ISBN: 978-3-95450-230-1 ISSN: 2673-5350 doi:10.18429/JACoW-IBIC2021-WEPP33

THE FACET-II DATA ACQUISITION SYSTEM*

S. J. Gessner†, SLAC National Accelerator Laboratory, Menlo Park, CA, USA

Abstract

The Data Acquisition System (DAQ) at FACET-II is designed to address the challenge of collecting synchronized, time-stamped data from a variety of diagnostics spread throughout the kilometer-long linac and experimental area. The EPICS control system is used to read out data from devices at FACET-II via Channel Access (CA) over the network. This poses a problem for collecting image data at the 30 Hz beam rate. With image sizes ranging from 0.3-10 Megapixels, the data rate from a single camera can be as high as 0.6 Gbps and there are nearly 100 cameras deployed at FACET-II. Simultaneous image acquisition from just a few of these cameras would overwhelm the network. The FACET-II DAQ solves this problem by coordinating the camera IOCs to write their image data to network-attached storage (NAS) and then validating time-stamps to confirm synchronization.

INTRODUCTION

FACET-II [1] at SLAC uses EPICS [2] to control and retrieve data from diagnostics and devices spread throughout the kilometer-long accelerator and experimental area. The principle diagnostics for experiments at FACET-II are digital cameras. FACET-II hosts nearly 100 digital cameras and more than half of them are dedicated to the experimental area. Communication with the cameras is handled by the Channel Access (CA) network protocol. CA is the appropriate tool for monitoring and managing devices, but data rates are limited by the network. This poses a challenge for data acquisition (DAQ) because we require simultaneous images from multiple cameras corresponding to a single shot. The beam rate at FACET is 30 Hz and the average image size is over 1 Megapixel, equal to 2 MB/frame. FACET users commonly request datasets utilizing ten or more cameras, which yields a data rate of $\approx$ 5 Gbps. This rate can easily overwhelm the network, which utilizes both 1 GbE and 10 GbE switches, and may disrupt accelerator applications.

The FACET-II DAQ avoids the problem of acquiring data over the network by coordinating camera Input-Output Controllers (IOC) to save their data to a Network Attached Storage (NAS) drive located on the same switch as the camera servers. With this approach, the camera data never leaves the switch and network traffic is maintained at an acceptable level.

HARDWARE, LAYOUT, AND TOPOLOGY

The majority of cameras at FACET are AVT Manta and Mako GigE-interface devices running on Advantech Sky 8201 servers. Each server can host up to 12 cameras. FACET also uses scientific-CMOS cameras in the experimental area. All servers run areaDetector [3] IOCs. There are 7 camera servers along the linac and 9 camera servers located in Sector 20 (S20), the experimental area. The camera servers in S20 collect the majority of the data requested by users via the DAQ. A control server called facet-srv20 and NAS drive are also located in S20 and connected to the same switch as the camera servers. The NAS uses bonded ports to maximize data rate onto the drive. The topology of the network in S20 is shown in Fig. 1. The layout of the machine is shown in Fig. 2 and the distribution of the cameras is shown in Table 1.

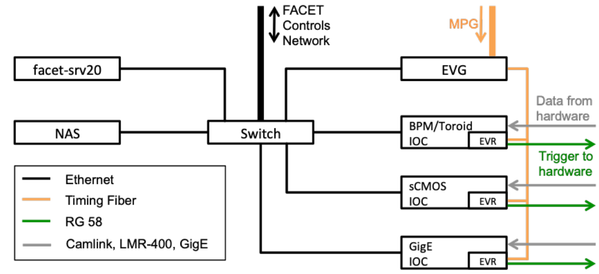

Figure 1: Network topology in FACET Sector 20. Note that the EVG has been relocated to the FACET injector in Sector 10.

DATA SOURCES AND SYNCHRONIZATION

In addition to camera images, the FACET-II DAQ also collects data from a variety of sources. This includes devices that update at the beam rate, such as beam position monitors (BPM), toroids, and photo-multiplier tubes (PMT), as well as devices that update asynchronously, such as magnet currents, motor positions, and temperatures. The beam-synchronous devices produce scalar data and are included in SLAC's beam-synchronous acquisition (BSA) infrastructure [4]. BSA devices store data in a rolling buffer with 2800 events. The length of the buffer was chosen for compatibility with the legacy SLC Control Program (SCP) [5]. The buffer exists on the IOC hosting the device, and each IOC can support multiple buffers with unique event definitions. For example, if the nominal beam rate at FACET is 30 Hz, an event code (EC) may be used to specify a lower rate (e.g. 10 Hz) or a rate subject to conditions (e.g. beam destination). Events are timestamped and tagged with a pulse ID number. The pulse ID is a rolling value that updates at 360 Hz. The maximum pulse ID is 131040, which is also a legacy of the SCP. This implies that the pulse IDs are not unique, but a buffer of length 2800 will not contain duplicate pulse IDs. BSA buffers are centrally managed by the Event Generator

* This work performed under DOE Contract DE-AC02-76SF00515.
† sgess@slac.stanford.edu

