

Streaming readout for next generation electron scattering experiments

Fabrizio Ameli¹, Marco Battaglieri^{2,3}, Vladimir V. Berdnikov⁴, Mariangela Bondi³, Sergey Boyarinov², Nathan Brei², Andrea Celentano³, Laura Cappelli⁵, Tommaso Chiarusi⁶, Raffaella De Vita³, Cristiano Fanelli^{7,8}, Vardan Gyurjyan², David Lawrence², Patrick Moran⁷, Paolo Musico³, Carmelo Pellegrino⁵, Alessandro Pilloni^{9,10}, Ben Raydo², Carl Timmer², Maurizio Ungaro², Simone Vallarino¹¹

¹ Istituto Nazionale di Fisica Nucleare, Sezione di Roma, 00185 Roma, Italy

² Thomas Jefferson National Accelerator Facility, Newport News, Virginia 23606, USA

³ Istituto Nazionale di Fisica Nucleare, Sezione di Genova, 16146 Genova, Italy

⁴ Catholic University of America, Washington, DC 20064, USA

⁵ Istituto Nazionale di Fisica Nucleare, CNAF, 40127 Bologna, Italy

⁶ Istituto Nazionale di Fisica Nucleare, Sezione di Bologna, 40127 Bologna, Italy

⁷ Massachusetts Institute of Technology, Cambridge, Massachusetts 02139-4307, USA

⁸ The NSF AI Institute for Artificial Intelligence and Fundamental Interactions, Massachusetts 02139-4307, USA

⁹ Università degli Studi di Messina, Dipartimento di Scienze Matematiche e Informatiche, Scienze Fisiche e Scienze della Terra, 98166 Messina, Italy

¹⁰ Istituto Nazionale di Fisica Nucleare, Sezione di Catania, 95123 Catania, Italy

¹¹ Istituto Nazionale di Fisica Nucleare, Sezione di Ferrara, 44122 Ferrara, Italy

Draft : October 9, 2022

Abstract Current and future experiments at the high intensity frontier are expected to produce an enormous amount of data that needs to be collected and stored for offline analysis. Thanks to the continuous progress in computing and networking technology, it is now possible to replace the standard ‘triggered’ data acquisition systems with a new, simplified and outperforming scheme. ‘Streaming readout’ (SRO) DAQ aims to replace the hardware-based trigger with a much more powerful and flexible software-based one, that considers the whole detector information for efficient real-time data tagging and selection. Considering the crucial role of DAQ in an experiment, validation with on-field tests is required to demonstrate SRO performance. In this paper we report results of the on-beam validation of the Jefferson Lab SRO framework. We exposed different detectors (PbWO-based electromagnetic calorimeters and a plastic scintillator hodoscope) to the Hall-D electron-positron secondary beam and to the Hall-B production electron beam, with increasingly complex experimental conditions. By comparing the data collected with the SRO system against the traditional DAQ, we demonstrate that the SRO performs as expected. Furthermore, we provide evidence of its superiority in implementing sophisticated AI-supported algorithms for real-time data analysis and reconstruction.

Keywords Streaming readout data acquisition, on-beam validation, PbWO calorimetry

1 Introduction

A new generation of electron scattering experiments is underway at the world-leading QCD facilities such as Brookhaven National Lab (BNL) and Jefferson Lab (JLab). New projects include the Electron Ion Collider (EIC) [1] at BNL, SOLID [2] and Moller [3] at JLab, and upgrades of the existing detectors in the two labs, sPHENIX [4] and CLAS12 [5], respectively.

All these experiments are characterized by modern detectors with millions of active readout channels and by an unprecedented data rate produced by high-luminosity operations of the accelerators. The ambitious scientific program at the *intensity frontier* of nuclear physics calls for a data acquisition system (DAQ) that can record the interesting events and filter out the unnecessary background. Advances in data manipulation algorithms, e.g. artificial intelligence (AI) and machine learning, open up new possibilities for (quasi-)real-time data processing, by providing an efficient tool to calibrate the detector while running and at the same time intelligently select and reconstruct the final state particles. To fully exploit this progress, it is necessary to

leave the triggered DAQ paradigm and move towards a more flexible software-based framework. In this scheme, all data is streamed from the detector to a data center where the entire detector's information can be analyzed and used for efficient data tagging and filtering. This framework is called *triggerless* or *streaming read-out (SRO)* DAQ.

