

Onsite data processing and monitoring for the Daya Bay experiment ^{*}

LIU Ying-Biao(刘颖彪)^{1;1)} HE Miao(何苗)^{1;2)} LIU Bei-Jiang(刘北江)^{1,2,3} WANG Meng(王萌)⁴
 MA Qiu-Mei(马秋梅)¹ QI Fa-Zhi(齐法制)¹ ZENG Shan(曾珊)¹

¹ Institute of High Energy Physics, Chinese Academy of Sciences, Beijing 100049, China

² Chinese University of Hong Kong, Hong Kong, China

³ Department of Physics, The University of Hong Kong, Pokfulam, Hong Kong, China

⁴ Shandong University, Jinan 250100, China

Abstract: The Daya Bay Reactor Neutrino Experiment started running on September 23, 2011. The offline computing environment, consisting of 11 servers at Daya Bay, was built to process onsite data. With the current computing ability, onsite data processing is running smoothly. The Performance Quality Monitoring system (PQM) has been developed to monitor the detector performance and data quality. Its main feature is the ability to efficiently process multiple data streams from the three experimental halls. The PQM processes raw data files from the Daya Bay data acquisition system, generates and publishes histograms via a graphical web interface by executing the user-defined algorithm modules, and saves the histograms for permanent storage. The fact that the whole process takes only around 40 minutes makes it valuable for the shift crew to monitor the running status of all the sub-detectors and the data quality.

Key words: multi-data-stream, data quality, Daya Bay

PACS: 29.40.Mc, 28.50.Hw, 13.15.+g **DOI:** 10.1088/1674-1137/38/8/086001

1 Introduction

The Daya Bay experiment is designed to determine precisely the neutrino mixing angle θ_{13} with a sensitivity better than 0.01 in the parameter $\sin^2 2\theta_{13}$ at 90% confidence level. Three experimental halls (EHs) were built to contain eight functionally identical antineutrino detectors (ADs) [1]. Before October 18, 2012, two ADs were located in EH1 and one in EH2 (the near halls). Three ADs were positioned in the far hall, EH3. After the installation of the last two ADs, 2 and 4 ADs are located in each near hall and the far hall, respectively. In each experimental hall there are also muon detectors, including the inner (IWS), outer (OWS) water shields, and a resistive plate chamber (RPC) tracker. With the 6-AD configuration, the Daya Bay collaboration reported the first observation of non-zero $\sin^2 2\theta_{13}$ [2], and has recently updated the result of $\sin^2 2\theta_{13} = 0.089 \pm 0.010(\text{stat.}) \pm 0.005(\text{syst.})$ [3] in a three-neutrino framework.

During physics data taking, events from different experimental halls are separately assembled in raw data format and stored into persistent files on the online disk

array by the data acquisition system (DAQ). Event rates for the 8-AD configuration for EH1, EH2, EH3 are about 1.2 kHz, 1.0 kHz, and 600 Hz, respectively. The total number of raw data files each day is about 320, approximately 1 GB volume per file. The online histogram presenter of the DAQ provides the hit-maps for sub-detectors. In addition, it is necessary to produce high-level histograms to monitor the running status of all sub-detectors and the data quality (DQ).

The Performance Quality Monitoring system (PQM) has been developed to use the offline software to process multiple data streams from the three experimental halls and generate more user-defined histograms, including some high-level ones containing the calibration and reconstruction information. It thus plays an important role in offline monitoring for the experiment. Even with the total event rate of around 3 kHz, the PQM can handle the whole process in around 40 minutes.

In this article, we present the offline computing environment at Daya Bay in Section 2. In Section 3, the PQM design is outlined. The performance and the information display of the PQM are shown in Section 4. We give a short summary in the last section.

