ISBN: 978-3-95450-222-6 ISSN: 2673-5350 doi:10.18429/JACoW-IBIC2020-FRAO05

HIGH PERFORMANCE DATA ACQUISITION FOR A MODERN ACCELERATOR*

G. Shen, N. Arnold, T. Berenc, J. Carwardine, E. Chandler, T. Fors, T. Madden, D. Paskvan, C. Roehrig, S. Shoaf, S. Veseli, Argonne National Laboratory, Lemont, USA

Abstract

Modern technology provides great potential to acquire large amounts of accelerator data, and the possibility to fine tune the particle beam. The wide use of embedded controllers, like field-programmable gate arrays (FPGA), enables collection of fast data from technical subsystems for monitoring, statistics, diagnostics or fault recording. This also presents a number of challenges related to the data acquisition and data management of accelerator data. As a part of the APS Upgrade project, a general purpose data acquisition (DAQ) system is under active development. The DAQ system interfaces with a number of technical subsystems to provide time-correlated and synchronously sampled data acquisition for commissioning, performance monitoring, troubleshooting and early fault detection. This paper will present the status update for the DAQ system, as well as its use cases at APS.

INTRODUCTION

State-of-the-art embedded controllers (microcontrollers, SoCs, FPGAs, DSPs, etc.) have a plethora of resources to implement a high level of functionality tightly coupled to the technical system equipment. A common use of these resources is to utilize large memory buffers to collect fast data for statistics, diagnostics or fault recording. Each embedded controller may contain several gigabytes of memory for such purposes. This presents a challenge: how does one collect, transfer, and utilize this large amount of data from numerous controllers without affecting normal operations? This challenge must be considered early on in the design cycle of a modern accelerator to ensure the solution is well integrated into the control system.

The original design for the APS-U DAQ system was presented in [1]. This paper describes several updates made during the final design phase.

APS-U DAQ SYSTEM

As presented in [2] and illustrated in Fig. 1, APS-U control system is based on EPICS [3], which is widely used in the accelerator control community. The DAQ system is marked with a red dashed line, and its main components are illustrated in Fig. 2.

Wherever possible DAQ IOCs will directly interface with the technical system hardware where the acquisition is performed. In cases where direct connection to the technical system hardware is not feasible, a system specific "DAQ Front-end" will be used between the technical system and DAQ IOC. Each DAQ IOC will be somewhat customized for a given technical system but will also utilize a common framework for capturing and transferring time-series data. This framework will support the transmission of data across dedicated subnets, either directly to services for continuous data collection, or to services responsible for distribution of data to storage, external services, or applications prepared to accept and analyse this data.

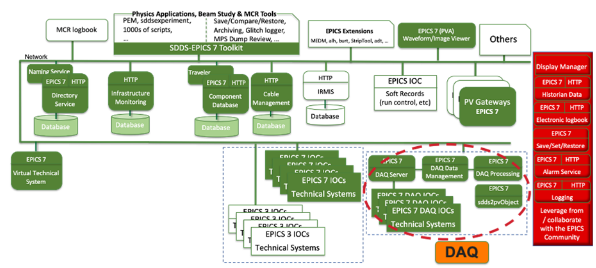

Figure 1: APS-U control system architecture and its DAQ system.

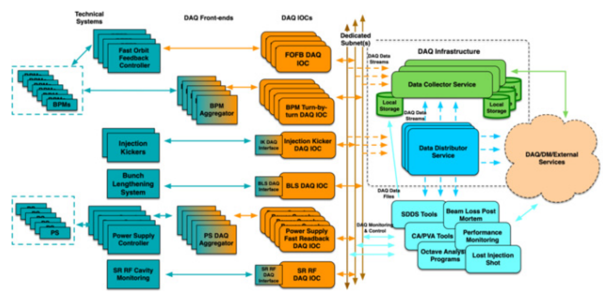

Figure 2: APS-U DAQ system main components.

Key aspects of the Data Acquisition (DAQ) System include:

- Capability to acquire time-correlated synchronously sampled data from several subsystems at various sample rates and correlate this data to within one beam revolution (3.6 μs) or better
- Most DAQ data includes a timestamp for each sample acquired allowing immediate plotting of data from various systems onto a common time-axis
- Support for continuous acquisition or triggered acquisition limited only by available storage
- The ability to route the data to any number of applications
- Use of PV Access [4,5] DAQ objects to encapsulate numerous signals to ensure data synchronicity
- Scalability by partitioning the heavy traffic on dedicated and multiple subnets and servers

* Argonne National Laboratory's work was supported by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences, under contract DE-AC02-06CH11357.
† gshen@anl.gov

ISBN: 978-3-95450-222-6 ISSN: 2673-5350 doi:10.18429/JACoW-IBIC2020-FRAO05

- Separation from operational systems (networks, processors, servers) to allow troubleshooting & enhancements during user operations

R&D WORK

There has been a significant amount of progress on the DAQ system during the APSU R&D phase. Five prototype DAQ IOCs were deployed, as well as a number of services and utilities for saving, viewing and processing data. This demonstrated time-correlated acquisition across different subsystems with different data sampling frequencies. Figure 3 illustrates an example for DAQ data collected from two different subsystems: BPM turn-by-turn (TBT) and FOFB DAQ systems.

