

NERSC's QCD Library is a Model Resource

Five years ago, when the first libraries of lattice QCD data were established at NERSC, the idea of making such information easily accessible to researchers at institutions around the world was a bold new approach. In fact, "The NERSC Archive" format, as it is colloquially known, has become a standard for the lattice storage format, which made exchanging lattices between groups much more efficient. The result is that lattice QCD researchers can avoid the time consuming and computationally intensive process of creating their own lattice data by drawing on the research of others.

Now, the concept is going global. With the appearance of Grid-based technologies, the lattice gauge theory community is developing a new protocol and tools for sharing lattice data across the Grid. The International Lattice Data Federation (<http://www.lqcd.org/ildg/wiki/tiki-index.php>) has been formed as part of the SciDAC initiative in the U.S. and similar large-scale computing projects in Europe, Japan, and Australia. Sometime this year the original lattice archive at NERSC will join similar mirrors around the world allowing access via Web service. Researchers will be able to query the archive database and examine the resource catalog of all available lattices in the world, then seamlessly transfer this data to their computing resources distributed on the Grid.

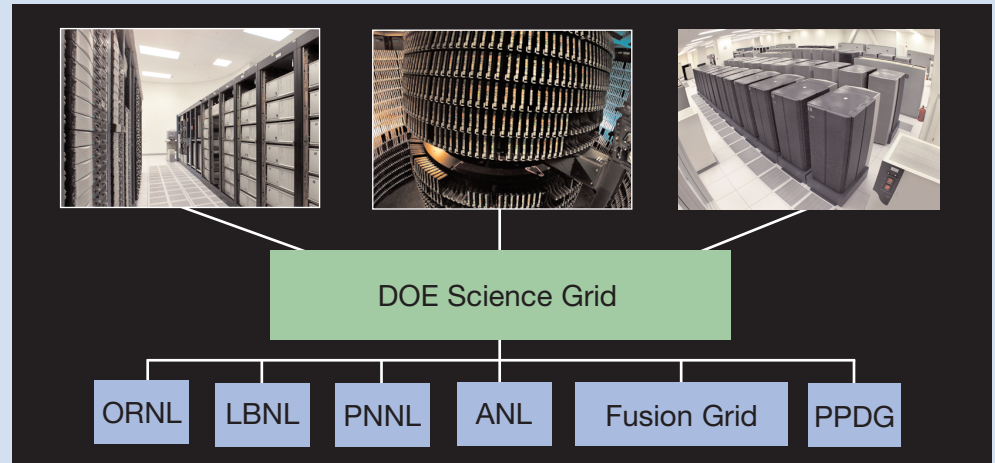
Behind all this are hundreds of gigabytes of nothingness, all created by lattice gauge theorists, the particle physicists who model the behavior of quarks. The data, called "lattices," are snapshots of the way quantum chromodynamic (QCD) fields, which mediate the forces between quarks, fluctuate in a vacuum. Each lattice is a four-dimensional

(continued on page 2)

Introducing NERSC News

NERSC News is a new publication highlighting recent achievements by staff members in Berkeley Lab's National Energy Research Scientific Computing Center Division. Initially, NERSC News will be distributed every other month via email and may be freely distributed. NERSC News is edited by Jon Bashor, JBashor@lbl.gov or 510-486-5849.

All NERSC Production Systems Now on the Grid



With the new year came good news for NERSC users – all of NERSC's production computing and storage resources are now Grid-enabled and can be accessed by users who have Grid applications using the Globus Toolkit.

As a result, Grid users now have access to NERSC's IBM supercomputer ("Seaborg"), HPSS, the PDSF cluster, the visualization server "Escher," and the math server "Newton." Users can also get support for Grid-related issues from NERSC's User Services Group.

"Now that the core functionalities have been addressed, our next push is to make the Grid easier to use and manage, for both end users and the administrators who handle applications that span multiple sites," said Steve Chan, who coordinated the Grid project at NERSC.

One of the major challenges faced by NERSC's Grid team was installing the necessary software on Seaborg, which operates using a complicated software stack. Additionally, the system needed to be configured and tested without interfering with its heavy scientific computing workload. By working with IBM and drawing upon public domain resources, Seaborg became accessible via the Grid in January.

"With Steve Chan's leadership and the technical expertise of the NERSC staff, we crafted a Grid infrastructure that scaled to all of our systems," said Bill Kramer, general manager of the NERSC Center. "As the driving force, he figured out what needed to be done, then pulled together the resources to get it done."

To help prepare users, center staff presented

tutorials at both the Globus World and Global Grid Forum, as well as presented training specifically for NERSC users.

When authenticated users log into NIM (NERSC Information Management), they are now able to enter their certificate information into the account database and have this information propagated out to all the Grid enabled systems they normally access, superseding the need to have separate passwords for each system. The next step is for NIM to help non-credentialed users obtain a Grid credential.

Because the Grid opens up a wider range of access paths to NERSC systems, the security system was enhanced. Bro, the LBNL-developed intrusion detection system which provides one level of security for NERSC, was adapted to monitor Grid traffic at NERSC for an additional layer of security.

The long-term benefits, according to Kramer, will be that NERSC systems are easier to use while security will be even higher.

"Of course, this is not the end of the project," Kramer said. "NERSC will continue to enhance these Grid features in the future as we strive to continue providing our users with the most advanced computational systems and services."

This accomplishment, which also required the building of new access servers, reflects a two-year effort by staff members. Here is a list of participants and their contributions:

Nick Bathaser built the core production Grid servers and has been working on making the myproxy server very secure.

