing and Visual Area Networking—on the grandest scale possible," says Bishop. "We welcome challenges like these. And we're delighted to be here."

"This analysis is very computationally challenging, but it will give us unprecedented new information about the origin and present state of our universe."

—Stephen Hawking,
Professor,
Cambridge University

Corporate Office
1600 Amphitheatre Pkwy.
Mountain View, CA 94043
(650) 960-1980
www.sgi.com

North America 1(800) 800-7441
Latin America 1(650) 933-4637
Europe (44) 118.925.75.00
Japan (81) 3.5488.1811
Asia Pacific (65) 771.0290

3350 (12/06/2002)

JI4081