The DAQ for the Single-Phase DUNE Prototype at CERN

Roland Sipos

***Abstract*— DUNE will be the world's largest neutrino experiment and is due to take data in 2025. This paper describes the data acquisition (DAQ) system for one of its prototypes, ProtoDUNE-SP, which started taking data in Q4 of 2018. ProtoDUNE-SP also breaks records as the largest beam test experiment yet constructed and is a fundamental element of CERN's Neutrino Platform. This makes ProtoDUNE-SP an experiment in its own right and the design and construction have been chosen to meet this scale. Due to the aggressive construction and commissioning timescale, off-the-shelf electronics have been chosen to meet the demands of the experiment where possible. The ProtoDUNE-SP cryostat comprises two primary subdetectors: a single-phase liquid-argon time projection chamber (LAr TPC) and a companion photon detector system. The TPC has two candidate readout solutions under test in ProtoDUNE-SP: RCE (ATCA-based) and Front-End Link EXchange (FELIX) (PCIe-based). Fermilab's artdaq is used as the data handling software framework for the experiment. Custom timing and trigger electronics and software are also described. Triggering and lossless compression will take the 430 Gb/s of data from the front end and reduce it to 20-Gb/s bandwidth to permanent data storage in CERN's EOS infrastructure.**

***Index Terms*— Data acquisition (DAQ), detector, field-programmable gate array (FPGA), high-throughput, liquid-argon time projection chamber (LAr-TPC), online computing, trigger.**

I. Introduction

THE ProtoDUNE-SP [1] time projection chamber (TPC) volume and its primary components are illustrated in Fig. 1. The active volume is 6 m high, 7 m wide, and 7.2 m deep (along the drift direction). The active volume consists of three central cathode plane assemblies (CPAs) and six anode plane assemblies (APAs). This represents 4% of a DUNE Far detector single-phase module. The anode planes on either side of the detector are constructed of three adjacent APAs that are each 6 m high by 2.3 m wide. Each APA consists of a frame that holds three parallel planes of wires, two induction planes and one collection plane. The induction wires are oriented at a stereo angle of 35.7° with respect to the collection wires to enable 3-D reconstruction of collected ionization tracks. The wire pitch for all wire planes is approximately 5 mm, and each APA holds a total of 2560 wires. The front-end cold electronics, mounted onto the APA frames and immersed in liquid-argon (LAr), continuously amplify and digitize the induced signals on the sense wires at 2 MHz, and transmit these waveforms to the data acquisition (DAQ) system. From the DAQ, the data are transmitted through storage buffers to disk, then to the central CERN Tier-0 Computing Center, and finally to other partner sites for processing and analysis.

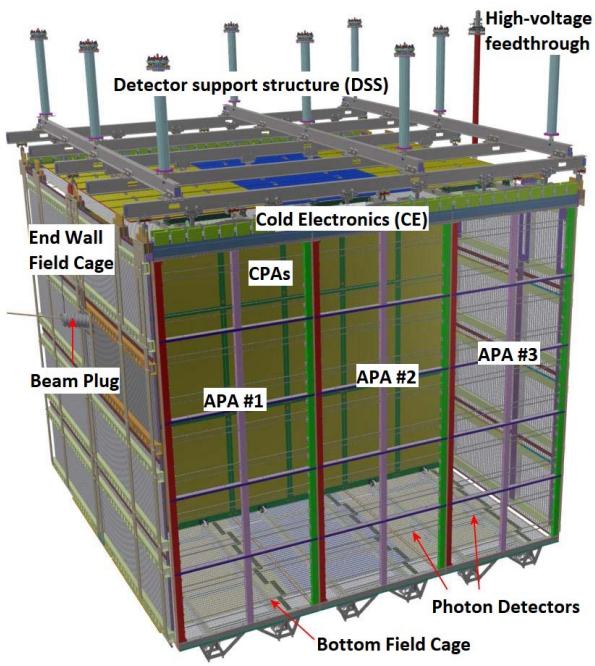

Fig. 1. Particle detecting elements and major components of the ProtoDUNE-SP TPC.

A. Detector Environment

The detector is located on the CERN SPS beam line, which allows for an exposure to charged particle beams in the momentum range of 0.5–7.0 GeV/c. The detector is located on the surface, leading to an additional high cosmic ray flux.

Considering these environmental aspects, the DAQ system has strict requirements. It needs to be a single, scalable system across all subdetectors, supporting different interfaces, without differentiation between subdetector types from an operations point of view. The DAQ system needs to be externally triggered in order to limit the amount of collected data and needs to support partitioning to allow for detector component specific triggering and data handling conditions. The collaboration decided not to employ zero suppression, as one of the goals

Manuscript received December 19, 2018; revised February 26, 2019; accepted March 5, 2019. Date of publication March 20, 2019; date of current version July 16, 2019.

