NERSC NEWS
NATIONAL ENERGY RESEARCH SCIENTIFIC COMPUTING CENTER
BERKELEY LAB
Office of Science
U.S. DEPARTMENT OF ENERGY
JULY 2007

Analyze This

A PERFORMANCE ANALYSIS SOFTWARE CREATED BY NERSC TO BE USED IN NSF CENTERS

The National Science Foundation (NSF) recently approved a proposal that will deploy a nimble performance evaluation tool, developed by NERSC's David Skinner, on all major NSF supercomputers.

The software, Integrated Performance Monitoring (IPM), analyzes the performance of HPC applications and identifies load balance and communication problems that prevent them from running smoothly and achieving high performance. IPM is easy to deploy and use in systems with thousands or tens of thousands of processors, making it a good tool for petascale computing.

Skinner, leader of NERSC's Open Software and Programming Group, developed IPM in 2005. Since then, the software has won fans beyond NERSC, including the San Diego Supercomputer Center, the Center for Computation and Technology at Louisiana State University, the Swiss National Supercomputing Center (CSCS) and the DOD's Army Research Laboratory.

continued on page 2

Expanding Praises

A $500,000 PRIZE SHARED BY RESEARCHERS WHO USED NERSC FOR ACHIEVING A COSMOLOGY BREAKTHROUGH

The 2007 Gruber Cosmology Prize was awarded this month to researchers who relied on NERSC for producing the scientific breakthrough that led to the recognition.

The $500,000 prize, given by the Peter and Patricia Gruber Foundation, represents the latest kudo received by two international teams that independently discovered the accelerating expansion of the universe. The discovery, announced by the teams in 1998, refuted a well-accepted theory that the pace of the expansion was slowing.

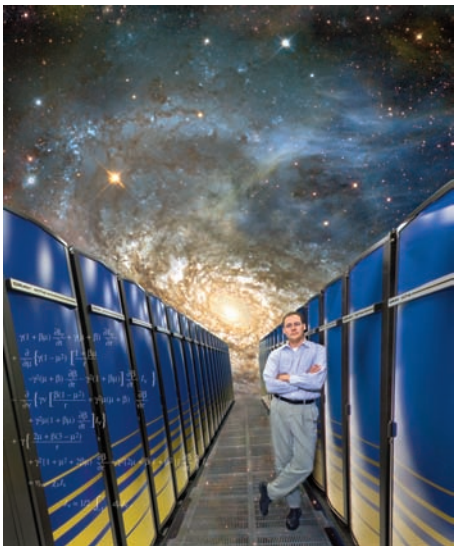

Peter Nugent's supernova simulations helped lead to a breakthrough in cosmology.

Berkeley Lab's Saul Perlmutter led one of the teams and used NERSC to analyze supernova data. Science magazine hailed the work, as part of the Supernova Cosmology Project, by Perlmutter and his team as the "Breakthrough of the Year for 1998."

"The Gruber Cosmology Prize is a great honor for all of us in both teams," says Perlmutter, who continues to lead the Supernova Cosmology Project today. "It's rare that a scientific prize is able to include a very large number of those in the community whose work actually became the underpinnings of the discovery being celebrated. Yet with the exception of a few notable contributors, that's the case here — and that's terrific."

Peter Nugent, an astrophysicist in Berkeley Lab's Computational Research Division, was responsible for running the

continued on page 4

Move Up the Scale

SCIENTISTS BENEFIT FROM CPU-HOUR REIMBURSEMENT PROGRAM TO IMPROVE CODE PERFORMANCE

NERSC's large-scale reimbursement program has provided nearly six million computing hours so far this year to 21 projects that have taken advantage of the opportunity to scale their runs in preparation for using the new Cray XT4 system.

The program has set aside nine million computing hours on Seaborg this year so far to encourage researchers to improve their codes for large runs on the new Cray, Franklin, once it enters production later this year.

The incentive has attracted projects from a variety of scientific disciplines, including astrophysics, life sciences, fusion, chemistry and climate research. Scientists say the program has enabled them to pinpoint and resolve issues before running jobs on Franklin.