ISBN: 978-3-95450-230-1 ISSN: 2673-5350 doi:10.18429/JACoW-IBIC2021-WEPP33

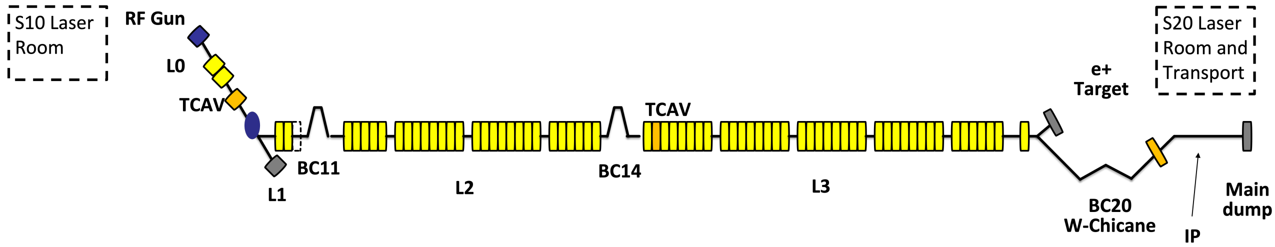

Figure 2: Layout of the FACET-II accelerator and experimental area.

Region	S10 Laser	Injector	L1	L2	BC20	S20 Laser	IP Area	Spectrometer	Total
Cameras Deployed	7	3	6	2	2	21	18	12	76
Cameras Planned	10	5	10	6	2	21	24	16	94

Table 1: Distribution of Digital Cameras at FACET-II

(EVG). Configuring an event buffer on the EVG configures the event buffer on all BSA devices at FACET.

The camera IOCs are not part of the BSA infrastructure because they cannot hold 2800 images in RAM. However, each IOC has an Event Receiver (EVR) which receives timing information from the EVG [6]. The IOC tags the camera image with pulse ID. The timestamp and pulse ID are included with the image when requested over CA and are embedded in the image metadata when written to disk. The pulse ID is used to compare and temporally align images across different cameras and with the BSA data.

DAQ SOFTWARE

The DAQ software is written in MATLAB R2020A and includes a GUI interface created with MATLAB's App Designer tool. The GUI allows users to interact with the FACET DAQ API [7]. The GUI prompts users to select the desired data rate (Event Code) and number of shots. The maximum number of shots is determined by the amount of RAM required store images on the camera servers with a cap set at 2800 (same as the BSA buffer). Users select cameras to include in the DAQ, organized by their region in the accelerator and experimental area. The user can include BSA and non-BSA devices, also organized by region. The BSA and non-BSA device lists are easily edited by users, so new devices can be included in the DAQ as needed. The DAQ has a single-step mode which acquires data for the given number of shots. It also has scan modes in which a parameter is varied and data is acquired for each parameter setting. The GUI interface allows the user to select one or two dimensional scans, but API is written such that n-dimensional scans are possible by flattening the scan region into a 1D array indexed by a single step number.

The DAQ contains various validation and "safe-fail" methods. At the start of a DAQ run, CA is used to verify the PVs in the BSA and non-BSA lists. If a CA call to a PV in the list does not return a value (either because it is mis-specified or the IOC is offline), it is removed from the list. The DAQ checks the connection status of the cameras before proceeding. Temporary camera failures are common due to the high-radiation environment. The DAQ identifies these failures when present and attempts to reboot the affected cameras. An "Abort" method exists so that the user can preemptively end the DAQ run.

The DAQ is run from the control server (facet-srv20). The control server and all camera servers mount the NAS drive. At the start of a DAQ step, the control server creates a folder on the NAS for the new dataset, and subdirectories for each of the cameras. Next, the DAQ configues the NDFileTIFF areaDetector plugin with the desired number of shots and the location to save the data. At this point, the camera triggers are disabled before data taking starts. The DAQ then configures the BSA buffer with the desired EC to match the cameras. The data taking begins when a command is sent via CA to the EVG to open the buffer and camera triggers are re-enabled. The DAQ starts a timer based on the number of images requested and the event rate. The camera IOCs start writing data to disk as soon as they start receiving frames. At the same time, the BSA buffers fill with data. Both processes are local to the IOCs and do not pass data over the network. During this time, the DAQ acquires non-BSA data at 1 Hz using CA. The DAQ also checks for user aborts during this time.

When the timer ends, a command is sent to the EVG to close the BSA buffer. The DAQ waits for all image data to be written disk. Writing data to disk may continue after data taking has ended if the image buffers are backed up. The DAQ also checks for camera failures that may have occurred during the DAQ step. Once image writing is complete, the DAQ retrieves the BSA buffers via CA and stores them in a MATLAB data structure. If the DAQ is running a scan, it sets the next parameter value and the process repeats.