The author is with CERN, 1211 Geneva, Switzerland (e-mail: roland.sipos@cern.ch).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TNS.2019.2906411

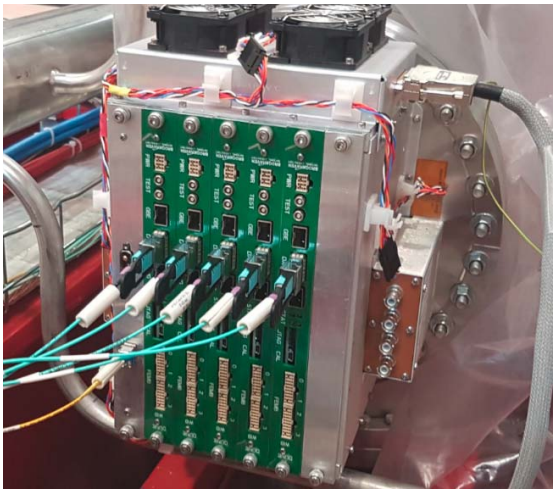

Fig. 2. Five WIBs (each reading out four FEMBs) are connected to one APA through its cold electronics flange; connections to the readout systems are optical to ensure electrical insulation.

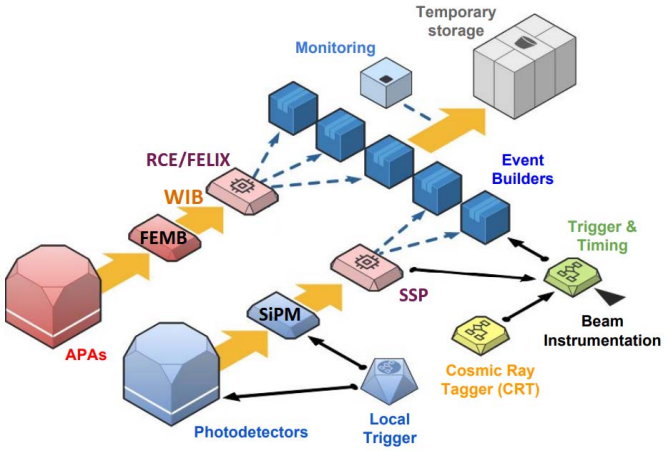

Fig. 3. Overview of the DAQ system, its interconnections, data flow, timing, and trigger signals, and the interfaces to the front-end electronics and offline computing systems. The TPC readout systems (RCE and FELIX) receive data from the WIBs, which is then compressed and selected based on trigger information. Triggered data are then sent to the event building farm and subsequently stored to disk. In parallel, a sample is prepared for online monitoring. Triggers are formed from inputs from the beam instrumentation, photon detection, and CRT systems.

is to acquire an unbiased set of data in order to study suppression methods for DUNE. Another goal is to reach a lossless compression factor of 4, based on signal-to-noise ratio estimates derived from MicroBooNE studies [2]. The DAQ will operate without dead-time under normal conditions, so continuous storage of overlapping ranges of data in space and time needs to be provided.

II. DAQ System

Due to the extremely tight construction and commissioning schedule (data taking is scheduled for Q4 2018), the DAQ has been designed to use commercial off-the-shelf (COTS) components. Hence, most of the implemented solutions are based on existing frameworks and generic technologies. A custom board was developed (as discussed in Section II-D) that accepts inputs from beam instrumentation, cosmic ray tagger (CRT), and photon detectors, and from these inputs assembles a single global trigger in the form of a timestamp that is distributed to the system. The readout system also heavily relies on large buffers to maximize the amount of beam particles that can be analyzed in each spill.

A. Front-End Electronics

There are two different modules that represent the input for the DAQ system, one to read the TPC's APAs and another to gather data from the photon detectors. The WIB interfaces between the cold electronics and the DAQ, with real-time diagnostic possibilities. The cold electronics consist of the ProtoDUNE front-end motherboards (FEMBs), each with four low-voltage differential signaling (LVDS) streams of 56 bytes per 500 ns. This results in a data stream of ~3.6 Gb/s per FEMB. Each WIB multiplexes data from four FEMBs and is connected to the DAQ and slow control systems via optical links. Each APA is electrically insulated inside the cryostat and is only connected to the detector ground through the cold electronics at its own cold electronics flange, where five WIBs receive the APA's data as shown in Fig. 2.

Two systems are used to receive data from the WIBs. The baseline system (discussed in Section III-A) is based on reconfigurable computing elements (RCE) [3], which are housed in industry standard ATCA shelves and are used to read out 5 of 6 APAs. The alternative system (described in Section III-B) is based on the front-end Link EXchange (FELIX) [4] technology and is used to receive the data from the remaining APA.

