NERSC NEWS

News from the National Energy Research Scientific Computing Center

October 2005

NERSC Develops Archiving Strategies for Genome Researchers

When researchers at the Production Genome Facility at DOE's Joint Genome Institute found they were generating data faster than they could find somewhere to store the files, let alone make them easily accessible for analysis, a collaboration with NERSC's Mass Storage Group developed strategies for improving the reliability of storing the data while also making retrieval easier.

DOE's Joint Genome Institute (JGI) is one of the world's leading facilities in the scientific quest to unravel the genetic data that make up living things. With advances in automatically sequencing genomic information, scientists at the JGI's Production Genome Facility (PGF) found themselves overrun with sequence data, as their production capacity had grown so rapidly that data had overflowed the existing storage capacity. Since the resulting data are used by researchers around the world, ensuring the data are both reliably archived and easily retrievable are key issues.

As one of the world's largest public DNA sequencing facilities, the PGF produces 2 million files per month of trace data (25 to 100 KB each), 100 assembled projects per month (50 MB to 250 MB), and several very large assembled projects per year (~50 GB). In aggregate, this averages about 2000 GB per month.

In addition to the amount of data, a major challenge is that way the data are produced. Data from the sequencing of many different organisms are produced in parallel each day, such that a daily

(continued on page 2)

NERSC News

NERSC News, highlighting achievements by staff and users of DOE's National Energy Research Scientific Computing Center, is published every other month via email and may be freely distributed. NERSC News is edited by Jon Bashor, JBashor@lbl.gov or 510-486-5849.

KamLAND Uses PDSF, HPSS to Find First Geo-Neutrinos

In a paper which was featured on the cover of the July 28, 2005 issue of Nature, an international group of researchers reported the first observation of geologically produced anti-neutrinos. The observation is giving scientists new insight into the interior of our planet.

While the "geo-neutrinos" were detected at the KamLAND facility in Japan, most of the data was stored on NERSC's High Performance Storage System (HPSS) and analyzed using the PDSF cluster at NERSC. Together, these systems allowed scientists to find the scientific equivalent of a needle in a very large haystack.

KamLAND records data 24 hours a day, seven days a week. This data is shipped on tapes from the experimental site to LBNL, where it is read off the tapes and stored in the HPSS at NERSC. KamLAND records about 200 GB of data each day and HPSS currently has more than 250 TB of KamLAND data stored, making KamLAND the second-largest user of NERSC's HPSS system.

The KamLAND experiment, located in a mine in Japan, is a 1 kiloton liquid scintillator detector that was built to study anti-neutrinos coming from Japanese nuclear reactors, which are about 200 km from the detector. KamLAND is the first reactor experiment that observed the disappearance of electron anti-neutrinos from the reactor to the detector. Last year, the experiment also showed that the energy spectrum has a distortion typical of neutrino oscillation and measured the so-called mass-splitting, a key parameter in neutrino oscillation.

During dedicated production periods at NERSC, the KamLAND data are read out of HPSS and run through the reconstruction software to convert the waveforms (essentially oscilloscope traces) of about 2,000 photo-multiplier tubes (PMTs) to physically meaningful quantities such as energy and position of the event inside the detector. This reduces the data volume by a factor of 60-100 and the reconstructed events are stored on disk for further analysis.

"The event reconstruction requires a lot of computing power, and with over 600 CPUs, PDSF is a great facility to run these kinds of analysis," said Patrick Decowski, an LBNL physicist who works with NERSC staff on the project. "PDSF has been essential for our measurements."

With the data on disk, specialized analysis programs run over the reconstructed events to extract the geo-neutrinos and perform the final analysis. PDSF is also used for various simulation tasks in order to better understand the background signals in the detector.

"The whole analysis is like looking for a needle in a haystack — out of more than 2 billion events, only 152 candidates were found," Decowski said. "And of these, 128 — plus or minus 13 — are background events."

Forty years ago, the late John Bahcall proposed the study of neutrinos coming from the sun to understand the fusion processes inside the sun. The measurement of a persistent deficit of the observed neutrino flux relative to Bahcall's calculations led to the 2002 Nobel Prize for Ray Davis and the discovery of neutrino oscillation.

Today, anti-neutrinos are being used to study the interior of the Earth, which is still little known. The deepest borehole ever drilled is less than 20 km in depth, while the radius of the Earth is more than 6000 km. While seismic events have been used to deduce the interior makeup of the Earth's three basic regions — the core, the

(continued on page 2)

KamLAND Researchers Rely on PDSF, HPSS (continued from page 1)

mantle and the crust — there are no direct measurements of the chemical makeup of the deeper regions.