"We have participated in the Scaling Reimbursement program in order to tackle the problem of testing the predictive power of new empirical force fields for biomolecular simulation. Our system of interest is the Abeta peptide and various sub-peptides which are associated in the formation of amyloid plaques in Alzheimer's Disease,"

continued on page 3

Earth Citizen

BERKELEY RESEARCHER WINS AWARD FOR TRANSFORMING A GLOBAL CLIMATE MODEL

Inez Fung, a climate expert and NERSC user, recently received the 2007 CCSM Distinguished Achievement Award for her role in shaping a modeling system that enables scientists to carry out comprehensive analyses of the earth's climate.

Fung, a researcher at Berkeley Lab's Earth Sciences Division and a UC Berkeley professor, was given the plaque

continued on page 4

Space, a New Frontier

AN UPGRADE TO NERSC'S STORAGE TECHNOLOGY WILL PROVIDE MORE ROOM FOR CUTTING-EDGE SCIENCE

NERSC has installed new tape drives that provide 2.5 times more capacity and four times the performance of the previous tape drives, a move to meet scientists' increasing demand for more storage.

After months of work, Wayne Hurlbert and his colleagues in the Mass Storage Group recently put to work the Titanium 10000A cartridges, made by Sun Microsystems. The upgrade is part of an ongoing effort to expand the storage capability at NERSC, where scientists increasingly are carrying out large-scale experiments and more detailed simulations. For example, the DOE's Joint Genome Institute received 1.5 million computing hours and 900,000 storage resource units this year for its sequencing work.

Each cartridge can hold 500 gigabytes of uncompressed data. With compressed data, each tape can hold 750 gigabytes or more. As a result, NERSC's 44,000 tape cartridge capacity can now provide 22 petabytes of storage for uncompressed data and 33 or more petabytes with compression.

The table below shows efforts by the Mass Storage Group to improve storage capabilities over the years. All the storage technology came from the company StorageTek, which was bought by Sun in 2005.

Technology	Year Deployed	Data Rate	Cartridge Capacity	Total Storage (uncompressed)	Total Storage (compressed)
T9840A	1999	10 MB/s	20 GB	0.88 PB	1.3 PB
T9940A	2002	30 MB/s	60 GB	2.6 PB	3.9 PB
T9940B	2003	30 MB/s	200 GB	8.8 PB	13.2 PB
T10000A	2007	120 MB/s	500 GB	22 PB	33 PB

IPM *continued from page 1*

IPM overcomes shortcomings exhibited by other performance analysis software, Skinner said. For example, IPM has low overhead and requires no source code modifications, making it easy for researchers to use. Its fixed memory footprint also ensures that running the software won't negatively impact the applications being profiled.

"An understandable application performance profile is something that all researchers using parallel computing resources should expect in-situ via a simple flip of a switch. It should not require additional effort of changing their code" Skinner said.

In the NSF proposal, Skinner reviewed real-life cases of DOE and NSF supercomputer centers using various performance monitoring tools. The principal investigators for the project are Kathy Yelick of UC Berkeley and Berkeley Lab, Allan Snavely and Nick Wright, both of the San Diego Supercomputer Center.

"Some means of doing performance analyses are quite invasive and disturb the application one is trying to study; others are more lightweight but don't provide adequate information to researchers to improve their codes. Some require all users of a system to actively participate in the profiling activities; others are more passive, operating in the background. Some scale to thousands of tasks and some do not," said Skinner.

The comparisons enabled the researchers to convince NSF to deploy IPM in all of its supercomputer centers. NSF has awarded $1.58 million for the project, which is scheduled to begin in August 2007 and take place over three years.

Part of the project will focus on expanding IPM's capabilities, such as broadening the scope of what is profiled and improving data analysis, Skinner said. The software is available under an open source software license and can run on major supercomputer architectures today: IBM, Linux clusters, Altix, Cray X1, Cray XT4, NEC SX6 and the Earth Simulator.

Learn more about IPM at http://ipm-hpc.sf.net.

Reimbursement

continued from page 1

said Nicolas Lux Fawzi, a UC Berkeley scientist on a research team led by Teresa Head-Gordon at Berkeley Lab. "We have used the CPU time in the program to demonstrate that we can run large parallel simulations on 1,024 processors using the replica exchange technique to generate the complete equilibrated ensemble for our system at a range of temperatures.

"We have very much enjoyed the chance to work with NERSC as part of the scaling program. We've received some excellent input from the NERSC staff on how to evaluate and improve the scaling of the code."