The TPC is supported by a photon detector system (PDS). Each PDS module consists of a bar-shaped light guide and a wavelength-shifting layer. The readout system of the PDS is a silicon photomultiplier (SiPM) signal processor (SSP) prototype module, which is a high-speed waveform digitizer with 12 channels per module. Each channel contains a fully differential voltage amplifier and a 14-bit, 150-megasamples per second (MSPS) analog-to-digital converter (ADC) with 2.1-Gbps data output. The timing is obtained applying signal processing techniques to the leading edge of the SiPM signal, using the on-board Artix field-programmable gate array (FPGA). It has deep data buffering 13 μs and operates with no dead-time up to 30 KHz per channel. A total of 24 SSPs serves to read out the photon detector modules in all six APAs.

B. Scope and Requirements

The physics requirements of ProtoDUNE-SP are the primary drivers of the DAQ system requirements. The front-end electronics and assumed bandwidth and storage requirements from the online and offline computing systems impose additional constraints. The DAQ system is shown in Fig. 3 along with its interfaces to the cold electronics, beam instrumentation, and offline computing systems.

The data rate from the electronics is dominated by the TPC data. The TPC readout features 15 360 channels as presented in Table I.

TABLE I
TPC READOUT ELECTRONICS COMPONENTS AND QUANTITIES

Element	Quantity	Channels per element
TPC	1	15,360
APA	6	2,560
FEMB	120, 20 per APA	128
WIB	30, 5 per APA	512

The SPS supercycle is not constant but varies between ∼30 and ∼60 s. For estimation, the average data rate is calculated with 24-s SPS cycles in the following manner.

1) The TPC's continuous data rate is ∼430 Gb/s.
2) Two 4.8-s extractions are performed in each 48-s SPS supercycle, which leads to a 20% beam duty cycle.
3) Then, the average TPC data rate is ∼430 Gb/s × 0.2, which results in ∼86 Gb/s.

The maximum bandwidth from the ProtoDUNE-SP online system to CERN's OpenSource Storage (EOS) [5] is 20 Gb/s. Therefore, the DAQ system must reduce the data by a significant fraction before they are sent to offline storage. This is achieved by a combination of data selection and compression.

1) The baseline trigger rate during the SPS spill is taken to be 25 Hz.
2) The size of triggered events with a 5-ms time window around the trigger timestamp is ∼60 MB, if a compression factor of 4 is achieved.
3) This leads to a 60 MB × 25 Hz = 1.5 GB/s (12 Gb/s) data rate during a spill to local temporary storage buffers.
4) Combining this with the beam duty cycle produces an average data rate toward the EOS long-term storage of 12 Gb/s × 0.2 = 2.4 Gb/s.

A combination of externally triggered events and self-triggered events makes up the PDS data. The external triggers come from the beam instrumentation via the trigger system at a rate of 25 Hz. The resulting data rate amounts to 118 Mb/s. The self-triggered data are induced by cosmic rays. The cosmic rate of 10 kHz is assumed, totaling 1106 Mb/s. Therefore, the PDS is estimated to send data at a combined rate of ∼1.2 Gb/s from the 24 SSPs.

The DAQ is equipped with sufficient temporary storage space for raw data files. The system can store up to 3 days worth of data taking, assuming 25 Hz of in-spill triggers and a compression factor of 4, with 300 TB. In order to have a large safety margin, 700 TB of usable disk space is installed.

C. Timing System

The timing system must:

1) provide a stable and phase-aligned master clock to all DAQ components;
2) receive external signals (including triggers) into the ProtoDUNE clock domain and time-stamp them;
3) distribute synchronization, trigger and calibration commands to the DAQ system;
4) and conduct continuous checks of its own function.

In addition, the timing system acts as a data source for its readout software, providing a record of timing signals received, distributed, or throttled. An overview and the connection with the trigger system can be seen in Fig. 4.

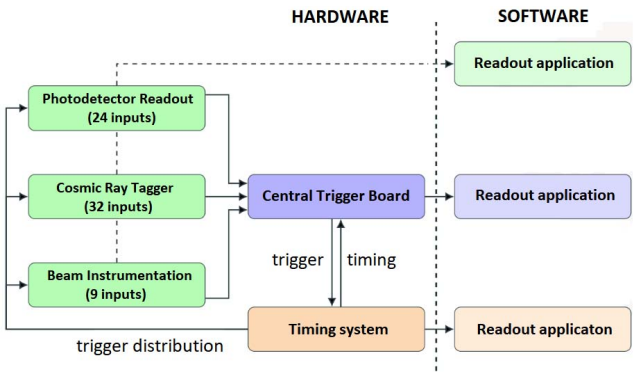

Fig. 4. Timing system provides a stable and phase-aligned master clock to all DAQ components and distributes the trigger signals over the same fabric as the clock and calibration signals.