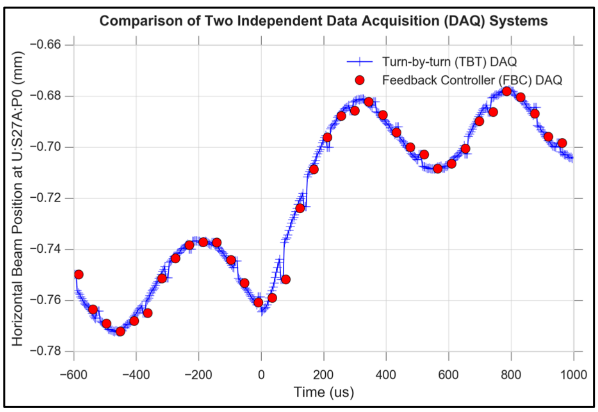

Figure 3: DAQ data example from different subsystems.

In order to provide users with the ability to easily view the DAQ data in real time, we developed an EPICS7 data viewer (see Fig. 4), a scope like application that receives streaming data over EPICS7 PV Access protocol. The viewer allows users to select and display several EPICS7 channels, as well as to perform simple processing in real time (FFT, histogram PSD, etc.).

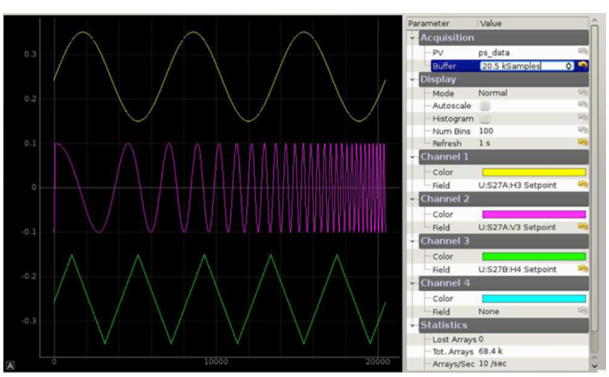

Figure 4: Scope-like-application for APS-U DAQ system.

Over the last few years, these R&D systems have been used extensively for machine studies, diagnostics and troubleshooting. One such example was recent work on the identification and suppression of a troubling 147 Hz vibration source in the APS storage ring. The DAQ system was used to capture 400 channels of continuous data for 20 seconds which was used to identify the sector where the vibration was occurring. Inserting shims to tightly secure a vacuum chamber mount eliminated the vibration [6].

DESIGN UPDATES

One significant architecture change that was made during the final design phase was introduction of the double-sector servers for running turn-by-turn and power supply IOCs. The initial plan involved running a linux-arm TBT DAQ IOC on a micro-TCA FPGA card, and a linux power supplies DAQ IOC hosted on a micro-TCA CPU. The fast orbit feedback IOC was hosted on a linux server. The revised plan calls for hosting all three IOCs on a local linux server (every double-sector), but also requires an FPGA-based data aggregator for both turn-by-turn and power supplies subsystems.

The main advantages of this approach are the common platform used for the three major DAQ subsystems, and also the cost-effectiveness of commodity hardware.

The primary hardware components of the DAQ system are illustrated in Fig. 5. A server will be positioned at each double sector to receive and process the data from three nearby technical systems: Power Supply Fast Monitoring (PS DAQ) data, Fast Orbit Feedback (FOFB) data and BPM TBT data. These double-sector servers isolate much of the "raw data" to local switches and also provide the first level of significant storage for storing long duration acquisitions. Other "one-of" DAQ systems (e.g. LFB, TFB, BLS, …) are connected where most convenient.

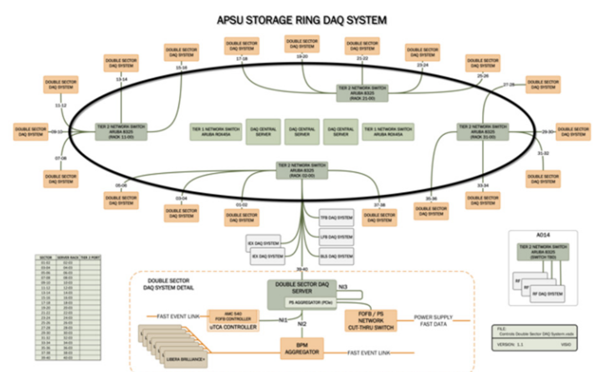

Figure 5: APS-U DAQ hardware architecture.