2 Streaming Readout DAQ

Due to the improved performance of CPUs and computer networks, the FPGA-based trigger scheme can be partially replaced by SRO DAQ. By removing the hardware trigger and performing the full online data reconstruction, it provides precise selections of (complicated) final states for further high level physics analysis. For example, see the current effort at CERN [6]: in preparation for the high luminosity run, expected in 2029, both ALICE and LHCb experiments are planning to take data in streaming mode starting from 2022 RUN3 [7, 8] and test some heterogeneous hardware platforms, such as GPUs, for online data processing during the run [9]. In a triggerless data acquisition scheme, each channel that exceeds a threshold (as implemented on the front-end board) is labeled with a time-stamp and then transferred, regardless of the status of the other channels. A powerful online CPU farm, connected by a fast network link (usually optical fiber) to the front-end electronics, receives all data samples, reorganizes the information ordering hits by time, includes calibration constants, and finally applies algorithms to find specific correlations between reconstructed hits (a.k.a. a software trigger), keeping and storing only filtered events. The advantages of this scheme come from the following: using fully reconstructed hit data to define a high-level event selection condition, implementing online algorithms in a high-level programming language, and gaining the ability to easily upgrade the system configuration and accommodate new requirements. Furthermore, the system can be scaled to match different experimental conditions (either unexpected or foreseen in a planned upgrade) by simply adding more computing (CPUs) and/or data transfer (network switches) resources. FPGAs (or similar fast front-end data concentrators) will still be used in a SRO DAQ system, not to make decisions concerning which events to select, but to handle “low-level” tasks such as adding timestamps or canalizing the data. The pipeline can then be extended to include calibrating and tagging events in the stream. The following will review some examples of how SRO can benefit current and future experiments at electron machines.

2.1 Real Time Processing: Calibration, Selection, and Tagging

A significant benefit of SRO systems is the ability to access data from all detectors when making decisions on what must be kept and what can be safely discarded in an event. A natural extension of this concept is to analyze the data as much as is reasonable during acquisition in order to aid and expedite the physics analysis downstream. Two techniques are particularly useful for this. One is the generation of detector calibrations in (quasi-)real time. The other is event tagging. It is with these that SRO starts blurring the line between data acquisition and offline analysis.

SRO systems must be designed to keep up with the peak detector rates that are determined by the accelerator's full operating luminosity. In actual experiments, accelerators do not operate continuously at full luminosity. Beam trips, fill depletion, configuration changes, etc. all contribute to downtime. Writing the data stream (after triggering and filtering) to a large, temporary disk buffer will smooth out these fluctuations on a time scale corresponding to the size of the buffer. Thus, there is a minimal latency between when data is streamed into the buffer and data is streamed out. This latency can be utilized to obtain detector calibrations from the streamed data before it is sent to storage. Furthermore, with the calibrations available, full or partial reconstruction of the data may be performed prior to storage, further reducing the resource requirements downstream while allowing events to be *tagged* as being of interest to certain physics reactions. Figure 1 illustrates this concept. The reconstruction/tagging stage shown in figure can be implemented on a large HTC or HPC facility, effectively making it part of the SRO data stream. By way of a quantitative example: existing triggered systems at JLab produce experimental data at peak rates of 1–10 GB/s with approximately a 60% efficiency of the accelerator. A buffer that gives a 72hr latency would need to be $(10 \text{ GB/s}) \times 0.60 \times (72 \text{ hr}) = 1.6 \text{ PB}$. The calibration and reconstruction/tagging processes can be done asynchronously with each other while conceptually forming a single, chained stream with multiple data reduction stages.

2.2 SRO examples

Here we consider two specific examples, the new EIC and an upgraded CLAS12 detector at JLab, to discuss benefits of an SRO scheme to the physics program of the two experiments.

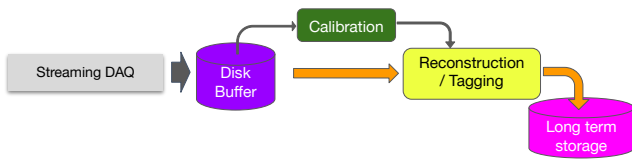


Fig. 1 Streaming DAQ systems must keep up with the peak rate of the accelerator. However, a large buffer can smooth out fluctuations and downtimes so downstream consumers of the stream need only keep up with the average rate. This gives the added benefit of a time delay that can be used to generate calibrations. Conceptually, this extends the stream all the way into the offline reconstruction.