An FPGA-based master unit receives a high-quality clock (provided in ProtoDUNE by a GPS-disciplined oscillator) and external signals from the trigger system and SPS accelerator. The master unit multiplexes synchronization and trigger commands, along with arbitrary command sequences generated by software, into a single encoded data stream, which is broadcast to all timing endpoints, and decoded into separate clock and data signals. Endpoints receive:

1) phase-aligned clock from external or FPGA phase-locked loop (PLL);
2) synchronous trigger and timing signals, aligned to ∼1 ns across the system;
3) setup and control data in timing packet format.

The same high-quality clock is used for the beam instrumentation timing such that both systems remain synchronized. A uniform phase-aligned cycle counter, updating at the ProtoDUNE system frequency of 50 MHz, is maintained at all endpoints, allowing commands to take effect simultaneously at all endpoints regardless of cable lengths or other phase delays. The system uses duplex links, allowing all endpoints to be regularly interrogated during system operation to verify correct operation and reception of timing commands.

D. Trigger System

The central trigger board (CTB) is housed in the ProtoDUNE DAQ racks as shown in Fig. 5. It assembles information from the beam instrumentation, photon detectors, and the CRT, and broadcasts trigger decisions to the timing network. The CTB is designed around the MicroZed system-on-chip (SoC) board, equipped with a Xilinx Zynq7020. The motherboard implements the hardware interface with different systems; the FPGA implements trigger logic and interfaces with timing; and finally, the CPU and software elements manage the FPGA configuration and communication with DAQ software. The CTB has the following inputs for ProtoDUNE.

1) CRT provides 32 trigger signals.
2) Photon detectors provide 24 trigger signals.
3) Beam instrumentation provides nine input signals.

The firmware is designed to be suitably flexible such that the exact trigger selection can be decided at configuration time

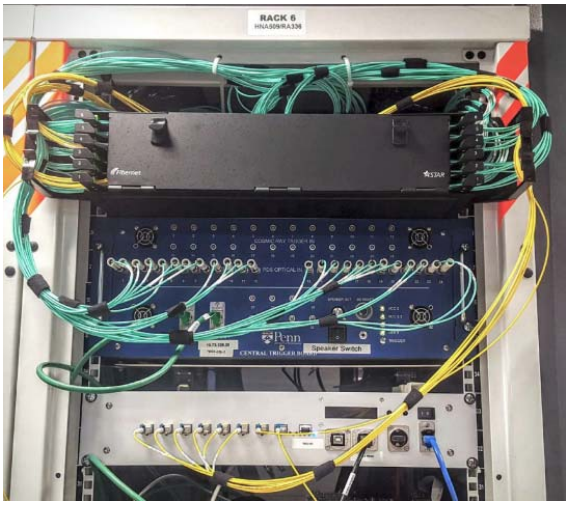

Fig. 5. CTB in the DAQ rack; electrically isolated from the cryostat.

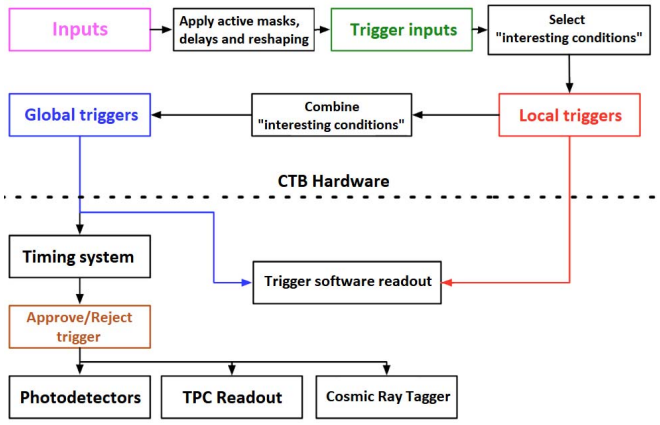

Fig. 6. CTB is designed to receive triggers from various subsystems and combine these into a global trigger based on a configurable input mask (or more sophisticated algorithm, if desired). It provides functionality to globally time-stamp triggers, keep event counts and provide information with trigger type and error conditions.

and does not require new firmware. It will be able to decide on a trigger in well under 1 μs. The general outline and logic of the trigger system are shown in Fig. 6.

III. Readout Systems

There are several flavors of off-detector readout systems in ProtoDUNE.

1) The readout system for the photon detectors as well as for the trigger and timing systems are fully software-based. Dedicated readout applications are responsible for the trigger selection for these systems.
2) For the TPC readout, two solutions are being implemented, to be evaluated and compared: the ATCA-based RCE and the PCIe-based FELIX solution. The aim is to identify strong and weak points of both solutions.