(continued on page 2)

Production Systems (continued from p.1)

Scott Campbell developed the Bro modules that are able to analyze the GridFTP and Gatekeeper protocols, as well as examine the certificates being used for authentication.

Shane Canon built a Linux Kernel module that can associate Grid certificates with processes, and helped deploy Grid services for the STAR and ATLAS experiments.

Steve Chan led the overall project and coordinated the effort, as well as developing and porting much of the actual software that created these abilities.

Shreyas Cholia built the modifications to the HPSS FTP daemon so that it supports GSI, making it compatible with GridFTP. He has also been working on the Globus File Yanker, a Web portal interface for copying files between GridFTP servers.

Steve Lau works on the policy and security aspects of Grid computing at NERSC.

Ken Okikawa has been building and supporting Globus on Escher to support the VisPortal application.

R.K. Owen, Clayton Bagwell and **Mikhail Avrekh** of the NIM team have worked on the NIM interface so that certificate information gets sent out to all the cluster hosts.

David Paul is the Seaborg Grid lead, and has been building and testing Globus under AIX.

Iwona Sakrejda is the lead for Grid support of PDSF and developed training for all NERSC users, as well as supporting all the Grid projects within the STAR and ATLAS experiments.

Jay Srinivasan wrote patches to the Globus code to support password lockout and expiration control in Globus.

Dave Turner is the lead for Grid support for Seaborg users. He recently helped a user, Frank X. Lee, move a few terabytes of files from the Pittsburgh Supercomputing Center to NERSC's HPSS using GridFTP.

QCD Library (continued from p.1)

array ($28^3 \times 96$, say) of four 3×3 complex matrices representing these fields in a tiny box of space measuring about 2 femtometers on a side ($1 \text{ fm} = 10^{-15} \text{ m}$) and extending about 10^{-22} seconds in time. Lattice gauge theorists generate many such lattices and then use these to compute quantum averages of various physical quantities under investigation.

To the right is a visualization of one of these lattices, showing instantons, fluctuations of the gluon field which are topologically non-trivial and strongly effect the dynamics of the quarks. Red and blue blobs represent instantons and anti-instantons, which twist with opposite topology. In fact, it is believed that much of the mass of the neutron and proton is related to these excitations. The full four-dimensional visualization of these fluctuations can be viewed at <http://cliodhna.cop.uop.edu/~hetrick/timeseq1135.mpg> as an MPEG movie.

The most time-consuming part of lattice gauge theory is the production of these lattices, which must be done so that the fluctuations are truly representative of the underlying laws of quantum field theory, including, for example, the probability that quarks and anti-quarks pop out of the vacuum and then

annihilate each other, producing more gluons. Running on a computer capable of performing several hundred gigaflop/s, this process typically takes from hours to days per lattice. But to complete a study of a particular quantity, several hundred to thousands of lattices must be generated. The bottom line? Creating enough lattices for meaningful research can require teraflop-years of computer time.

Once generated, however, these lattices can be used by many researchers who might each be interested in studying different phenomena. Measuring quantum averages on the ensemble can be done quickly on local clusters or even workstations. The obstacles to sharing lattices in the past included the size and location of the data. Typically, lattices are generated at a supercomputing center, and access is restricted to those with accounts at the centers. Thus it is impossible for a researcher in France, say, to copy the files to their local account.

In 1998, Jim Hetrick at the University of the Pacific in California, Greg Kilcup at The Ohio State University, and NERSC's mass storage group developed the first Web-based lattice repository, thereby giving other scientists access to lattice data stored on HPSS at NERSC. The lat-


NERSC Helps Make Strides Toward Open Science Grid

In addition to making all of its production systems available via the DOE Science Grid, NERSC is also participating in the Open Science Grid initiative. Late last year, NERSC's 400-processor PDSF cluster was integrated into a broad scientific grid currently connecting 27 sites in the U.S. and South Korea.

Driven by the high energy physics community to help prepare for participation in multi-institution experiments on the Large Hadron Collider at CERN, the Open Science Grid (OSG) aims to tie together the scientific grid infrastructure of DOE, the National Science Foundation and other agencies in the United States.

As a lead in to the OSG, the Grid2003 project was launched as a collective effort by the U.S. Physics Grid Projects (DOE-funded PPDG, NSF-funded iVDGL, NSF-funded GriPhyN) and the software and computing projects of the US-ATLAS, US-CMS, the LIGO and SDSS experiments. Getting all 27 sites, including NERSC, integrated into Grid2003 entailed many teleconferences, countless emails, extensive diagnosing of cryptic error messages and very long hours.

At NERSC, Shane Canon installed and configured grid-monitoring software, then configured the appropriate security settings for the grid middleware so that PDSF could participate effectively and securely in this shared grid fabric. Iwona Sakrejda installed and configured the grid middleware and applications packaged for Grid2003 on PDSF.



A visualization of a QCD lattice, with red and blue blobs representing instantons and anti-instantons.

tice archive (<http://qcd.nersc.gov>) allows researchers around the world to browse and download lattices of interest. Hetrick and Kilcup maintain the archive and contribute the lattices of their research groups and contributions from others.

This archive was the first of its kind in the lattice gauge theory community and quickly became very popular. It has served terabytes of data to researchers around the world who would either not have been able to carry out their investigations, or would have first had to spend many months generating their own lattice data.

For more information, contact Nancy Meyer at NLMeyer@lbl.gov.