2.2.1 SRO for JLab CLAS12-HI-LUMI

The CEBAF Large Acceptance Spectrometer for operations at 12 GeV beam energy (CLAS12) is used to study electro-induced nuclear and hadronic reactions. This spectrometer provides efficient detection of charged and neutral particles over a large fraction of the full solid angle. Cherenkov counters, time-of-flight scintillators, and electromagnetic calorimeters provide identification of the scattered electron and of produced hadrons. Fast triggering and high data-acquisition rates (event rate up to 30 kHz, data rate up to 800 MB/s) allow operation at a luminosity of $10^{35} \text{ cm}^{-2}\text{s}^{-1}$. These capabilities are being used in a broad program to study the structure and interactions of nucleons, nuclei, and mesons, using polarized and unpolarized electron beam and targets, for beam energies up to 11 GeV. The detector is undergoing to an upgrade aiming to increase the operational luminosity by a factor of two. In particular the tracking system, currently based on Drift Chambers, will be replaced by a faster tracker with higher granularity such as GEMs, MicroMegas or the novel μ Rwell detectors. The current trigger-based DAQ performance can be increased, reaching an event rate close to 100 kHz. For further performance improvement, a full streaming DAQ is under consideration as part of the detector upgrade. The new scheme should overcome other present limitations, such as triggering on neutrals (i.e. photons from π^0 decay, by imposing a sharp cut on the invariant mass), or triggering on kaons (after performing a crude particle identification). Improvement is also expected in trigger purity, currently constrained by a significant hadronic background (mainly hadronic showers initiated by pions). More sophisticated algorithms, that make use of the shower distribution and the combined information from the calorimeter and the Cherenkov detector, will result in a cleaner selection of the desired reaction channels. The data rate from the CLAS12 detector in DAQ streaming mode is estimated to be on the level of 50 GB/s, and a data reduction factor 10 or higher is required to decrease the recorded

event rate to a level that is compatible with current data storage technology.

2.2.2 SRO for BNL Electron Ion Collider (EIC)

The detectors planned for the future EIC at BNL will be among the few major collider detectors to be built from scratch in the 21st century. A truly modern EIC detector design must be complemented with an integrated readout scheme that supports the scientific opportunities of the machine, improves time-to-analysis, and maximizes the scientific output. A fully SRO design delivers on these promises. In particular, the EIC is expected to measure different reactions with an electron (at least) in the electromagnetic calorimeter. On-line calibration to compensate for e.m. shower energy leakage and gain variation, implementation of sophisticated AI-supported clustering algorithms for better reconstruction of nearby tracks, and improvement of e.m.-hadron shower discrimination will result in a better resolution, better electron/ π^0 discrimination, and higher hadron background rejection. Moreover, the flexibility of the SRO will allow setting dedicated ‘triggers’ to pin down rare processes, e.g. exclusive kaon-rich reactions, among the leading electromagnetic production. The EIC Yellow Report [1] states that SRO is the chosen option for the collider.

3 SRO validation and on-beam tests

Despite the conceptual simplicity of a triggerless DAQ, a realistic implementation with the specific detector readout is necessary to demonstrate the expected performance. The sophisticated combination of suitable front-end electronics, network facilities, and CPU-based algorithms requires a significant effort to identify or develop (if not available yet) the best option for each element, set up and test the whole scheme, and compare the results with traditional approaches.

In the following, we will describe in detail the components of Jefferson Lab’s SRO DAQ prototype (both front-end and back-end) and the results obtained from on-beam tests during an opportunistic measurement campaign performed at the lab in 2020. This prototype is designed to serve as a template for future SRO systems.

3.1 Detector setup

On-beam tests were performed with two setups of increasing complexity. In the first experiment, a 3×3 PbWO₄ matrix was placed downstream of a secondary

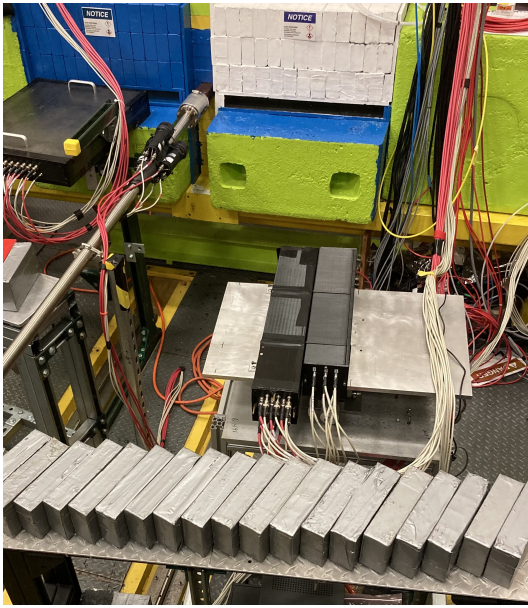


Fig. 2 The PbWO₄ EIC EM calorimeter prototype (on the left) installed in the Hall-D downstream of the Pair Spectrometer (PS).

electron/positron beam generated by the primary photon beam in Hall-D. Leptons were identified and their energy measured by a Pair Spectrometer tagger (PS) [10, 11]. This enabled a precise determination of the energy and a quantitative comparison between SRO and triggered clustering algorithms. In the second experiment, we measured the inclusive π^0 electroproduction with the Hall-B CLAS12-FT detector. Detailed Geant4 simulations of the FT provided a realistic estimate of the detector acceptance and efficiency. The π^0 mass was used as a reference to identify the channel and provide a quantitative comparison to the expected rates.

3.1.1 Hall-D

On-beam tests used an EIC calorimeter prototype consisting of a 3×3 matrix of PbWO₄ crystals. Each crystal, $2.05 \text{ cm} \times 2.05 \text{ cm} \times 20 \text{ cm}$ in size (corresponding to $\sim 22X_0$) was read out by a 19 mm diameter photomultiplier tube (R4125) with a custom HV base and active divider.