In regard to the TPC readout systems, two firmware variants with different data formats exist (targeting RCE and FELIX) that can be uploaded onto the WIBs' FPGAs. For all readout systems, the software application that interfaces with the event building farm is the so-called BoardReader, shown in detail in Section IV-A. This application type is tailored to read data from specific detector electronics, prepare event fragments, and send them to the event building system.

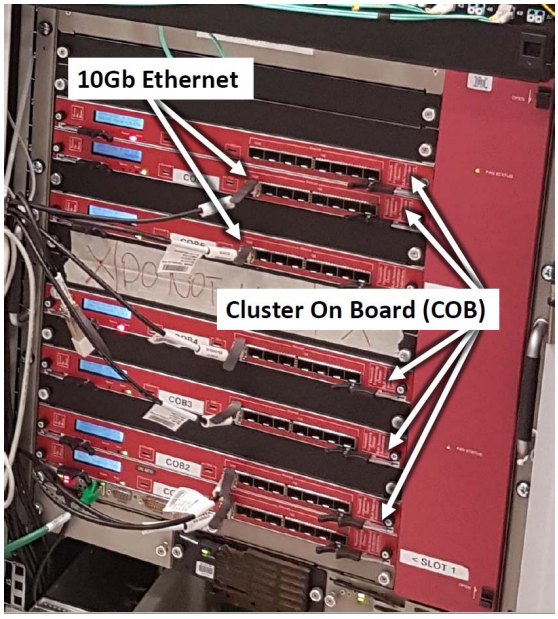

Fig. 7. RCE readout ATCA crate, with seven COBs installed.

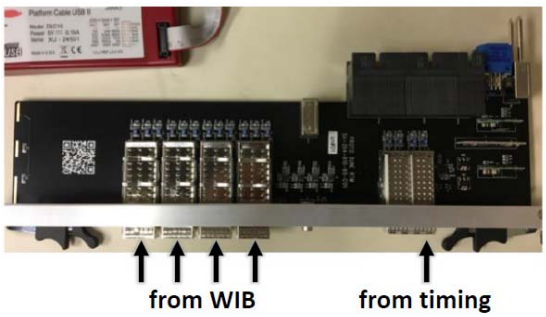

Fig. 8. ProtoDUNE RTM for the RCE readout. This application-specific board uses a set of QSFP transceivers to receive the data from the WIB and an small form-factor pluggable (SFP+) optical interface for communication with the timing and trigger distribution system.

A. RCE Readout

The reconfigurable cluster elements (RCE) [3]-based readout is a full meshed distributed architecture, based on networked SoC elements on the ATCA platform, developed by the SLAC National Accelerator Laboratory. Being the baseline solution, it reads five APAs, for a total of 12 800 channels. The ATCA crate mounted in the DAQ racks is shown in Fig. 7.

Each cluster on board (COB) hosts nine processing elements, each with dual-core ARM A9 processors and 1-GB DDR3 memory. The COB also carries a data processing daughter board with dual Zync 045 FPGAs. The rear transition module (RTM) at the back of the COB supports the experiment-specific interfaces and is shown in Fig. 8.

This particular approach focuses on early data processing, with tightly coupled custom firmware and software implementations. The ProtoDUNE version of the RCE will accept digitized data from the WIBs over optical fiber. The data are sent to a set of processing units where it is checked for

Fig. 9. FELIX BNL-711 PCIe card.

errors, time-stamped, aggregated, formatted, compressed using arithmetic probability encoding, and buffered. Then data are selected and sent to the back-end DAQ over TPC/IP upon the receipt of an external trigger. An output data rate of ∼1 Gb/s per RCE can be sustained. A compression factor of 4 or more is foreseen but contingent on noise.

B. FELIX Readout

One APA (2560 channels) will be read out using the FELIX system. The FELIX [4] is a project initially developed within the ATLAS collaboration at CERN. Its purpose is to facilitate the development of high-bandwidth readout, needed for the high-luminosity LHC, presently planned to start in 2026. Some detector upgrades will already make use of the DAQ system in Run 3 of the LHC, starting in 2021.

The motivation for FELIX is the desire to move away from custom hardware at as early a stage as possible, and instead employ commodity servers and networking. The FELIX design is based on a shared firmware/software solution. A PCIe card is used to stream input data arriving from the detector front ends to a circular memory buffer in a host PC using a continuous direct memory access (DMA) transfer (with fixed 4-kB block size); "publisher" software running on the host PC routes the data to multiple output destinations using network interface cards. The used BNL-711 card with Xilinx Kintex Ultrascale and 48 optical links (MiniPODs) is shown in Fig. 9.

