ISBN: 978-3-95450-193-9 doi:10.18429/JACoW-ICALEPCS2017-THPHA195

THE CERN n_TOF FACILITY DATA ACQUISITION SYSTEM

A. Masi, A. S. Almeida, R. Cevenini, E. Chiaveri, M. Donze, S. Gilardoni, A. Giraud, R. Losito, D. Macina, F. Marazita, P. Peronnard, Á. Prieto, L. Tassan Got, CERN, Geneva, Switzerland

Abstract

n_TOF is a pulsed neutron facility at CERN which studies neutron interactions as function of the energy. Neutrons are produced by a pulsed proton beam from the PS directed to a lead target. In a typical experiment, a sample is placed in the neutron beam and the reaction products are recorded. The typical output signals from the n_TOF detectors are characterized by a train of pulses, each one corresponding to a different neutron energy interacting with the sample. The Data Acquisition System (DAQ) has been upgraded in 2014 and is characterized by challenging requirements as more than hundreds of 12 or 14-bit channels at a sampling frequency of 1 GS/s and 1.8 GS/s acquired simultaneously every 1.2 s for up to 100 ms. The amount of data to be managed can reach a peak of several GB/s. This paper describes the hardware's solutions as well as the software's architecture developed to ensure the proper synchronization between all the DAQ machines, the data's integrity, retrieval and analysis. The software modules and tools developed for the monitoring and control of the n_TOF experimental areas and the DAQ operation are also detailed.

INTRODUCTION

The CERN neutron time-of-flight facility n_TOF [1] features a white neutron source produced by spallation through 20 GeV/c protons produced by the PS and impinging on a lead target. The facility, aiming primarily at the measurement of neutron-induced reaction cross sections, was operating at CERN between 2001 and 2004, and then underwent a major upgrade in 2008 with the so called Experimental ARea (EAR) 1. During the CERN long shutdown 2013-14, n_TOF constructed a new experimental area (EAR2) above the spallation source, 10 times closer than EAR1 to the lead target. This allowed increasing the neutron flux by about 40 times. A whole renovation of the DAQ was also done, focused both on the software and the hardware, by increasing the number of readout electronic channels to cover EAR2 and the new physics challenges.

THE N_TOF DAQ REQUIREMENTS

The typical n_TOF DAQ measurement application consists in the digital acquisition of the detectors output signals to perform time domain analysis. The signal shape is a pulse train where each pulse amplitude is proportional to the energy of the neutron interaction products, the rise/falling time to the detector type and the integral to the neutrons flux. The acquisition time window for each PS pulse corresponds to the whole neutron energy range, i.e. from a gamma flash detected when the proton beam impinges on the target until approximatively 100 ms for thermal neutrons (the slowest ones) which travel a 185 m path to the EAR1. This also corresponds to the longest time window since the travel path to the EAR2 is only of 20 m, which gives an acquisition time window of about 10 ms. The acquisition is triggered by the PS timing signal and the data samples should be transferred to the memory of the host controller before the next acquisition starts again 1.2 s later. Data are stored permanently on the CERN Advanced Storage (CASTOR) [2]. The architecture is conceived to be modular with different ADC multichannel cards, distributed on several chassis equipped with a host controller running Linux CentOS 7 (DAQ units).

The key parameters of the digitizers are:

- The **ADC front-end** to match all the possible signals from the different detectors (bandwidth and amplitude). A variable input gain and a bias adjustment is preferable to use the entire ADC dynamic range.
- **ADC resolution and sampling frequency**: the detector signal noise floor gives an indication of the ADC resolution needed whilst the signal shape, rise time and duration, specify the minimum sampling frequency needed.
- **On board memory size**: there should be enough local memory on the digitizer to store the entire acquisition time window per channel at the maximum sampling frequency.
- **Card Data interface**: the digitizer data interfacing has to transfer on time all the channels' samples related to the maximum time window to the host controller memory, before the next PS cycle (1.2 s).

THE N_TOF DAQ COMPONENTS

The global DAQ hardware architecture for each experimental area is depicted in Figure 1.