The prototype was installed in Hall-D downstream of the pair spectrometer (PS), as shown in Figure 2. Electron-positron pairs are produced by the primary photon beam interacting with 750 μm beryllium converter. Lepton pairs are deflected in a 1.5 T dipole magnet and detected using two layers of scintillation counters positioned symmetrically around the photon beam line. Each arm consists of 8 coarse counters and 145 high-granularity counters. The high-granularity hodoscope is used to measure the lepton momentum; the position

of each counter corresponds to the specific energy. Each detector arm covers the lepton momentum range 3–6.2 GeV/c. The energy resolution of the PS is estimated to be better than 0.6%. The position of the prototype was surveyed and aligned with respect to the beam line and the center of the pair spectrometer magnet, such that the lepton beam’s spot is focused on the center row of the prototype, perpendicular to the front face of the crystals.

3.1.2 Hall-B

The JLab SRO DAQ system, which is also expected to be used in upcoming CLAS12 high luminosity operations, was validated on one of the two electromagnetic calorimeters present in CLAS12. The EM calorimeter provides the trigger for most, if not all, of the processes of interest in electron machines by tagging the scattered electron. In CLAS12 there are two different calorimeters: the FT-Cal and the FD-Cal. The FT-Cal covers polar angles (with respect to the beam axis) $2.5^\circ < \theta < 5.0^\circ$, and azimuthal angles $0^\circ < \phi < 360^\circ$. The FT-Cal, located 1.85 m downstream the production target, is made of PbWO₄ crystals read by APDs for a total of 332 channels. The FD-Cal, a sampling calorimeter made of lead and plastic scintillator bars and read by PMTs, covers the region $5.0^\circ < \theta < 35.0^\circ$. It is split into six identical sectors, each covering $\sim 60^\circ$ in ϕ , for a total of ~ 3000 channels.

The limited number of channels, the compactness, and the full coverage of the azimuthal angle range together make the FT-Cal an ideal detector for testing the SRO DAQ. Moreover, the FT-Cal is part of a larger system (the CLAS12 Forward Tagger or FT) consisting of a scintillator-tile hodoscope (FT-Hodo) and a MicroMega tracker (FT-TRK) located upstream of the calorimeter. The FT-Hodo, composed of 232 plastic scintillator tails, is used to distinguish neutrals from charged particles and, in turn, identifies photons. The FT-TRK, consisting of two layers of MicroMega detectors, allows a precise determination of the electron coordinates. Figure 3 shows the detector with the different components. The whole FT is a small-scale detector that provides charged/neutral particle detection and identification, representing a simple but complete template of the entire CLAS12. For the aforementioned reasons we focused our tests on the FT-Cal. A detailed description of the detector, including on-beam performance, is reported in Ref. [12].

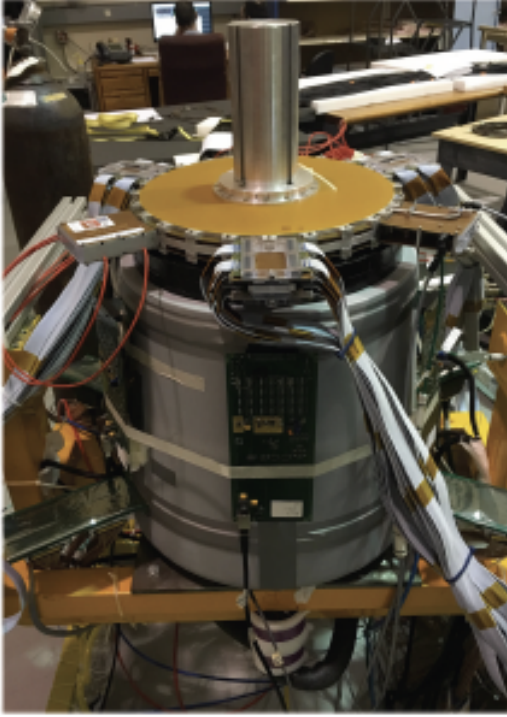


Fig. 3 The CLAS12 Forward Tagger (FT) with details of the PbWO calorimeter (FT-Cal), the plastic scintillator tile hodoscope (FT-Hodo) and the MicroMega tracker (FT-TRK).

4 JLab SRO DAQ

On-beam tests were performed using different combinations of front-end electronic boards (INFN-WaveBoard and JLab fADC250), the CEBAF Online Data Acquisition framework (CODA), the Triggerless Data Acquisition System (TriDAS) (for the back-end) and the high-level analysis framework JANA2.