For ProtoDUNE, CPU-based compression routines will be used in the FELIX design (in contrast to FPGA-based compression in the RCEs). This task can be also offloaded from the CPU, with the possible hardware acceleration provided by Intel QuickAssist (QAT) [6] technology.

IV. DAQ Software Layer

The DAQ software layer mostly depends on existing solutions and frameworks in order to meet the aggressive schedule.

A. Data-Flow and the Artdaq Framework

The data-flow software for the ProtoDUNE-SP DAQ is primarily responsible for acquiring the data from the readout electronics, packaging it appropriately, and storing it in files that are local to the DAQ cluster. It is also responsible for other

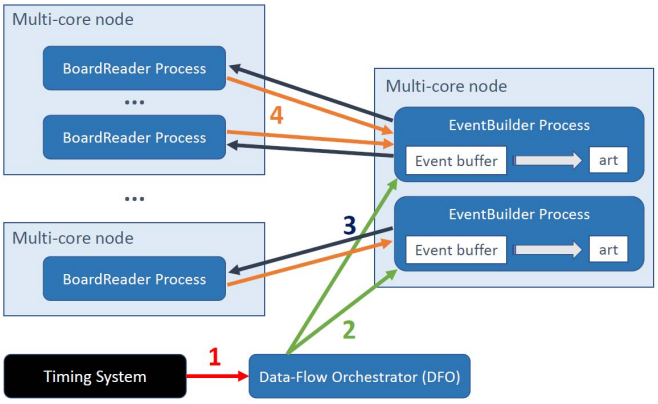

Fig. 10. Data-flow in the artdaq framework. 1: trigger messages are received from the timing system. 2: the DFO forward the trigger messages to the EventBuilders. 3: EventBuilders request data fragments from BoardReaders by trigger timestamp. 4: Data fragments are routed from the BoardReaders to the EventBuilders on the same partition.

functions, including the delivery of configuration parameters to the front-end electronics and real-time monitoring of the data quality and the performance of the DAQ system. It is not responsible for the transfer of the raw data files to permanent storage. The design of the ProtoDUNE DAQ data-flow software is based on artdaq [7], which is a DAQ toolkit developed at Fermilab.

The artdaq framework features the following processes.

1) *BoardReader*—Instances of this process are responsible for communicating with the detector electronics.
2) *EventBuilder*—Instances of this process are responsible for assembling complete events and optionally processing them in art [8].
3) *DataLogger*—Instances of this process are responsible for logging the data and serving events to real-time data quality monitoring (DQM) processes.
4) *Data-Flow Orchestrator*—It is responsible for queuing triggered events and load-balancing them among the EventBuilders.

The underlying hardware will be fully based on COTS servers running Linux CentOS7 and on a 10-Gb/s switched network. Overall, 10–15 servers are sufficient to implement the complete data-flow for ProtoDUNE (BoardReaders and data flow orchestrator, EventBuilder, storage). Fig. 10 shows a possible data-flow based on artdaq.

Specific applications have additional functionalities:

1) SSP BoardReader: The SSP hardware is responsible for recording the waveforms from the photodetectors and making them available to downstream components. The SSP is capable of triggering internally on photon signals, on an external trigger source, and at fixed time intervals, and can also generate trigger signals for other detector components. Using an on-board FPGA, it can calculate metadata related to the observed pulses (heights, widths etc.), which may also be sent in lieu of the actual waveform data in order to reduce data volumes. In turn, the BoardReader program is responsible for receiving data from an SSP unit, packaging it into fragments in time corresponding to those from the other detector systems, and sending it on to the EventBuilder to produce full events. The SSP BoardReader also does the trigger selection, taking

the trigger request from the EventBuilder and assembling a 5-ms window of any internal trigger activity around the external trigger. It also deals with the configuration of the SSP hardware and with state transitions to/from the control system.

2) Timing and Trigger BoardReaders: Besides their specific function on the system, the timing and trigger systems will also provide event data. The timing system will provide the timestamp information of accepted triggers, dead-time information, and so on. The trigger system will provide information about the specific inputs that contributed to trigger decisions. The information from these systems will be dealt with by dedicated BoardReaders, to be folded into the ProtoDUNE-SP events, and, in addition, may be continuously gathered, irrespective of the trigger decisions.

3) RCE BoardReader: As the multiplexed data from the WIB enter the RCE FPGA fabric, it is demultiplexed and buffered into per-channel, fixed-time-length chunks (for instance 512 or 1024 ticks). These chunks are compressed and written to the DRAM where the RCE processor waits for a trigger (also handled by the FPGA) to arrive. Upon a trigger, the processor sends data for a fixed window in time, including pretrigger and posttrigger time chunks for all channels, to the RCE BoardReader applications.