Received 11 September 2013, Revised 27 September 2013

^{*} Supported by NSFC (11105168) and 973 Program (2013CB834303)

1) E-mail: liuyb@ihep.ac.cn

2) E-mail: hem@ihep.ac.cn

©2014 Chinese Physical Society and the Institute of High Energy Physics of the Chinese Academy of Sciences and the Institute of Modern Physics of the Chinese Academy of Sciences and IOP Publishing Ltd

2 Daya Bay offline computing environment

2.1 Offline computers

There are 11 servers in the offline computing environment at Daya Bay, including a file server, a data transfer server, an offline database server, a web server, two PQM servers (pqm1-2), and five servers forming user farms (farm1-5), as shown in Fig. 1. pqm1 is usually run as the PQM server, in which the main PQM control script is running, and pqm2 is a backup machine. The PQM servers together with the user farms, which are all blade servers with 8 CPU cores per server, form a computing cluster. The CPUs of pqm1, farm1, farm2 are Intel (R) Xeon (R) E5506@2.13 GHz, and those of pqm2 and the other user farm machines are Intel (R) Xeon (R) E5420@2.50 GHz. Portable Batch System (PBS) is implemented to allocate computational jobs among the available computing resources. The 16 cores of the PQM servers are dedicated to the PQM jobs, and are sufficient during normal physics data taking. The 40 cores of the user farms can be shared with PQM jobs in the case of special data taking with higher event rates. The functions of the other servers will be described in Section 2.2.

The total disk volume of the file server is around 25 Terabytes (TB). 17 TB are for storing raw data files from the DAQ, 5 TB for onsite users, 2 TB for offline softwares and 1 TB for storing ROOT [4] files and histogram pictures produced by PQM jobs.

2.2 Data flow

The offline computing environment communicates with the online systems to process onsite data. Here we briefly introduce the related online system parts, as shown in Fig. 1. These are the DAQ system, the Daya Bay Detector Control System (DCS), and the online database. When the DAQ closes a raw data file, the information, including the starting and end timestamps of the file, the run number, the file number and so on, are dumped into the online database to generate a new record. The values of the high-voltage supplies to the photo-multiplier tubes (PMTs) and RPC modules, the detector temperature, the water levels of the water shields, etc., which are provided by the DCS, are also recorded in the online database.

To prevent interruptions to the DAQ, a raw data file is copied to the offline file server after the data transfer server detects a new record in the online database. Copies of raw data files are stored in the file server for more than 1 month. The offline database server is responsible for copying the information that is necessary for offline analysis from the online database to the offline database. Besides the online information, there are cal-

ibration constants in the offline database, such as PMT gains, PMT time offsets, and energy scales of the ADs, which are updated after careful validation. The PQM server is responsible for the implementation of the functionality of the PQM for monitoring the detector performance and the DQ. This will be described in detail in Section 3. The web server is used to display user-defined histograms in analysis algorithm modules.

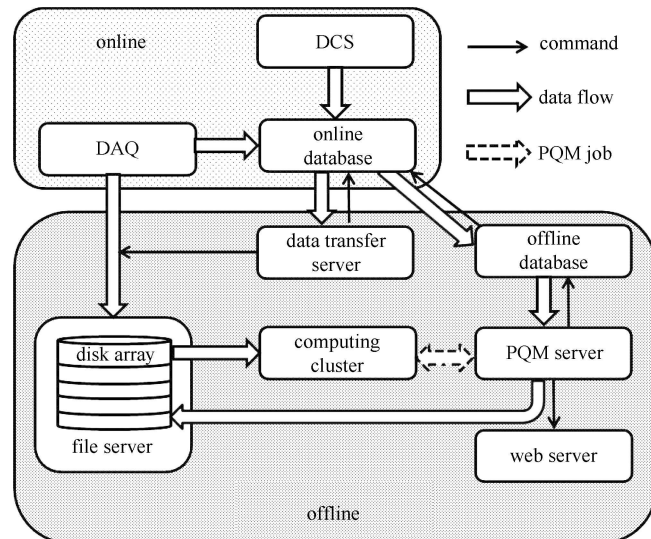


Fig. 1. Schematic diagram of the offline computing environment at Daya Bay.

2.3 Offline software

The Daya Bay offline software (NuWa) [1], using Gaudi [5] as the underlying software framework, is available in the offline computing environment. NuWa, which employs Gaudi's event data service as the data manager, provides the full functionality required by simulation, reconstruction and physics analysis. Job modules are used to configure simulation and analysis tasks. Specifically, job modules are scripts which can add analysis algorithms and tools to the job, and configure algorithms, tools, and services used by the job. In the NuWa framework, reconstruction algorithms have been developed to reconstruct the energy and the vertex of the antineutrino event from the charge pattern of the PMTs. One can find the details of the calibration and reconstruction methods in Ref. [1]. By default, the latest calibration constants in the offline database are loaded into the reconstruction algorithms, reflecting the best understanding of the detectors.