Franklin, which arrived at NERSC earlier this year, has more than 19,000 processors and can deliver 10 times more computing power than any existing NERSC system. Franklin can provide a sustained performance of at least 16 trillion calculations per second, with a theoretical peak speed of more than 100 teraflop/s.

The reimbursement program is open to all NERSC users and requires researchers to run 1,024- to 1,500-processor jobs on Seaborg, depending on whether they have participated in the program in the past. The target is to carry out a run on at least 2,418 CPUs. Each project can get a maximum reimbursement of 500,000 hours this year. Scientists also have to use the Integrated Performance Monitoring (IPM) software to gather performance information about each run.

"The quantum Monte Carlo methods developed in the Lester group are naturally amenable to parallel computing. Historically, our production jobs have run on several hundred processors with near perfect parallel efficiency," said Brian Austin, a researcher in a group led by William A. Lester, Jr., a UC Berkeley chemistry professor and Berkeley Lab researcher. "The advent of near-petascale computers such as Franklin will bring jobs with thousands of processors into the norm. In this regime, subtle changes to our mode of parallel communication have dramatic effects that were unnoticeable at previous scales: communication time increased from 2 percent to almost 50 percent as the number of processors increased from 512 to 2,048. The reimbursement program has been essential to our exploration and resolution of these issues."

Don Lamb, a University of Chicago scientist who leads a supernovae research project, also is an active participant in the reimbursement program. Lamb is a recipient of this year's Innovative and Novel Computational Impact on Theory and Experiment (INCITE) awards, a program by the DOE Office of Science to support large-scale projects in national labs, universities and industry.

Lamb's research team members said the IPM wasn't easy to use initially, but they overcame those issues with the support of the NERSC staff.

"We had to pay more attention to exit codes issued by FLASH than we had previously, since non-zero exit codes force IPM to throw away all output. But the NERSC staff was helpful and understanding of IPM issues, never letting missing IPM output interfere with reimbursement. Where available, the IPM output supplied interesting profiling information," said Carlo Graziani, a researcher on Lamb's team.

Other scientists whose projects are among the top 10 recipients of reimbursed hours are George Vahala (fusion plasma), Doug Toussaint (quantum chromodynamics), Stephen Gray (chemistry – nanoscale electrodynamics), Cameron Geddess (accelerator physics), Wei-li Lee (fusion plasmas) and Paola Cessi (climate research).

More information about the reimbursement program can be found at http://www.nersc.gov/hypermail/all-announcements/0755.html.

SPOTLIGHT

EVERYTHING YOU KNOW IS WRONG

John Shalf and **Erich Strohmaier** last month spoke at the International Supercomputing Center, which broke the attendance record this year by attracting more than 1,200 people from 47 countries to the conference in Dresden, Germany.

John Shalf

Erich Strohmaier

Shalf, head of the Science-Driven System Architecture Team at NERSC and an invited speaker, gave a talk on "Overturning the Conventional Wisdom for the Multicore Era: Everything You Know is Wrong." In an abstract about the talk, Shalf said, "As the pace of processor clock rate improvements continues to slow, the microprocessor industry has moved to doubling the number of processor cores per chip every 18 months. Consequently, the path towards realizing petascale computing depends on riding a wave of exponentially increasing system concurrency. This is leading to reconsideration of interconnect design, memory balance, and I/O system design because our entire software infrastructure is built upon assumptions that are no longer true. This will have dramatic consequences for the design of future HPC applications and algorithms."

Strohmaier, a researcher at NERSC and the Computational Research Division at Berkeley Lab, presented the semi-annual Top500 list of the world's fastest supercomputers. Strohmaier is a co-founder of Top500, which issued its first ranking in 1993. More information about Top500 can be found at http://www.top500.org/.

Strohmaier and Shalf also organized a discussion called "Energy Consumption of Supercomptuers" at the conference. Participants traded ideas on using new performance and power-related metrics to rank supercomputers. Strohmaier noted, in an abstract about the discussion, that "For decades, the notion of 'performance of computer' has been synonymous with 'raw speed' as measured in FLOP/S. This isolated focus has ultimately led to supercomputers that consume egregious amounts of electrical power. Other performance metrics have been largely ignored, e.g., power efficiency, space efficiency, reliability, availability, and usability. As a consequence, the total cost of ownership (TCO) of a supercomputer has increased extraordinarily. To

continued on page 5

Gruber Prize

continued from page 1

simulations on NERSC's 512-processor Cray T3E-900 back then. The computing power enabled Nugent to analyze data from 42 supernovae, including the simulation of 10,000 exploding stars at varying distances. Nugent compared with data from the simulations with those from observations to check for errors.