Additional DAQ servers are in the main Computer Room. These "DAQ Central Servers" are used for data aggregation and user applications.

The components involved with data flow and processing of the DAQ data are numbered & identified in Fig. 6. These are the DAQ interface to the technical system (1), the acquisition and timestamping of the collected data (2), the DAQ IOCs (3), the communication medium and protocols between the IOCs and the services (4), the data collection and distributor services (5), and finally the client API and CLI interfaces for applications that consume the data (6). Also shown is the "processing stack" (7) that illustrates the layers of computing power available for data processing and manipulation. The following

9th Int. Beam Instrum. Conf. IBIC2020, Santos, Brazil JACoW Publishing
ISBN: 978-3-95450-222-6 ISSN: 2673-5350 doi:10.18429/JACoW-IBIC2020-FRAO05

sections give additional details and suggested implementations of these components.

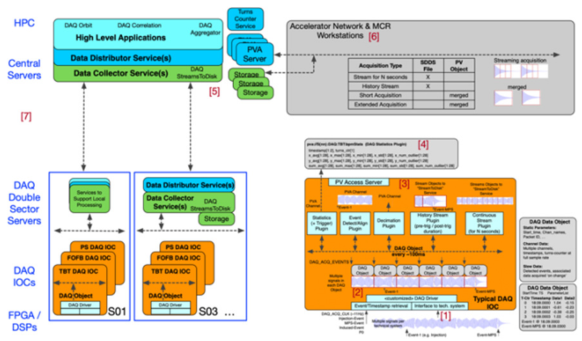

Figure 6: APS-U DAQ system data flow and processing.

As indicated in the bottom-right corner of the Fig. 6, the DAQ IOCs use Area Detector framework [7], which gives us enormous amount of flexibility in terms of processing plugins. The driver gathers data from the technical subsystem, assembles DAQ objects that are then passed through various processing and streaming plugins. The streaming plugins in particular are responsible for actually sending data directly to the data collector service. The history streaming plugins keep circular buffers of data and enable the system to collect data before and after a significant event happens. This can be incredibly useful tool for troubleshooting things like causes of beam dumps or other problems.

Note that all DAQ IOCs have a PVA server that is used to expose DAQ objects on the PV Access channel which can be retrieved using standard EPICS7 clients or APIs.

In addition to various high-level monitoring and processing applications, there are two DAQ services deployed in each double sector, which are 1) data collector service and 2) the distributor service respectively. The data collector service is responsible for storing data to local storage, while the distributor service forwards data to various client applications or other DAQ, data management, or external services.

DATA ACQUISITION MODES

Figure 7 shows a slightly bigger picture of the DAQ IOC and illustrates several DAQ acquisition modes.

The start of each DAQ object is synchronized across IOCs using the fast event system. Each DAQ object contains ~90ms of fast sampled data, slow data that changes infrequently (compared to the fast sampling rate) and static parameters of interest.

The first acquisition mode is continuous streaming of DAQ objects to services and applications. The second one is data acquisition with history, where DAQ objects are stored into a set of circular buffers and streamed to the collector service after a significant preconfigured event happens. The third mode is event based; it involves detecting specific events, aligning data to those events, and presenting this data to clients via designated PV Access channels.

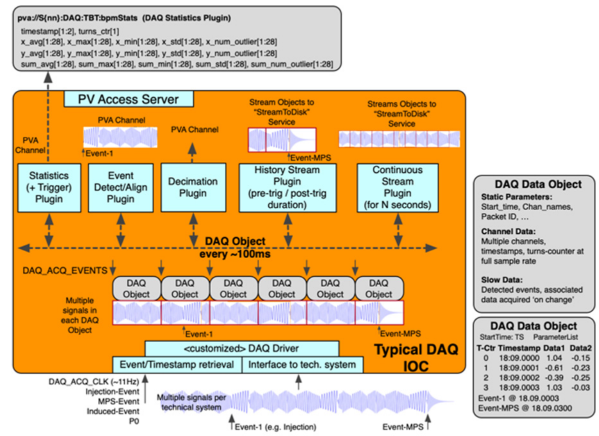

Figure 7: DAQ IOC and acquisition modes.

PERFORMANCE CONSIDERATIONS

As mentioned above, there will be 20 double-sector DAQ servers distributed around the storage ring, each hosting the DAQ IOCs for the turn-by-turn BPM DAQ, the power supply DAQ and the fast orbit feedback DAQ. Because of the numerous channels and fast acquisition, both the turn-by-turn DAQ and power supply DAQ include FPGA-based “aggregators” which receive streams from the technical system hardware and repackage it for efficient processing in the DAQ IOC. Note that just those three double sector subsystems taken together represent more than 13000 channels of data that can be collected simultaneously (see Table 1).