The NuWa auto-build system is implemented using the Bitten [6] plug-in for trac [7], which is an enhanced wiki and issue tracking system for software development projects. One of the user farms serves as the Bitten slave and builds NuWa automatically when the code is updated. A series of pre-defined jobs for validation testing

follow the build. The validated build can then be manually set for PQM jobs if it contains a necessary upgrade of reconstruction or analysis algorithms.

3 Implementation of the PQM

The PQM is designed to efficiently process all events in all raw data files as soon as possible. It communicates with the offline database to get the latest raw data file list and submits jobs to PBS to process them in parallel. At the end of each job, some selected histograms are printed for web display. The shift crew can compare the new histograms with standard ones to report sub-detector status and DQ.

3.1 Data processing chain

To manage the data flow of the PQM as shown in Fig. 2, a Python control script has been developed and runs in background mode on the PQM server. When the DAQ closes a raw data file, the offline database is updated with a new record copied from the online database. A copy of the file is transferred to the file server in less than 5 minutes, assuming a size of ~ 1 GB. When the data transfer is complete, the status of this file is set to TRANSFERRED in both online and offline databases.

The control script queries the offline database at a

fixed time interval (10 seconds). If a new record is found and the data status returned by the query shows that the corresponding raw data file is already on the file server, a job for processing the new file is submitted to PBS. The PQM job uses a recently manually tagged NuWa version to reconstruct data with the latest calibration constants in the offline database, runs analysis algorithms, and fills the results into user-defined histograms. All of the histograms are dumped into a ROOT file at the finalization step of the algorithms, which is then merged with the accumulated ROOT file for the same run. After that, a C++ histogram printer is executed, taking the merged ROOT file as input to print selected histograms into figures. The control script also checks if the job is finished at 10-second intervals by detecting an empty TXT file created at the end of the job. For finished jobs, the control script transfers the corresponding figures to the PQM disk for web display (see Section 4). When the run ends and all the corresponding raw data files are processed, the control script saves the accumulated ROOT file on the PQM disk for permanent storage.

3.2 PQM analysis modules

Currently, four modules have been developed for the PQM to analyze data and fill user-defined histograms in the NuWa framework: one for histograms of electronics

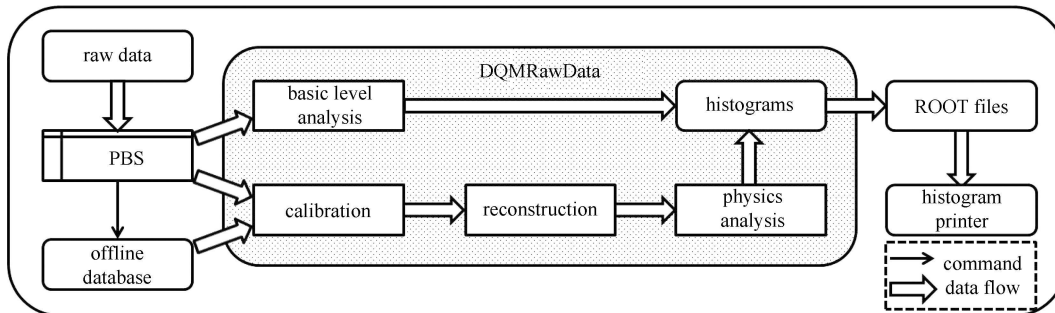


Fig. 2. The data flow of the PQM job.