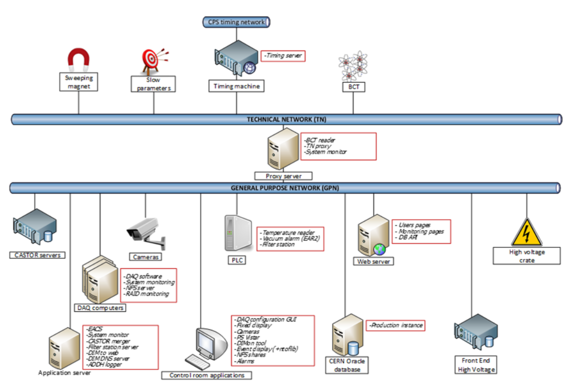

Figure 1: n_TOF DAQ hardware architecture.

The DAQ Units

Each DAQ unit hosts several high sampling Data Acquisition Cards (DAC) as well as the high writing speed local storage able to sustain the raw data bandwidth related to a maximum acquisition window of 100 ms. It guarantees a data buffer for 3 days of acquisition in nominal operation

ISBN: 978-3-95450-193-9 doi:10.18429/JACoW-ICALEPCS2017-THPHA195

conditions, to overcome network traffic issues in the data transfer to CASTOR.

The Hardware

The DAC are SPDevices ADQ412DC and ADQ14DC-4C. These cards feature the best trade off price performance with respect to the requirements set up by the experiments. Table 1 summarizes the main characteristics of the two models.

Table 1: n_TOF DAQ Cards Characteristics

Requirements	**ADQ412**	**ADQ14**
Resolution	12-bit	14-bit
Full Scale Input Range (FSR)	0.1-5Vpp	0.05-5Vpp
Sampling frequency & N. of channels	1.8GS/s 4-ch 3.6GS/s 2-ch	1GS/s4-ch
On board Memory Size	175MS/ch	256MS/ch
Bandwidth (-3dB)	1.3GHz	400MHz
Adjustable input bias	+-100% FSR	
Bus Interface	PCIe GEN2x8	
Trigger	Internal/External	
Scientific Linux drivers	Yes	

Nowadays the DAQ counts 12 ADQ412 and 21 ADQ14 for a total of 33 cards split in 6 DAQ units among the two experimental areas. Figure 2 shows the hardware architecture of a typical DAQ unit.

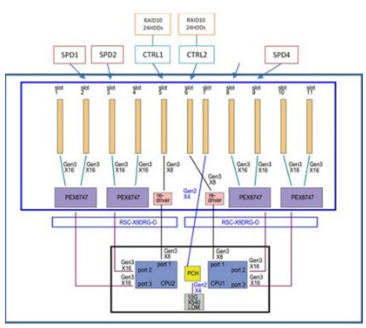

Figure 2: DAQ unit hardware architecture.

The chassis is based on a Supermicro chassis with 8 PCIe Gen3x16 slots with the following components:

- Two Intel Xeon E5-2650 CPUs with 8 cores and 4 memory channels each
- 384GB of DDR3-PC12800
- 48x 900GB SAS Hard Disk Drives divided in 2 stripes of 24 HDDs configured as RAID10
- Up to 6 DAC
- 2 Adaptec RAID controllers (CTRL1-2) configured in software RAID0.

In order to optimize the data transfer from the card to the memory and from the memory to the disks, each card is assigned to a different memory channel whilst two memory channels on the other CPU are used to dump the data to the local storage. Below is a summary of the recorded performances:

- 3.2GB/s data rate from card's buffer to chassis memory. The DAQ card has a Gen2x8 PCIe interface.
- Parsing of raw data (reorganization of the 12-bit data into 16-bit) is done at 800MB/s by the CPU
- Local storage data rate is 6GB/s peak