4) FELIX BoardReader: The BoardReader implementation for the FELIX-based readout integrates a network messaging layer that subscribes to the FELIX publisher application. The implementation focuses on flexibility, as the topology of the queues and links is scalable by the BoardReader configuration. In order to achieve the required performance, particular care has been put into the implementation, avoiding dynamic memory allocation. Internal elements and functionalities strictly avoid memory copies. Every link has dedicated subscriber threads that populate single producer single consumer queues, using a lock-free implementation. Multithreading features ensure the proper synchronization of the trigger matching threads and also to comply with the internal state machine of artdaq. The compression library relies on the DEFLATE [9] algorithm, with possible hardware acceleration.

B. Run Control

The run control software for the experiment must:

1) provide process control functionality; launch and terminate DAQ applications;
2) provide a finite-state-machine (FSM), to ensure the proper sequence of process control steps;
3) be the interface to operators;
4) provide support to include, exclude, and partition a given part of the system.

The run control software is based on the Joint COntrols Project (JCOP) [10] extension for the WinCC-OA supervisory control framework by Siemens. It is in common use at LHC experiments, and it is officially supported by CERN. The JCOP system, interfaced with the artdaq system, satisfies the above-mentioned requirements and carries with it the added benefit of expertise and support available at CERN. Monitoring metrics for process statuses, errors, alerts, and trends is developed with such a system for monitoring both the data-flow and detector status. JCOP is also used for the slow control system, providing a uniform interface to detector operation. The run configuration utilities for the DAQ provide:

1) configurations for DAQ system architecture (describing the data flow), DAQ hardware, and DAQ software;
2) run numbers, as well as the assignment of a user-defined configuration to each run;
3) partitioning support;
4) an UI for configuration generation/viewing;
5) and a means to look up previous run configurations.

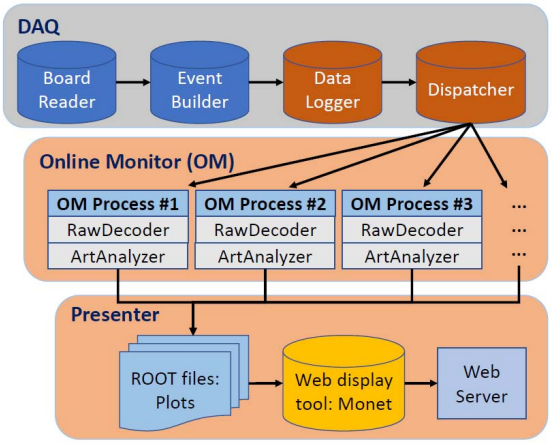

Fig. 11. data-flow in the online monitoring system, highlighting the path of events after they are received from the DataLoggers.

C. Online Monitoring

A fraction of the data taken is also streamed to the online monitoring computers. The online monitoring functions to provide a quality check on the data. It is designed for rapid feedback and processing of an event and should take no more than O(10) s. Higher level reconstruction (which can take minutes) is reserved for offline DQM. A nonexhaustive list of the types of quantities to be measured is TPC noise, photon dark/cosmic count, rate and types of triggers, error counts, hit maps, and so on. The online monitoring has a defined data-flow as shown in Fig. 11 and as follows.

1) *Dispatcher* processes receive events from DataLoggers and route them to the online monitoring processes.
2) *RawDecoder* processes unpack the raw data to perform low-level analysis.
3) *ArtAnalyzer* processes perform high-level analysis of the unpacked events.
4) Histograms are saved and propagated to a web display tool (*Monet*) that is inherited from the LHCb experiment's DQM system [11].

D. Backpressure

The DAQ system is required to be able to sustain the requested rate of 25 Hz during an SPS spill and is capable of making use of the interspill time to absorb any backlog that may have built up during the spill. Backpressure is not expected at the level of the on/near detector readout electronics, since the system has been dimensioned such that for the

TPC (and for the SSPs), data will be read-out continuously (based on self-triggering).

1) Trigger Rate Spikes: Triggering with beam is subject to variations in rate that may cause an oversubscription of DAQ resources. In order to prevent this from happening, artificial dead-time will be managed inside the trigger board logic. For example, simple dead-time handling will be introduced, excluding the firing of the trigger at an interval closer than a predefined fraction of the readout window (5 ms); additional complex dead-time logic in the trigger board may forbid accepting more than a fixed number of triggers during a time interval (e.g., a maximum of 200 triggers during spill, a maximum of 50 triggers during interspill; at 25 Hz over a 4.8-s spill, 125 events are expected on average per spill).