Table 1. The histograms produced by analysis modules in PQM for monitoring the detector performance and the DQ.

detector unit	level	histograms
PMT	basic	mean of ADC, TDC and preADC, RMS of ADC, TDC and preADC, Δ ADC, dark noise, dark rate and TDC vs channel ID, hit rates, ADC sum vs. trigger types, ADC sum, number of blocked triggers vs. run time, number of channels, trigger types
	reconstruction	2-dimension vertex distribution, distribution of vertex vs. energy, distribution of event rate vs. vertex, distribution of energy
RPC	basic	map of patch rate, map of module rate, map of channel multiplicity, distributions of time between module and system triggers, maps of layer efficiency, distribution of layer efficiency, electronics error types vs. time, maps of channel hit count, channel hit count vs. channel ID, distribution of number of layers per readout, distribution of number of modules per system trigger, maps of layer singles rate, distribution of layer singles rate, module and system trigger rates vs. time.

channel information, one for histograms of calibrated PMT information for the ADs and water shields, one for the histograms of the reconstruction for the ADs, and one for the RPC histograms. The histograms displayed by the web interface are shown in Table 1. A detailed description of the variables listed in Table 1 can be found in Refs. [1, 8, 9]. This provides the flexibility to add more algorithm modules as well as to fill more user-defined histograms.

Since a different configuration of detectors is adopted for EH3, all the analysis modules take the strategy of creating histograms dynamically for each detector, and even for each PMT channel, to reduce the output file volume and improve the processing efficiency. If one event read from the raw data file is the first event for the current detector, a new set of histograms for this detector is created. Otherwise, the event is analyzed and filled into existed histograms.

4 Performance and information display

The PQM is required to provide more information than the online histogram presenter provided by the DAQ about the running status of all sub-detectors in the 3 EHs. The data transfer from the DAQ takes less than 5 minutes. The job can be submitted to PBS about 10 seconds after the data transfer is finished. Running NuWa job modules takes around 30 minutes. In the meanwhile, some user-defined histograms produced by job modules are saved in ROOT files. They are printed and published via the web interface in around 2 minutes. The job failure rate is less than 1% during normal data taking. The typical latency of the PQM plot display is 40 minutes after a new raw data file has been closed by the DAQ. The most time-consuming part is running the calibration and reconstruction algorithms. Since all algorithms are configurable in the PQM job, it is easy to remove them and show only the basic-level information of the detectors to reduce the latency to less than 20 minutes. This special configuration of PQM jobs was usually implemented during the commissioning period, since the onsite experts could see the detector status as soon as possible.

A web interface, as shown in Fig. 3, runs from the web server to display the selected user-defined histograms. On the main page (the left panel of Fig. 3), two options are provided for users: one to show histograms for the current runs at each site, named ‘realtime’, the other to show histograms for all previously processed runs, named ‘history’. For the latter, one can choose the run number in a drop-down menu listing the site ID and run type, i.e. ‘Physics’, ‘FEEDiag’ and so on. On the ‘realtime’ page (the right panel of Fig. 3), site buttons are available for different sites: ‘DayaBay’, ‘LingAo’ and ‘Far’.

The ‘realtime’ page refreshes at a fixed time interval of 1 minute. For each site, the user can also use detector buttons to view histograms for different sub-detectors. On the ‘history’ page, there are only detector buttons, because this page is for the run selected by the user, and the specific run is for only one site.

Several example histograms are given here. Fig. 4 shows the number of blocked triggers in one event as a function of time. Generally, if a cosmic-ray muon passes through the detector and induces a shower with extremely large energy deposit, overflow of the electronics occurs, and the local trigger and the corresponding data package can be blocked. The blocked trigger number is recorded and read out. If the fraction of the blocked trigger number over the total trigger number is larger than a pre-defined value, the shift crew should report to the DQ working group (DQWG).



Fig. 3. Web interface for the PQM.

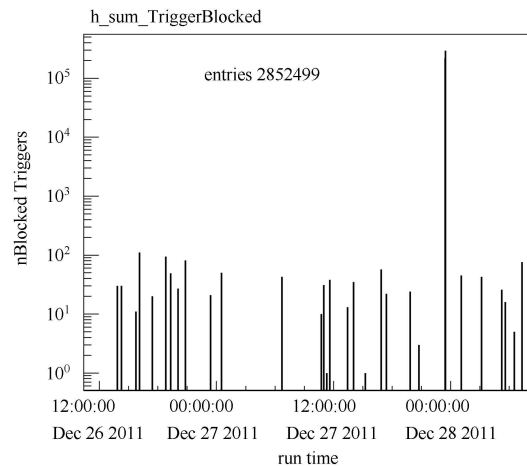


Fig. 4. Number of blocked triggers in one event vs. run time.

