KamLAND Data Processing

M.P. Decowski for the KamLAND Collaboration

The KamLAND collaboration collects on average about 160 GB of experimental data a day. The data is written to LTO[1] tapes in Japan and shipped to the US for data reconstruction and analysis. When the data arrives in the US, it is copied into a HPSS[2] storage system at the NERSC computing facility.

During the initial phase of the analysis, KamLAND had access to a large, 350-slot IBM LTO tape-robot (located at the Oakland NERSC facility) to copy the data from LTO tape to the HPSS system. Up to 60 tapes a week could be copied into HPSS due to the close network proximity of the LTO library to the HPSS system. However, this system became unavailable to the collaboration at the end of 2002. We therefore purchased two smaller, 23-slot LTO libraries and build a tape-copying system at LBNL. The current system is fully operational and transfers up to 35 tapes per week to HPSS.

KamLAND has copied over 75 TB of data into HPSS over the past year. The data was initially uncompressed, but a sophisticated new compression scheme was developed at LBNL in the fall of 2002 to compress the data to 30% of its original size. The compression system was deployed at the experimental site in December 2002. However, at the same time additional 20-inch PMTs were commissioned, which increased the raw amount of data. At the present time, with compression and the 20-inch data, KamLAND is writing about 140 GB of data a day. We are currently working on methods to reduce the data volume further, to under 100 GB a day.

It typically takes about 6 to 8 weeks from the date the data is recorded until it becomes available in HPSS.

KamLAND started using the NERSC PDSF facility in the summer of 2002. PDSF operates by pooling computing resources from different experiments, where each experiment contributes resources based on its usage. The initial KamLAND analysis period made use of temporarily unused computational capacity at PDSF.

Production runs at the PDSF cluster at NERSC process the data and reconstruct physical quantities such as the event vertex and energy. This information is stored in ROOT based data summary (DST) files. An additional pass correlates the events in the DSTs and produces mini-DSTs (so-called 'coincidence files') which are also stored in a ROOT compatible format. The overall data reduction from raw data (stored in HPSS) to DST and mini-DST level is around a factor of 100. The large data reduction allows collaborating institutions to copy the mini-DSTs to a local system for further (interactive) analysis.

The following tasks are performed during a typical data production run:

1. *Time sorting (Event-building).* The waveform signals from the individual photo-multiplier tubes are acquired asynchronously and have to be sorted in time order before the data processing can begin.

2. *Signal Processing.* This step involves waveform signal conditioning and pedestal subtraction.

3. *Time and charge processing.* This processing step extracts the pulses from the waveforms and records the time and charge of each pulse. Several pulses can be extracted from one waveform.

4. *Event reconstruction.* Once the time and charge information is available, the event vertex and event energy are reconstructed for low energy events. Muon tracks are reconstructed for muon events (about 1 event in a 100 low energy events is due to a muon). This step produces DST files.

5. *Event correlation.* This step involves correlating events in the DST files with each other in order to find the neutrino event signature and to veto backgrounds. The data is written to mini-DST files.

6. *Physics analysis.* Here, the output from event reconstruction and event correlation is analyzed. Cuts to the data are studied and applied. Physical processes such as the backgrounds and finally the anti-neutrino signal are extracted.

Steps 1-5 in the above list are typically run in a batch job, while step 6 is performed interactively.

At the present time, the time-sorting and event-building (step 1) is done at the experimental site and shipped together with the asynchronous data on tapes to the US.

The collected data is typically reconstructed in a production run once every two months.

---

[1] Linear Tape Open
[2] High Performance Storage System