An important measurement to understand the Earth is the measurement of the heat-flux coming from within. These measurements show that the Earth produces somewhere between 30 and 45 TW of heat. Two important sources of heat generation are the primordial energy released from planetary accretion and latent heat from core solidification. However, it is believed that radiologically produced heat (heat from radioactivity) also plays an important role in the Earth's heat balance, contributing perhaps half of the total heat.

Neutrinos can help in the understanding of the Earth's internal structure and heat generation. Three important isotopes that are part of current Earth models — potassium, uranium and thorium — produce electron anti-neutrinos in their radioactive decay. These neutrinos (so-called geo-neutrinos) only interact with the surrounding Earth material very weakly and almost all of them reach the surface of the Earth.

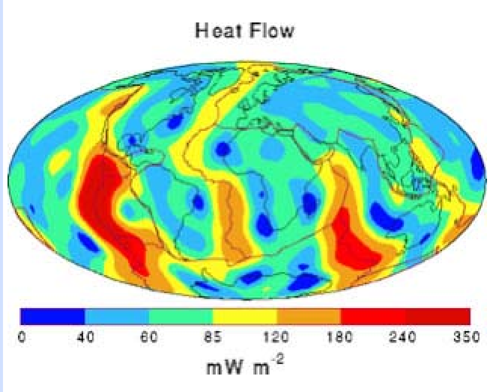

Earth's conductive heat flow is estimated to be about 31 terawatts, nearly half of which comes from the Earth's interior. Radioactivity is known to account for some of this heat, but there has been no accurate means of measuring radiogenic heat production.

However, occasionally they do interact with normal matter, and by building a large device that can detect them, something can be learned about the abundance of these isotopes. This allows scientists to study part of the composition of the Earth and most importantly, provide an estimate of the amount of heat produced through radioactive decay. The research is a multinational effort, as shown by the fact that the Nature article represented the work of 87 authors from 14 institutions spread across four nations.

Archiving Genome Data (continued from page 1)

"archive" spreads the data for a particular organism over many tapes.

DNA sequences are considered the fundamental building blocks for the rapidly expanding field of genomics. Constructing a genomic sequence is an iterative process. The trace fragments are assembled, and then the sequence is refined by comparing it with other sequences to confirm the assembly. Once the sequence is assembled, information about its function is gleaned by comparing and contrasting the sequence with other sequences from both the same organism and other organisms. Current sequencing methods generate a large volume of trace files that have to be managed — typically 100,000 files or more. And to check for errors in the sequence or make detailed comparisons with other sequences, researchers often need to refer back to these traces. Unfortunately, these traces are usually provided as a group of files with no information as to where the traces occur in the sequence, making the researcher's job more difficult.

This problem was compounded by the PGF's lack of sufficient online storage, which made organization (and subsequent retrieval) of the data difficult and led to unnecessary replication of files. This situation required significant staff time to move files and reorganize file systems to find sufficient space for ongoing production needs; and it required auxiliary tape storage that was not particularly reliable.

Enter NERSC's Archiving Expertise

Staff from NERSC's Mass Storage Group and the PGF agreed to work together to address two key issues facing the genome researchers. The most immediate goal was for NERSC's High Performance Storage System to become the archive for the JGI data, replacing the less-reliable local tape operation and freeing up disk space at the PGF for more immediate production needs. The second goal was to collaborate with JGI to improve the data handling capabilities of the genome sequencing and data distribution processes.

NERSC storage systems are robust and available 24 hours a day, seven days a week, as well as highly scalable and configurable. NERSC has high-quality, high-bandwidth connectivity to the other DOE laboratories and major universities provided by ESnet.

Most of the low-level data produced by the PGF are now routinely archived at NERSC, with ~50 GBs worth of raw trace data being transferred from JGI to NERSC each night.

The techniques used in the developing the archiving system allow it to be scaled up over time as the amount of data continues to increase — up to billions of files can be handled with these techniques. The data have been aggregated into larger collections which hold tens of thousands of files in a single file in the NERSC storage system. This data can now be accessed as one large file, or each individual file can be accessed without retrieving the whole aggregate.

And not only will the new techniques be able to handle future data, they also helped when the PGF staff discovered that raw data that had been previously processed by software that had an undetected "bug." The staff were able to retrieve the raw data from NERSC and reprocess it in about 1? months, rather than go back to the sequencing machines and produce the data all over again — which would have taken about six months. In addition to saving time, this also saves money — a rough estimate is that the original data collection comprised up to 100,000 files/day at a cost of $1 per file, which added up to $1.2 million for processing six months' worth of data. Comparing this figure to the cost of a month and a half of staff time, the estimated savings are about $1 million—and the end result is a more reliable archive.