Figure 5 presents the reconstructed energy spectrum for all triggers for an AD. The peaks of ^{40}K (1.46 MeV), ^{208}Tl (2.61 MeV), and the gamma peak from neutron capture on Gd (~ 8 MeV) are labeled in fixed positions in the histogram. The red line for ^{208}Tl is set to ~ 2.75 MeV due to the energy non-linearity which is not corrected in the analysis algorithm. A relative shift between the red lines and the observed peaks indicates problems in the data, e.g. incorrect energy scales read from the offline central database or unexpected behaviour of the PMTs.




Fig. 5. The reconstructed energy distribution of all events for AD1.

The statistics for the efficiency of each RPC layer [8, 9] are displayed in Fig. 6. The efficiency is calculated in the algorithm of analysis modules for the PQM. The efficiency is a critical quantity to evaluate the performance of the RPCs used as an anti-coincidence detector for muons for the experiments. As shown in the figure, most of the RPC layers are working at an efficiency of about 90%. The performance of the RPCs can change due to the gas pressure or other unexpected reasons. The RPC experts and the DQWG are informed if the RPC layer efficiency distribution distorts greatly in comparison with the standard distributions.

There are also many other histograms to monitor the detector performance and DQ. The shift crew can compare them with the corresponding standard figures to check the running status of the sub-detectors and the DQ. In the case that irregularities are observed, the shifter should inform the related experts and report to the DQWG.

A supernova monitoring program is under development to join the Supernova Early Warning System (SNEWS) [10], in which individual neutrino-sensitive experiments send supernova (SN) burst datagrams to the SNEWS coincidence computer at Brookhaven National Laboratory (BNL). Supernova burst neutrinos are im-

portant in studying the dynamics of core collapse and supernova explosions, neutrino properties, and many other interesting studies. The three-hall configuration design at Daya Bay can significantly reduce false SN alarms caused by the spallation background. We will have on-line and offline triggers for monitoring SN bursts. Given the short latency time of the PQM, the offline trigger for SN bursts will be added to the PQM job.




Fig. 6. The efficiency of each layer for the RPCs.

5 Conclusions

The Daya Bay onsite data processing is running smoothly. The PQM, which has been developed for the Daya Bay experiment, has the ability to provide user-defined histograms to the shift crew with a latency of about 40 minutes. This satisfies the requirements for data taking. By comparing the histograms with standard distributions, potential DQ problems can be found. The PQM will contribute to the Daya Bay SN monitoring program in the near future.

The author, LIU Ying-Biao, would like to thank LI Gao-Song for the contribution to the ReconDataHistogram, Logan Lebanowski and ZHANG Qing-Min for the contribution to the RpcDataHistogram, and JI Xiang-Pan for the contribution to CalibDataHistogram.

References

- 1 AN F P et al. (Daya Bay collaboration). Nucl. Instrum. Methods A, 2012, **685**: 78
- 2 AN F P et al. (Daya Bay collaboration). Phys. Rev. Lett., 2012, **108**: 171803
- 3 AN F P et al. (Daya Bay collaboration). Chinese Physics C (HEP & NP), 2013, **37**: 011001
- 4 ROOT, An Object Oriented Data Analysis Framework. <http://root.cern.ch/drupal/>
- 5 Gaudi. <http://cern.ch/gaudi/>
- 6 <http://bitten.edgewall.org/>
- 7 <http://trac.edgewall.org/>
- 8 ZHANG Q M et al. J. Phys.: Conf. Ser., 2012, **396**: 022061
- 9 NING Z et al. J. Instrum., 2013, **8**: T03007
- 10 WEI H Y. Supernova Trigger in the Daya Bay Reactor Neutrino Experiment. <http://indico.ihep.ac.cn/getFile.py/access?contribId=54&sessionId=5&resId=0&materialId=slides&confId=2813>. In: The 15th Int'l Workshop on Advanced Computing and Analysis Techniques in Phys. Research, 2013