4.1 Front-end electronics

4.1.1 The INFN WaveBoard digitizer

The *WaveBoard* is a 12-channel digitizer designed for High Energy Physics experiments. Its main features are an extremely performant sampling architecture at a low cost per channel (compared to equivalent commercial boards), a versatile front-end which can be interfaced to different sensors, a flexible timing system with high resolution, a memory buffer large enough for temporary data storage, and self-triggering algorithms based on waveform analysis. Fig. 4 displays the multi-channel digitizer board.

The firmware of the digitizer board supports an auto-triggering mechanism particularly suited to the streaming readout environment: digitized waveforms,

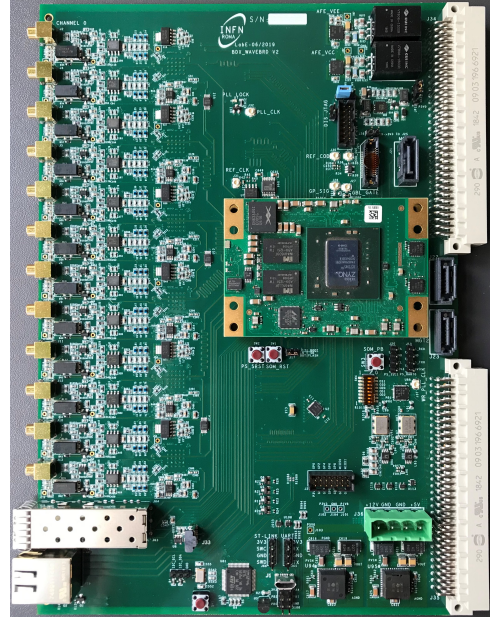


Fig. 4 The multi-channel digitizer board.

whose amplitude is greater than a programmable threshold, are stored in a fast FIFO, timestamped, and forwarded to the DAQ via a GbE link for the physics event search. A triggered system can be accommodated as well through dedicated firmware.

Data digitization is performed by 6 dual true differential ADCs from Texas Instruments. The ADC family members are pin to pin compatible, offering resolutions from 12 to 14 bits and sampling frequencies from 65 to 250 MHz.

The board is able to power the input sensors individually with a High Voltage (HV) of up to 100 V, provided by a dedicated HV linear regulator.

To align the system to an external time reference, a clock and a timing signal can be received by the *WaveBoard* on two dedicated U.FL connectors. The timing signal is fed directly to the FPGA and can be either a slow reference (e.g. a Pulse Per Second signal) or a digital timing protocol (e.g. NMEA, IRIG, etc.). This signal imposes the same phase to timing on different boards.

Data collection and manipulation is accomplished by a commercial System On Module (SoM) mezzanine board from Trenz Electronics, based on a Xilinx Zynq-7000 SoC. The SoM is hosted by the board via proper connectors. The Zynq-ARM processor runs a Linux distribution, making it substantially simpler to interface with the board. Physics data, stored in DDR memory by the auto-triggering algorithms implemented in the programmable logic, are transmitted over the GbE link using a TCP/IP protocol. Thanks to a Direct Memory

Address (DMA) architecture, the available data rate can be as high as 900 Mbit/s.

4.1.2 The JLab fADC250 digitizer and VTP board

The FADC250 (Fig. 5) is a similar digitizer (250 MHz sampling rate, 16 channels) designed at JLab and used in many experiments for general purpose triggered readout of detectors. The firmware has been adapted for streaming readout by utilizing the VXS serial links, which were previously used for trigger outputs. The FADC250 supports a total output bandwidth of up to 20 Gbit/s, though we currently utilize 10 Gbit/s. The FADC250 firmware detects pulses using a software-defined threshold. When a threshold crossing is found, the pulses are integrated in a programmably-sized window, a pedestal is subtracted, and a gain is applied. The result is a calibrated charge and time for the found pulse, which is sent over the VXS interface to the next stage (VTP). When running in this mode, the FADC250 supports up to 30 MHz of hits per channel.

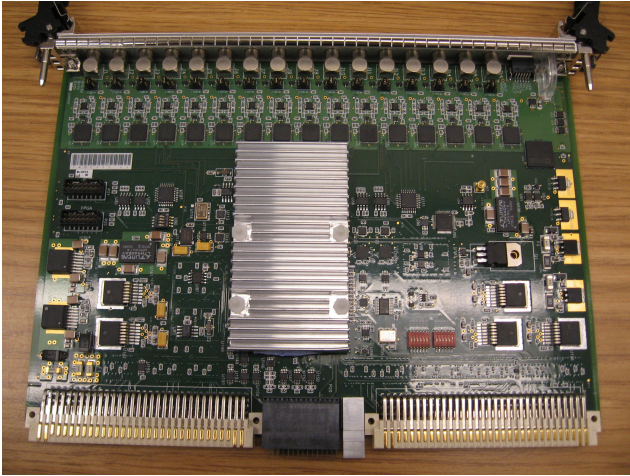


Fig. 5 JLab FADC250 Board. A VME/VXS style, 16 channels, 250 MHz 12bit ADCs. Readout can be over VME (200 MB/s) or through VXS serial links (up to 20 Gbit/s).