The Low-Level Acquisition Software

The low level software is organized in a single process multi-threaded application. In the main process, the state machine that takes care of the decoding and execution of the commands received through the experiment middleware is implemented. The data acquisition is organized in two separated threads: the Acquisition and the Writer. The former waits until each card is triggered and starts transferring the raw data from the memory card to a specific memory buffer system as soon as the acquisition is finished. The memory system allocation is performed at the beginning of the run. A pool of buffers is allocated for each card to avoid data losses even in the worst case of several consecutive triggers interleaved by only 1.2 s. The Writer thread instead performs the parsing and applies a zero-suppression algorithm on the data in memory for each card; after that, it takes care of writing the compressed data to the disk at low priority. Figure 3 depicts how the memory system is managed. Writer and Acquisition threads have their own indexes pool. Each time a trigger occurs, the acquisition threads pushes the corresponding index to the writer threads through a synchronized queue and then increments the index for the next trigger. Since the number of allocated buffers is limited, an overlap check is done. If all the buffers are used, an error is raised and the acquisition is stopped.

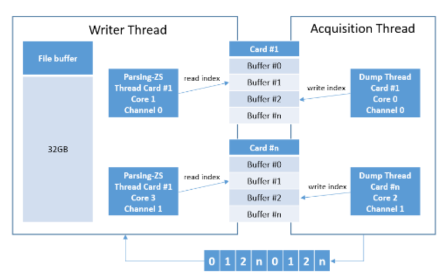

Figure 3: DAQ unit memory management.

The Timing Machine

The DAQ is synchronized with the operation of the PS accelerator, for example by the proton beam extraction towards the n_TOF target through the timing machine. A PCI industrial computer equipped with the CTR card receives information through the timing network [3], particularly the telegram of the n_TOF proton extraction, and translates it into a digital trigger with ns accuracy provided to each experimental area DAC trigger input. A fan out buffer is used to distribute the trigger signal to ensure the same timing delay for each card. The acquisition trigger was adjusted precisely with respect to the proton extraction from

ISBN: 978-3-95450-193-9 doi:10.18429/JACoW-ICALEPCS2017-THPHA195

the PS to start the acquisition right before the fastest neutrons start being detected.

The High Voltage Crate

Caen SY4527 High Voltage power supply mainframes are used to supply the detectors in each experimental area. A crate can contain several cards each one with up to sixteen 3 kV/3 mA High Voltage channels that can be controlled via TCP/IP using a proprietary protocol. Several parameters can be changed on each channel as current trip threshold, overvoltage, undervoltage, ramp up/down time.

The Infrastructure Monitoring & Control

The neutron flux production as well as the physics experiments rely on the control and monitoring of devices and/or critical experimental area parameters. In particular:

- Sweeping magnet, a resistive magnet used to eliminate charged particles from the neutron beam. Its 600 A power supply is controllable via the CERN accelerator middleware and a proper FESA class [4].
- Filter box, a motorized mechanical support to position in the neutron beam up to 8 different filters to absorb neutrons at specific energies. It is controlled with a Siemens PLC and the control software communicates with the high-level applications through TCP/IP using the S7 protocol
- Experimental area temperatures and ventilation door switches are acquired by a second Siemens PLC controlled via TCP/IP and the S7 protocol
  Several additional environmental data as experimental area pressures or target cooling parameters are instead acquired through the CERN infrastructure middleware DIP [5].

THE DAQ SOFTWARE ARCHITECTURE

The DAQ software architecture as well as the software components are shown in Figure 4.

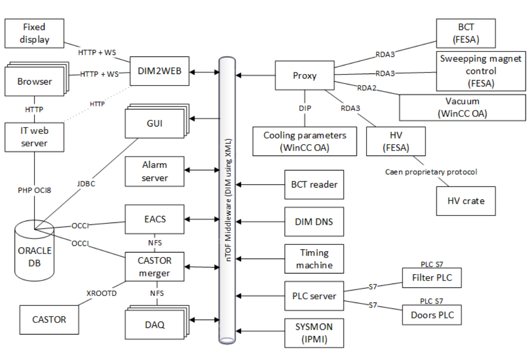

Figure 4: nTOF DAQ software architecture.