2) Transient Noise Bursts and Event Building/Aggregation Limitations: Transient noise bursts may increase significantly the size of data corresponding to a trigger. In order to prevent data loss, a BUSY software message may be sent to the trigger, when buffer occupancy reaches a certain threshold. When reaching a comfortable buffer occupancy, the BUSY signal shall be cleared. Temporary failure of one or more data flow components (a network link, an EventBuilder node/process, an Aggregator, a pool of disks, etc.) may cause the DAQ performance to be degraded. This will be noticed by the data-flow aggregator as lacking EventBuilder nodes to assign events to. Similar to noise bursts, a BUSY message shall be issued, once queue lengths exceed a defined value.

3) Storage Limitations: The DAQ writes raw data to a temporary storage area, waiting for the data to be transferred to EOS. A long-term outage of the connection to EOS or its performance degradation may cause the storage area to progressively fill up. This type of issue builds up in hours/days. No transient BUSY mechanism can be applied in this case. From the analysis of back-pressure sources, it is apparent that in ProtoDUNE there is no requirement for a fast feedback to the trigger system to stop the generation of triggers. Therefore, a software-based messaging system is well suited for throttling the trigger. Nevertheless, provisions are made at the timing master level to be able to introduce a hardware signal indicating busy conditions, in case it turns out to be needed.

V. Conclusion

This paper describes the ProtoDUNE single-phase DAQ system, which meets the requirements of challenging conditions and described constraints. In order to meet the extremely tight construction and commissioning schedule, most of the implementations are based on COTS hardware solutions and mature software tools and frameworks. As the DAQ needs to be externally triggered, a custom timing and trigger system was developed, which aggregates inputs from beam instrumentation, muon tagger, and photodetectors and assembles a single global trigger for the system. The system provides two TPC readout solutions, in order to identify strong and weak points of each one. The DAQ is designed and tested to handle the effects of backpressure, trigger rate spikes, and transient noise bursts, and the data flow orchestrator provides features to handle event building and aggregation limitations at higher trigger rates. ProtoDUNE-SP started taking data in Q4 of 2018 and it is the largest test beam experiment constructed to date. The DAQ system faced a wide range of challenges and it successfully proved its design principles.

Acknowledgment

The author would like to thank groups external to DUNE, who contributed to the development of the DAQ system; CERN IT for the support on the DAQ servers and network, the FNAL artdaq Team for providing the software framework and the JCOP Team for the run control and operational monitoring support. The author would also like to thank the ATLAS collaboration for the assistance with the development of the FELIX-based readout and Neutrino Platform for the infrastructural support at EHN1.

References

[1] B. Abi *et al.* (2018). "The single-phase ProtoDUNE technical design report." [Online]. Available: https://arxiv.org/abs/1706.07081

[2] R. Acciarri *et al.* (2017). "Noise characterization filtering MicroBooNE liquid argon TPC." [Online]. Available: https://arxiv.org/abs/1705.07341

[3] R. Herbst *et al.*, "Design of the SLAC RCE platform: A general purpose ATCA based data acquisition system," in *Proc. IEEE Nucl. Sci. Symp. Med. Imag. Conf.*, Nov. 2014, pp. 1–10. doi: 10.1109/NSS-MIC.2014.7431254.

[4] J. Anderson *et al.*, "FELIX: A PCIe based high-throughput approach for interfacing front-end and trigger electronics in the ATLAS Upgrade framework," *J. Instrum.*, vol. 11, no. 12, 2016, Art. no. C12023.

[5] A. J. Peters, "EOS as the present and future solution for data storage at CERN," *J. Phys. Conf. Ser.]*, vol. 664, Aug. 2015, Art. no. 042042.

[6] Intel.com. (2018). *Intel QuickAssist Technology (Intel QAT)*. [Online]. Available: https://www.intel.com/content/www/us/en/architecture-and-technology/intel-quick-assist-technology-overview.html

[7] K. Biery, C. Green, J. Kowalkowski, M. Paterno, and R. Rechenmacher, "artdaq: An event-building, filtering, and processing framework," *IEEE Trans. Nucl. Sci.*, vol. 60, no. 5, pp. 3764–3771, Oct. 2013.

[8] C. Green, J. Kowalkowski, M. Paterno, M. Fischler, L. Garren, and Q. Lu, "The art framework," *J. Phys. Conf. Ser.*, vol. 396, Aug. 2012, Art. no. 022020.

[9] P. Deutsch, "DEFLATE compressed data format specification version 1.3," *Proc. RFC*, May 1996, p. 1951. [Online]. Available: https://tools.ietf.org/html/rfc1951

[10] Jcop.web.cern.ch. (2018). *JCOP Framework*. [Online]. Available: http://jcop.web.cern.ch/jcop-framework

[11] M. Adinolfi *et al.*, "LHCb data quality monitoring," *J. Phys. Conf. Ser.*, vol. 898, May 2017, Art. no. 092027.