Table 1: Number of Channels of APS-U DAQ System

Technical System DAQ	**# of Channels**
Fast Orbit Feedback	256 x 20
Power Supply Fast Monitoring	354 x 20
BPM Turn-by-turn	84 (typical) x 20
IX/EX Waveforms	8 x 2
Bunch Lengthening System	9 (typical)
Longitudinal Feedback	1
Transverse Feedback	2
SR RF (current)	6
SR RF (near term)	10 (typical) x 4
SR RF (future)	5 (typical) x 12
Booster RF (5)	14, 10, 10, 8, 8 (typical)
PAR RF (3)	9, 12, 10 (typical)
Total # of Channels	**> 13,000**

Several DAQ systems demand significant performance of the network and server components. Table 2 lists the anticipated data bandwidth for each DAQ system. A series of performance tests have been completed in order to

ISBN: 978-3-95450-222-6 ISSN: 2673-5350 doi:10.18429/JACoW-IBIC2020-FRAO05

ensure that all IOCs running on a DAQ double sector server can receive data from the technical systems at maximum rates.

The total double sector data rate is about 144MB/s, while the cumulative data rate for the entire DAQ system is over 3 GB/s. Note that generated Turn-By-Turn data could also include up to 13 additional signals per BPM, or up to 364 signals per double sector that were not counted in the Table 2 analysis. This could potentially result in almost 395 MB/s of additional TBT data per double sector, or almost 8GB/s of data for the entire system.

Table 2: DAQ System Anticipated Data Rates

Technical System DAQ	# of Channels	Sample Rate	Anticipated Data Rate (Per IOC)
Fast Orbit Feedback	256 (x 20)	22.6 kSPS	24MB/s
Power Supply Fast Monitoring	354 (x 20)	22.6 kSPS	32.3MB/s
BPM Turn-by-turn	84 (x 20)	271 kSPS	94MB/s
IX/EX Waveforms	8 (x 2)	4GSPS for 50ns	2.7MB/s (aperiodic)
Bunch Lengthening System	9	2.44MSPS	112.2MB/s
Longitudinal Feedback	1	352 MSPS	48MB/s per acquisition (2 acquisitions/s max)
Transverse Feedback	2	352 MSPS	48MB/s per acquisition x 2 (2 acquisitions/s max)
SR RF	6	271 kSPS	9.2MB/s

HARDWARE CONSIDERATIONS

Aggregators

Two of the DAQ systems, the TBT DAQ and the PS Fast Monitoring DAQ, require "aggregators" between the technical system and the DAQ servers. The platform for these aggregators will be determined during the final design phase.

DAQ Servers

For the purpose of estimating DAQ hardware needs, we assume that each DAQ server supports 1GB/s of combined (read/write) data throughput to/from the local storage. With today's hardware, this should be achievable with a mid-end Linux server: dual 6 physical core CPU with 128GB RAM, at least two 10Gbps networking interfaces, and local SSD storage in a RAID 0+1 (mirrored stripe set) with a PCI Express x4 controller card (each PCI-e lane supports 250MB/s data rate).

CONCLUSIONS

APS-U DAQ is a time-correlated data acquisition system. Among other things, its main features are the ability to acquire data from multiple systems at different sample rates, support for continuous data acquisition, and ability to collect multiple signals within a single IOC. During R&D phase, a significant amount of development has been done, and the system has already seen extensive use during machine studies for diagnostics and troubleshooting.

During final design phase we made several updates to the system design, and we are now actively working towards its deployment into production.

ACKNOWLEDGEMENTS

The authors would like to thank all colleagues at APS facility and APS-U project for their fruitful discussions.

REFERENCES

[1] S. Veseli *et al.*, "Data Acquisition System for the APS Upgrade", in *Proc. ICALEPCS'19,* New York, NY, USA, Oct. 2019, pp. 841-846. doi:10.18429/JACoW-ICALEPCS2019-TUDPP02

[2] G. Shen *et al.*, "High-Level Application Architecture Design for the Aps Upgrade", in *Proc. ICALEPCS'19*, New York, NY, USA, Oct. 2019, pp. 1436-1440. doi:10.18429/JACoW-ICALEPCS2019-WEPHA143

[3] EPICS, https://epics-controls.org

[4] EPICS4 documentation, http://epics-pvdata.sourceforge.net/literature.html

[5] PVA specification, http://epics-pvdata.sourceforge.net/pvAccess_Protocol_Specification.html

[6] L. Emery, pprivate communication, 2018.

[7] Area Detector documentation, http://cars.uchicago.edu/software/epics/areaDetectorDoc.html