The Database

The n_TOF Database (DB) is from Oracle and stores critical information for the DAQ operation and the subsequent data analysis. The information is grouped as follows:

- The hardware settings
- The operational information
- The organisational data

The hardware settings group all the DAQ channels parameters and configuration, the High Voltage supply channels parameters as well as the file names and related size to be transferred from the DAQ unit local storage to CASTOR. The DAQ operation is organized in physics runs, each one related to a specific experimental areas configuration (i.e. target used, possible neutrons filters, detectors configuration, HV channels and DAQ channels settings). Each run usually lasts several hours. The description and the identifier of each physics run as well as the information on the experimental area layout, the number of protons cumulated on the target per run, the operator in the shift logbook content as well as the run start and stop time belong to the operational information and are essential for data analysis. The organisational data are instead all the information needed for the shifts management and follow-up as data of each n_TOF collaboration member (i.e. DAQ user) as well as of the shifts covered by each one during the year.

The Graphical User Interface (GUI)

The GUI has been developed in Java, runs on the consoles in the n_TOF Control Room (CR) and is the only application used by the shifter to operate the DAQ. The main interface is shown in Figure 5. It is characterized by the following main functionalities:

- Experimental Area configuration: according to the physics experiment, the operator can specify the material of the sample under study, the type of radioactive source used to calibrate the detectors or select a particular filter to insert in the neutron flux; he can controls the sweeping magnet current and set the DAQ operation mode in BEAM or CALIBRATION; with the former, the DAQ trigger signal is received by the timing machine whilst the latter mode allows the generation of a trigger signal at the frequency specified by the user and not synchronized to the PS operation. This operation mode is used to calibrate the detectors before a physics run
- DAQ configuration
- Physics run description: the shifter inserts the title and the description of the run as plus update the logbook
- Operation follow-up: after launching the DAQ acquisition, the GUI helps the shifter to monitor the correct operation; the main DAQ and environmental alarms responsible to prevent the acquisition are displayed together with a sound. The list of the validated events and the related value of the cumulated protons are also updated online. In this case the operator can immediately check for network or synchronization problems.

The GUI directly accesses the n_TOF Database using JDBC to store and retrieve all the data mentioned above. The communication with the DAQ units is performed instead through the EACS.

ISBN: 978-3-95450-193-9 doi:10.18429/JACoW-ICALEPCS2017-THPHA195

The EACS

The Experimental Area Coordination System (EACS) is the DAQ coordinator. It connects to the DB through OCCI whilst to the GUI and each DAQ units through the n_TOF middleware. At the start up, it initializes each card with the latest configuration settings stored in the database. It receives commands and settings from the GUI and dispatches them to the related DAQ units. The changes are saved in the DB only if the latter respond correctly.

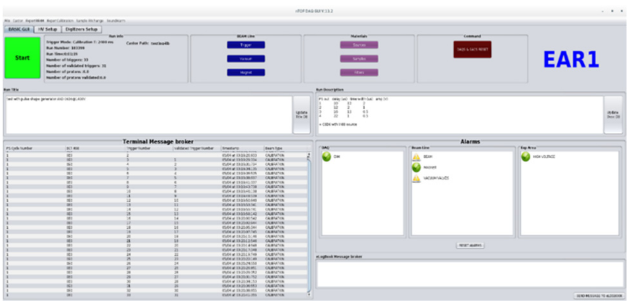

Figure 5: n_TOF DAQ GUI main panel.

The EACS exploits the crucial function to synchronize the acquisition data of all the machines. Indeed, at each trigger received, the data are considered valid only if all the active DAQ units have successfully completed the acquisition. The EACS is notified by the timing server through the middleware when a new trigger has been fired. It increments the event number and waits for the file validation of each DAQ units. The EACS is also responsible for the creation of the index files. They contain the information concerning the run, the cards' configuration, and the data of the event (i.e. beam intensity) as well as the slow data, for example the parameters important for the data analysis as experimental area temperatures and the HV channel values.