The VTP board (Fig. 6) is another custom JLab design that is used in conjunction with FADC250 modules when fast readout is needed. This is a fairly generic and flexible module (with several firmware implementations to match the requirements of different experiments) whose conversion to streaming readout was a simple and natural extension. The resources on the VTP provide reasonable serial connectivity between the VXS and the optical links, as well as significant buffering capability. See Fig. 7 for the board schematic.

The VTP collects streaming data from up to 16 FADC250 modules, storing the streams in DDR3 mem-

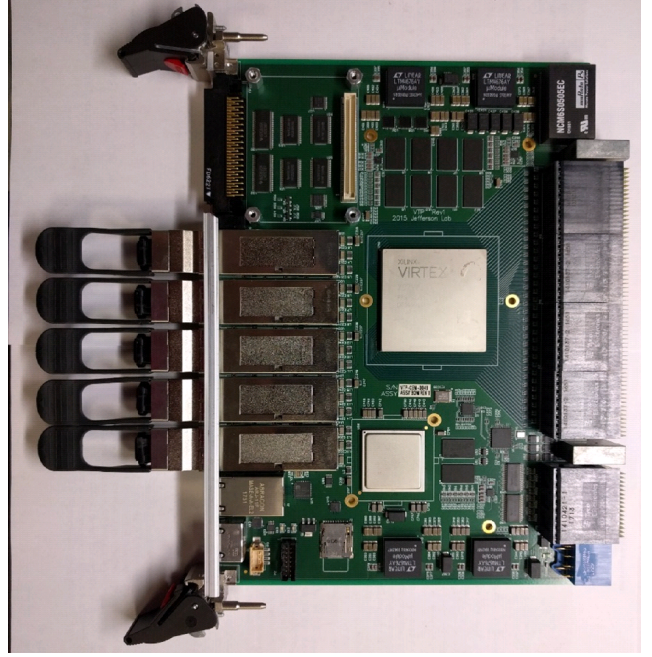


Fig. 6 VXS Trigger Processor (VTP) Board. This connects to all VXS front-end cards in a crate to perform readout and/or trigger functions. It provides multiple optical interfaces for sharing information between other VTP modules and for streaming data to the network.

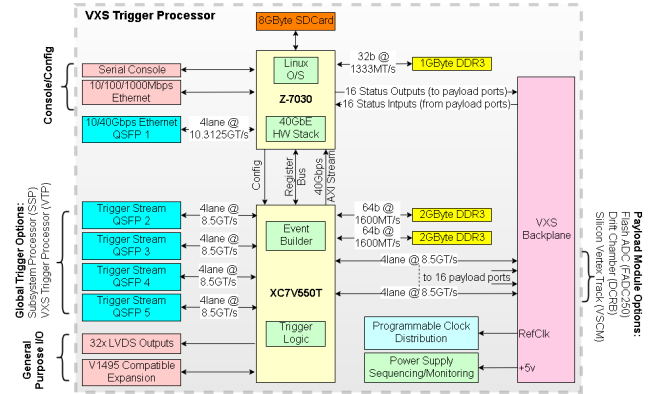


Fig. 7 VXS Trigger Processor (VTP) Diagram. Streaming data is collected over VXS from up to 16 payload front-end cards using up to 4 serial lanes at up to 8.5 Gbit/s per lane. This data is then buffered in DDR3 memory (connected to the XC7V550T) while the XC7V550T streams it over Ethernet using up to 4 fully hardware accelerated 10 Gbit/s TCP/IP Ethernet interfaces.

ory. FADC250 hits are packaged into 65 μ s frames and sent over 10Gbps Ethernet using the TCP protocol. Each frame contains data from up to 8 FADC250 modules. Currently, up to two 10 Gbit/s Ethernet links can run together to support up to 16 FADC250 modules. The 10 Gbit/s TCP link can sustain up to 8 Gbit/s of TCP data without creating significant backpressure. When the data rate is too high and creates TCP backpressure, the streaming data is allowed to back up into

the DDR3 memory before any is dropped. If this back-pressure is sustained, the VTP will drop data frames. When data is dropped, it will be a complete $65 \mu\text{s}$ frame that is lost; each frame contains a frame counter and timestamp so that the receiving end will know that this happened. To support high data rates without loss, ‘hot’ channels should be distributed over multiple VTPs, and additional 10 Gbit/s Ethernet links should be enabled. Currently, when running with 16 FADC250 modules (256 channels) and two 10 Gbit/s Ethernet links, the rate limit is around 150 MHz for hits (or about 2 MHz for average hits per channel), which is an extremely large value.

In the future we plan to expand this setup to provide more FADC250 data types and to enable other JLab VXS modules to use the same streaming readout concept, thereby providing a low-cost upgrade opportunity for future JLab experiments.