"It's so nice to have a result where you can sit down with a fifth grader and they can understand it. You don't have to teach them general relativity to explain to them how you measure distant supernovae and find out whether they are slowing down or speeding up," Nugent said. "As for the Gruber prize itself, it's wonderful that they are rewarding the whole team — that's a rare thing in the world of scientific prizes."

The Gruber prize also went to the High-Z Supernova Search Team, headed by Brian Schmidt of the Australian National University.

The announcement by both teams in early 1998 that the expansion of the universe is not slowing, as almost all astronomers had expected, but instead is expanding ever faster led to the realization that the universe is dominated by a mysterious something that stretches space and works against the mutual gravitational attraction of ordinary matter and energy. Dark energy, as it was soon named, is now believed to constitute some three-quarters of the density of the universe, with dark matter making up most of the rest.

Astronomers had first proposed using supernovae to measure expansion in the 1930s, but it was not until the mid-1980s that bright Type Ia supernovae were determined to be highly dependable "standard candles," meaning they are visible and can serve as a yardstick for measuring distances. By comparing the distance of these exploding stars with the redshifts of their home galaxies, researchers can calculate how fast the universe was expanding at different times in its history.

The Gruber Cosmology Prize citation underscores the teams' application of "new techniques that use supernovae exploding within distant galaxies to measure precise distances across a large fraction of the observable Universe."

The Supernova Cosmology Project, which continues to use NERSC resources, contributed several such innovations, including the key technique that, by 1994, allowed them to prove that they could schedule telescope time in advance and guarantee "supernovae on demand." To do this they used a new observing approach that reliably produced batches of high-redshift Type Ia supernovae, comparing pairs of matching images taken three weeks apart — the latter just before the new moon — and each pair including many target galaxies. Every observing session resulted in the identification of numerous Type Ia's caught while still growing brighter.

Learn more about the Supernova Cosmology Project at http://www.supenova.lbl.gov/public.

Inez Fung

continued from page 1

by the Community Climate System Model (CCSM) organization during its annual workshop in June. CCSM is an organization supported by the National Science Foundation and the DOE and run by the National Center for Atmospheric Research in Colorado.

The award recognized Fung's leadership in founding the Biogeochemistry Working Group in 1998 and steering its direction to coordinate research and develop biogeochemical components for the CCSM. Fung provided "the intellectual guidance and energy behind the transformation of the CCSM from a physical climate model into a comprehensive earth system model," said Peter Gent, Chair of the CCSM Scientific Steering Committee. "These capabilities allow us to explore a whole new range of critical interactions and feedbacks between atmospheric composition, ocean and land biosphere and climate."

The National Center for Atmospheric Research started the CCSM in 1983 as a global atmosphere model that focused primarily on the physics and mathematics of the earth in studying climate. Starting in the mid 1990s, however, scientists began to expand the model to include other elements that shape the global climate, such as oceans and sea ice. The CCSM then formed working groups within each discipline together with the Scientific Steering Committee to undertake long-term projects on further broadening the scope and applications of the model.

In her work leading the Biogeochemistry Working Group, Inez contributed to the launch of an interactive carbon cycle model within the CCSM and shepherded the new project (in collaboration with Oak Ridge National Laboratory and Lawrence Livermore National Laboratory) to prototype a new system to diagnose verisimilitude of the carbon simulations.

Inez Fung receives the 2007 CCSM Distinguished Achievement Award.

Fung said NERSC's support has been essential for her work at the CCSM.

"Computing at NERSC has been indispensable for testing new ideas about the climate system, and for demonstrating their significance by large numbers of long-inte-grations," Fung said.

She currently is a principal investigator of a project being carried out at NERSC. The project, "Carbon-Climate Interactions," aims to develop a more accurate modeling of carbon dioxide in the atmosphere and its impact on the climate. In previous research, Inez's team specified fossil fuel emission scenarios in their model. The current work will take into account new scenarios related to the growing population demand for food and water.

The research also will continue to broaden the biogeochemical scope of the CCSM by adding components such as water isotopes and mineral aerosols. Fung is the co-PI on a newly funded DOE project to assimilate large volumes of satellite observations of CO2 into the carbon-climate model to yield the first global observation of the geographic variations of atmospheric CO2. For more information about Fung's research, check out http://www.atmos.berkeley.edu/~inez.