The CASTOR Merger

The CASTOR merger is the application responsible for the re-building of the event and for the optimized transfer of the acquisition data from the DAQ units and the EACS to CASTOR. The experience acquired in the first year of operation has shown a serious degradation of the CASTOR files transfer performance with files of small sizes (i.e. less than 10 MB). A file size of a few GB represents a good trade-off to keep high performances both on the migration to and recall from CASTOR [6]. This implies that the files to be transferred should all ideally have a similar size or better than a data stream transfer has to be implemented. Although on each DAQ units at each event corresponds a different raw data file (with size ranging from hundreds of MB to a few GB depending on the number of channels activated and the level of compression), the application, once opened an xroot connection towards CASTOR, ensures a continuous data stream collecting the data through NFS from each DAQ unit and performing the merging of several raw data files. The final file on CASTOR is terminated and the connection is closed when the optimal size is reached (i.e. 4 GB) or a given number of files have been merged. The CASTOR merger is responsible to group together also all the header files on the EACS in order to write on CASTOR only one file per run. As much as optimizing the data transfer, the CASTOR merger also ensures that all the raw data are synchronized and grouped per experimental area. The information on the validated events are retrieved by the DB.

The Proxy

The entire DAQ system is on the CERN Global Purpose Network (GPN) whilst all the accelerators operational devices are on the Technical Network (TN), for safety and performance issue, completely separated by the GPN. The proxy machine allows the access to the infrastructure parameters as experimental area pressures or target cooling published by the related low-level systems through the CERN middleware CMW [7] or DIP [5]. These data are read by the EACS through the n_TOF middleware to prepare the events' files.

The n_TOF Middleware

A dedicated and proper middleware has been developed to fulfil the requirements of the n_TOF DAQ system as lightness, short latency on the commands/status communication, robustness. The experience granted on critical missions as the LHC Collimators low level control [7] suggested the use of the Distributed Information Management (DIM) system as a communication layer between the different software components. DIM is based on the Client/Server paradigm; Servers provide services to clients. For the n_TOF DAQ middleware XML has been used to format the service data in the so-called DIM data set, where the data are organized in properties and fields and the type of each field is specified. Taking into account the small quantity of data exchanged, XML has been preferred with respect to other solutions of data compression and serialization (i.e. Google protobuf) to enhance the diagnostic capability. DIM client/server communication has been integrated with a synchronous handshaking mechanism. A special command service coupled with an acknowledge service has been created. As soon as a new command is sent and received by the server, it sends back to the client the acknowledged with the proper command identification. The n_TOF middleware is based on a combination of DIM clients and servers deployed on the different components.

The WebMonitoring

The DAQ operation status and performance monitoring is ensured through proper customized web pages. Figure 6 shows the main DAQ monitoring page. It contains information on the proton beam (i.e. type of beam, primary or parasitic, and proton intensity), the DAQ status for each experimental area as well as a summary of the most important run information, such as the number of missed events for each run and the information on the shifter and shift leader. The web monitoring is enriched by a page to monitor in detail the status and the performance of the raw data files transferred to CASTOR as well as a dedicated web page to monitor the acquisition's software activity as well as the health parameters of each DAQ units (i.e. RAM, CPU and disk space usage). As depicted in Figure 4 the core of the web monitoring is the application named DIM2WEB. It collects all the significant information of the

ISBN: 978-3-95450-193-9 doi:10.18429/JACoW-ICALEPCS2017-THPHA195

DAQ status published by the different components through the n_TOF middleware and redirects directly on a web socket the requested DIM data set. The web pages front end in JavaScript interprets the information left in the XML format and displays them on the widgets. The web monitoring is accessible from anywhere. The front-end also takes care of identifying the web connection. If this comes from inside the CERN GPN, a direct connection via web socket with the DIM2WEB machine is established otherwise an HTTP request to CERN's n_TOF webserver is performed every 100 ms. On the back end, a PHP script retrieves the data for the web monitoring pages' front end through a HTTP request to the DIM2WEB.

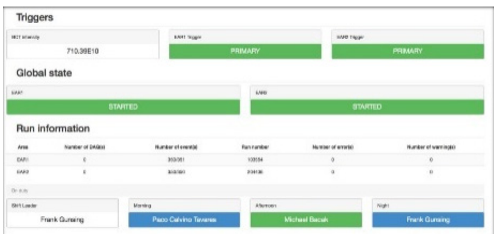

Figure 6: n_TOF DAQ main web monitoring page.