4.2 Run Control: CODA

The CEBAF Online Data Acquisition system (CODA) was designed to work with trigger-based readout systems. The key component is the Event Builder: it collects data from 100+ Readout Controllers (ROCs) and VXS Trigger Boards (VTPs), and builds events based on event number. Another important component is the Trigger Supervisor (TS), which synchronizes all components using clock, sync, trigger and busy signals. ROCs are reading front-end electronics over a VME bus, and VTPs are forming trigger decisions and reporting some trigger-related information. A detailed description of CODA can be found in [13].

To use the available front-end electronics in streaming mode, the role of the TS was reduced to clock distribution, and the Event Builder was replaced with a new SRO component and back-end software capable of gluing ROC information based on timestamp instead of event number. ROCs are not sending any data in that mode, rather they merely do initial configuration settings over the VME bus. All front-end electronics readout is performed by VTP boards over serial lines rather than the VME bus, which increases the bandwidth limit from about 2 to 20 Gbit/s for each of three participating electronics crates, with the possibility for a further increase to 40 Gbps if needed. New firmware was developed for VTPs to implement streaming mode.

Finally, an intermediate software layer, called *SRO Translator*, was introduced between the VTPs and TriDAS. It runs in online mode on its own server. It rearranges and reformats the data streaming from the VTPs so that the data can be ingested by TriDAS. The

SRO Translator and TriDAS communicate via TCP sockets over the network.

4.3 Triggerless Data Acquisition System (TriDAS)

The Triggerless Data Acquisition System (TriDAS) [14, 15] is a streaming readout software framework originally designed and implemented for the astrophysical neutrino detector prototype NEMO [16]. NEMO is a very large detector, distributed across a cubic kilometer beneath the Mediterranean Sea, and was deployed in stages over several years. A key requirement of TriDAS was to support data taking from the beginning, and to scale with the detector installation.

To fulfill this ambitious requirement, TriDAS leverages a modular design, which turned out to meet the needs of a beam-based experiment with minimal development effort. TriDAS is made of various software components implemented in C++11, each dedicated to a specific task in the data-processing chain. The HitMan-

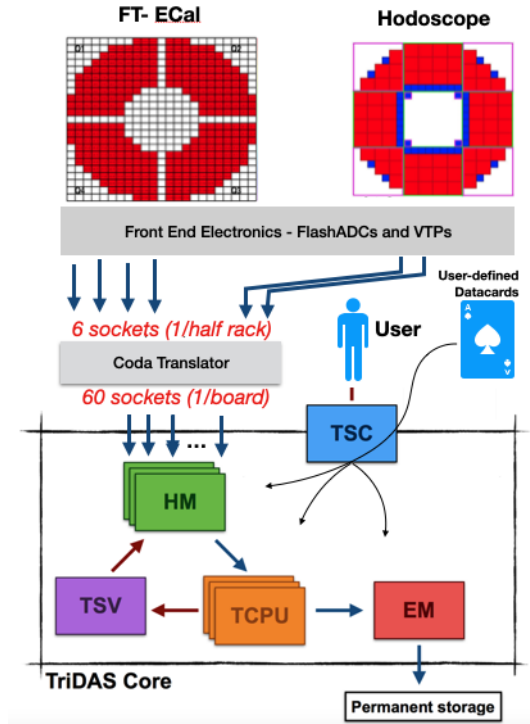


Fig. 8 The DAQ model, from the used CLAS12 subdetectors and front-end electronics to the TriDAS elements.

agers (HMs) represent the first data aggregation stage. They receive data streams from a predefined number of CODA translators, topologically corresponding to a sector of the detector. Each HM divides the collected data into a sequence of time-ordered bunches of data, called Sector Time Slices (STS), with an adjustable

width set to 50 ms. Sharing a common time reference, all HMs arrange their STSs according to the same intervals of time, which are referred to as Time Slices (TS).

The TriggerCPUs (TCPUs) receive the STSes assembled by all HMs corresponding to the same TS and apply the event building and the classification/selection algorithms to the data. Time slices are processed in parallel by multiple threads of the same TCPU and by multiple TCPU processes running on a CPU farm. The Level 1 events (L1) consist of all data within a programmable time window around a hit whose energy exceeded a certain threshold. Considering the typical sizes of the coincidence time window ($O(100 \text{ ns})$) and the TST ($O(10 \text{ ms})$), events spanning two adjacent STSes are neglected. Level 1 events identified within a TS are then fed to the L2 classification/selection algorithms that are implemented in separate binaries and specified in the run configuration file.

A token-based mechanism is at the core of the TriDAS SuperVisor (TSV) load balancing. Each TCPU thread owns a token that is given to the TSV on completion of the TS processing. The TSV maintains a pool of “free to use” TCPU threads which are then matched to the new time slices that are continuously assembled by the HMs.