SPOTLIGHT

raise awareness and facilitate discussion of these issues, developing new ranking systems that make use of alternative performance metrics would be helpful."

MAKING PETASCALE WORK

John Shalf and **Jonathan Carter** last month spoke at SOS11, a workshop that brought together national lab representatives, researchers and supercomputer vendors in addressing hardware, algorithm and software development challenges for petascale computing.

Shalf, head of NERSC's Science-Driven System Architecture Team, was invited to talk about how memory bandwidth affects multicore processor performance. Shalf's talk can be found at http://www.csm.ornl.gov/workshops/SOS11/presentations/j_shalf.pdf.

Carter, head of the User Services Group, gave a presentation about NERSC, describing the center's computing resources and the scientific achievements attained by its users. More information about Carter's talk is at http://www.csm.ornl.gov/workshops/SOS11/presentations/j_carter.pdf.

Oak Ridge National Laboratory hosted the workshop, which took place in Key West, Florida. Learn more about the workshop at http://www.csm.ornl.gov/workshops/SOS11/index.html.

SOFTWARE INNOVATION

David Skinner, who leads the SciDAC Outreach Center, spoke at Harvard University's Initiative Innovative Computing, an interdisciplinary research

David Skinner

and development center that focuses on using innovative computing tools to accelerate scientific discovery.

In his talk, which took place in June, Skinner presented resources available through SciDAC (Scientific Discovery through Advanced Computing). DOE funded the Outreach Center last year to create a clearinghouse for SciDAC activities. The center identifies and provides training opportunities and publicizes scientific accomplishments.

Skinner, who leads the Open Software and Programming Group at NERSC, addressed an audience of scientists who were interested in large-scale computing and workflow management over grids, such as the Open Science Grid.

Find Skinner's slide presentation at http://iic.harvard.edu/seminars/index.html.

NEW STAFF FOR SCIDAC OUTREACH

Andrew Uselton joins the Open Software and Programming Group as the first employee staffing the new SciDAC Outreach Center (http://outreach.scidac.gov). The outreach center is an initiative funded by the DOE Office of Science and headed up by David Skinner, head of the software group. The center provides information and support services for the SciDAC community. Uselton's responsibilities include establishing a web-based collaborative environment in which SciDAC funded research teams will have access to source code revision control, messages forums, mailing lists, task lists, an issue tracker and release management.

Andrew Uselton

Prior to joining NERSC, Uselton worked at Lawrence Livermore National Laboratory as a developer in the Production Linux Group in Livermore Computing. In 1999, he was the second employee to join the team newly formed to explore running commodity Linux clusters at the lab. His early work included "powerman," a scalable, distributed power control application (http://sourceforge.net/projects/powerman/). He also worked on the testing and debugging of "petal/frangipani," a parallel file system. More recently, Uselton was deeply involved in the testing and debugging of the Lustre parallel file system. He recently delivered a paper (http://www.linuxclustersinstitute.org/conferences/abstracts.html#IIpT2) recounting his experiences with Lustre during the Blue Gene/L deployment. In addition to the testing and debugging of Lustre, Andrew developed "mib," a parallel MPI I/O benchmark and visualization application (http://sourceforge.net/projects/mibtest), as well as "logjam," a cluster-oriented log abstraction tool (http://sourceforge.net/projects/logjam).

Uselton has a master's degree in computer science from Stony Brook University, where he published research in the areas of programming language semantics and concurrency theory. He received a bachelor's degree in physics from the University of Texas at Austin. For his undergraduate study, he focused on solid state physics and performed laboratory work with ultra-high vacuum equipment and Auger spectroscopy.

WHAT IS NERSC NEWS?

NERSC News publishes every other month and highlights the cutting-edge research performed using the National Energy Research Scientific Computing Center, the flagship supercomputer facility for DOE's Office of Science. NERSC News editor Ucilia Wang can be reached at 510 945-2402 or Uwang@lbl.gov. Find previous NERSC News articles at http://www.nersc.gov/ news/nerscnews.

DISCLAIMER

This document was prepared as an account of work sponsored by the United States Government. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or The Regents of the University of California. Ernest Orlando Lawrence Berkeley National Laboratory is an equal opportunity employer.