The Alarm Web Page

An alarm server application has been developed to verify that all the critical infrastructure parameters as well as the DAQ errors and warning are properly conveyed towards a dedicated alarms web display. It fetches the DAQ errors and warnings directly from the n_TOF middleware as well as the critical experimental area parameters published by the Proxy. Each parameter is compared with a defined warnings and errors threshold contained in a local configuration file. The alarms are published through dedicated DIM services subscribed by the DIM2WEB which takes care of web publications.

The File Viewer

Although the analysis of the data acquired is performed using proper routines developed per detector and experiment by the n_TOF collaboration and launched locally on the files recalled by CASTOR [8], a standalone application in C++ named "event display" has been developed to perform a pre-analysis of the data acquired. It enables the display of the detectors signals read both from CASTOR and/or each DAQ units together with the zero-suppression algorithm threshold set during the acquisition. This application is extremely useful during the experiment setup to verify the integrity and the noise level of the detectors signals before the raw data files are transferred to CASTOR.

The DIMon

The DIM Monitoring Tool (DIMon) allows the shifter to keep under control the status of the n_TOF middleware communication between each DAQ components (i.e. DIM services timeout) as well as the status of each software process. The operator is allowed to restart the related services or perform a software restart of the related process only if missing services are detected or a process has stopped. For the DAQ units, the network communication is also monitored and only if the machine is not accessible anymore the operator, logged in expert mode, can perform a hard reset.

The SYSMon

A special agent has been developed to monitor the health parameters of each n_TOF machines (i.e. DAQ units, Proxy and application servers). The agent reads information as CPU, RAM and local disk use as well as CPU, chassis temperatures and fans speed and publishes them through the n_TOF middleware. This information is retrieved by the DIM2WEB and published on the related web monitoring page.

DAQ OPERATIONAL EXPERIENCE

The new n_TOF DAQ system has been put in operation in 2015 and since then has worked reliably 24/7 for roughly ten months each year. The amount of data acquired and transferred to CASTOR for both experimental area is about 1 PB/year. The data transfer to CASTOR optimization implemented in the CASTOR merger has guaranteed to reach, during specific timeslot of heavy data acquisition (i.e., all channels activated at the maximum sampling frequency and without compression), the physical limit of the network bandwidth (i.e., 1.2 GB/s). The operation reliability has been improved thanks to a proper extensive commissioning phase performed during each shutdown period (i.e. Jan-March each year). The diagnostic and management tools developed ad hoc allowed to the operator in shift to prevent failure and/or to quickly resolve most of the issues without asking for an intervention of the experts. This has significantly reduced the n_TOF DAQ downtime. A good estimation of the system uptime can be obtained comparing the number of cumulated protons related to the successfully acquisitions with respect to the total number of protons extracted to the n_TOF target. This ratio gives for the n_TOF DAQ 2017 operation an uptime of 98.9 %.

REFERENCES

[1] The neutron time-of-flight facility at CERN, `https://ntof-exp.web.cern.ch`

[2] CASTOR, `http://castor.web.cern.ch`

[3] J. Serrano *et al.*, "Nanosecond level UTC timing generation and stamping in the CERN's LHC", in *Proc. ICALEPCS'03*, Gyeongju, Korea, Oct 2003, pp. 119-121.

[4] A. Guerrero *et al.*, "CERN front-end software architecture for accelerator controls", in *ICALEPCS'03*, Gyeongju, Korea, Oct 2003, pp. 342-344.

[5] DIP and DIM, `http://readthedocs.web.cern.ch`

[6] A. Masi *et al*, "n_TOF DAQ raw data file format proposal in CASTOR", unpublished.

[7] A. Masi and R. Losito, "LHC collimators low level control system", *IEEE Trans. Nucl. Sci.*, vol. 55, no. 1, pp. 333–340, Feb. 2008.

[8] P. Žugec *et al.*, "Pulse processing routines for neutron time-of-flight data", *Nuclear Instruments and Methods in Physics Research Section A (NIM A) 812,* pp. 134-144, Jan. 2016

ISBN: 978-3-95450-193-9 doi:10.18429/JACoW-ICALEPCS2017-THPHA195