When the scan or single-step is complete, the DAQ performs synchronization and quality control. The DAQ loops through all of the camera images and extracts the pulse IDs which are compared across cameras and to the pulse IDs in the BSA buffers. The DAQ assigns a timestamp and unique ID (the pulse IDs are not unique) to all data, and compiles

ISBN: 978-3-95450-230-1 ISSN: 2673-5350 doi:10.18429/JACoW-IBIC2021-WEPP33

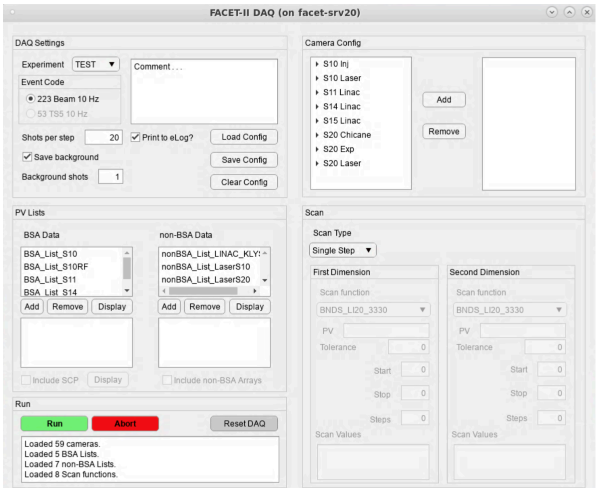

Figure 3: The FACET-II DAQ interface.

a list of common shots. In a typical DAQ with no camera failures, 90-100% of shots from all devices will be matched. The list of matched is added to the data structure, along with the paths to the image data on the NAS. The image data itself is not stored in the data structure. The DAQ ends by writing the data structure to the NAS as a MATLAB file and as an HDF5 file. The HDF5 format is convenient for analyzing data in Python, which is a popular language amongst FACET users. Finally, the DAQ writes a message to the FACET eLog which includes the location of the data and the success rate of the matching process.

Users perform online analysis of the data using the Data Analysis Tool (DAN). The DAN is a separate MATLAB GUI and API. The DAN includes image analysis, correlation plots, and a motivation indicator. For offline analysis, the data is transferred from the NAS to SLAC Central Computing by a daily cron job. We expect to collect roughly 25 TB of data per year with the FACET-II DAQ.

IMPROVEMENTS TO THE DAQ

There are several improvements planned for the DAQ, mostly related to the synchronization procedure. At present the DAQ sequentially starts the BSA buffer and then enables the trigger for each of the cameras with CA commands. The time required to send these command to the IOC is a few milliseconds at worst, and usually the cameras all start and stop taking data on the same event. A more efficient process is to create an EC that is used by both the BSA buffers and the cameras. This EC will be identical to the usual ECs that specify data rate, but can be globally enabled or disabled on the EVG. With this approach all devices participating in the DAQ will receive and identical set of triggers and see the same events.

Another improvement to synchronization is the planned switch to the event2 module which will provide more robust timestamping. The event2 module is aware of both trigger delay and camera exposure. This helps eliminate off-by-one errors and flags events when the frame was returned outside of the expected window. In addition, the event2 module returns the camera pulse IDs as an array and eliminates the time-consuming process of extracting pulse IDs one-by-one from image metadata.

Additional improvements to the DAQ include the capability to acquire BSA and non-BSA waveforms, as well as legacy SCP data. A number of devices such as linac BPMs, klystrons, and magnets are controlled by the SCP and are not available through EPICS. We plan to install a dedicated server in S20 for running the DAN.

CONCLUSION

We have developed a data acquisition system capable of recording synchronized from multiple cameras at high data rates. The DAQ enables plasma wakefield accelerator science at FACET, as well as machine learning studies and general accelerator R&D. The FACET-II DAQ will continue to evolve in order to meet the needs of the users.

REFERENCES

[1] V. Yakimenko *et al.*, "FACET-II facility for advanced accelerator experimental tests", *Phys. Rev. Accel. Beams*, vol. 22, no. 10, p. 101301, 2019.

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

[3] areaDetector, M. Rivers, ANL. `https://areadetector.github.io/master/index.html`

[4] K. H. Kim, S. Allison, T. Straumann, and E. Williams, "Real-Time Performance Improvements and Consideration of Parallel Processing for Beam Synchronous Acquisition (BSA)", in *Proc. ICALEPCS'15*, Melbourne, Australia, Oct. 2015, pp. 992–994.

[5] N. Phinney, "Report on the SLC Control System", in *Proc. PAC'85*, Vancouver, Canada, May 1985, pp. 2117–2120.

[6] J. E. Dusatko, S. Allison, J. Browne, and P. Krejcik, "The LCLS Timing Event System", in *Proc. BIW'10*, Santa Fe, NM, USA, May 2010, paper TUPSM083, pp. 379–383.

[7] FACET DAQ, S. Gessner, SLAC. `https://github.com/slaclab/facet-matlabTNG/tree/master/F2_DAQ`

[8] FACET DAN, H. Ekerfelt, SLAC. `https://github.com/slaclab/facet-matlabTNG/tree/master/F2_DAN`

[9] event2, B. Hill, SLAC. `https://github.com/slac-epics/event2`