The Event Manager (EM) collects the selected L2 events and then writes them to the so called Post Trigger (PT) file.

The user interacts with the system via the TriDAS System Controller (TSC). For the CLAS12 tests, a simple interface was custom-built around the TSC.

4.4 High level data processing (JANA2)

One requirement of SRO systems is to be able to filter the data stream so that only the small portion relevant to the physics being measured is ultimately stored. The JANA2 reconstruction framework was integrated with TriDAS for this purpose. The solution took the form of a TriDAS plugin which utilizes JANA2. That TriDAS plugin loads multiple JANA2 plugins, each of which provides a different trigger while sharing any underlying algorithms (e.g. calorimeter clustering). Each plugin reports its own *TriggerDecision* back to TriDAS. The *TriggerDecision* contains a 16 bit tag to identify which plugin produced it and a 16 bit value to record details of the trigger decision. Each plugin is free to assign its own meaning to the 16 bit value, with the constraint that nonzero values indicate TriDAS should keep the event. At least one nonzero trigger value is considered sufficient for keeping an event. All *TriggerDecision* objects are stored alongside each recorded event, allowing

later analysis of which triggers were active and their specifics for a given event.

An important design feature of the JANA2 framework is its on-demand algorithm execution. This can significantly reduce the amount of computation required to implement a sophisticated online trigger system. An algorithm in JANA2 (referred to as a “factory” in the framework’s parlance) operates by taking certain data objects as inputs and producing one or more data objects as an output. A user-defined *processor* activates an algorithm by simply requesting the type of data objects that the factory produces. Thus, for a given event, if a certain type of data object is never requested, then the corresponding algorithm that produces it is never run. This allows the user to organize the requests for data objects such that the least expensive algorithms are run first. If a keep/no-keep decision can be made based on their output, no further algorithms need to be run for that event. Figure 9 illustrates a flow chart of this. Having this fine-grained control built-in at the event level helps the end user optimize the system based on the specifics of their data, hardware, and trigger complexity.

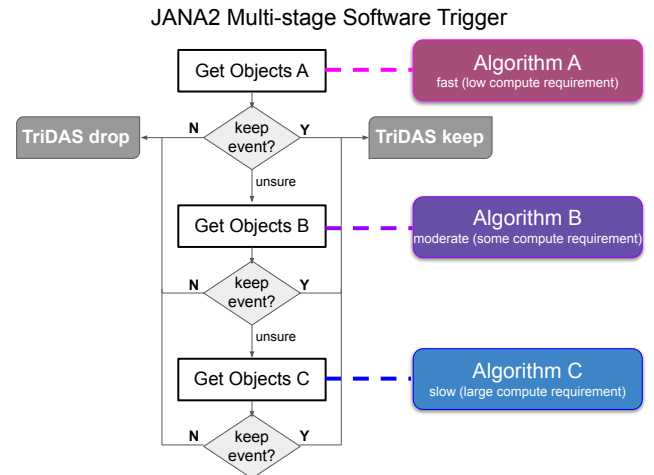


Fig. 9 Example of a configuration supported in JANA2 where event-by-event decisions are made by applying successively more expensive algorithms. As soon as a keep/no-keep decision is made, TriDAS is informed and additional algorithms need not be run, reducing the overall compute requirement for the online trigger(filter).

4.5 AI algorithms for SRO

The adoption of AI in nuclear and particle physics is accelerating and will be an essential part of future experiments such as the EIC [17].

AI encompasses all of the concepts related to the integration of intelligence into machines; unsupervised learning, in particular, is a family of algorithms capable of learning patterns from untagged data, i.e. without a training phase. This has been explored for the first time in near real-time reconstruction of clusters detected in the Forward Tagger calorimeter (FT-Cal) using TriDAS and JANA2.

A clustering algorithm inspired by the *Hierarchical Density-Based Spatial Clustering of Applications with Noise* (HDBSCAN) [18] has been implemented in the form of a plugin within the JANA2 framework. Such an algorithm is characterized by the following features: (i) it can be easily ported to other experiments; (ii) it formally does not depend on cuts during the cluster formation process, making it less sensitive to variations in experimental conditions during data-taking; (iii) it is able to cope with a large number of hits; (iv) it excels when dealing with challenging topologies and arbitrarily shaped clusters, varying cluster sizes and, remarkably, noise (i.e. hits identified as noise are not clustered); (v) it supports calculating the probability that a hit belongs to a cluster or is an outlier.

The last two features are desirable in many experiments and are not provided by simpler algorithms such as k -means [19], which is a semi-supervised approach that has also been implemented in the JANA2 framework during these studies.¹

HDBSCAN permits considering all of the hit-level information in the detector (e.g. spatial, time, and energy) and looking at the density of hits in that space of parameters. To do this, the algorithm leverages a metric called “mutual reachability” distance [18]. This metric combines the density at each point (hit) with the relative distance between two points, and is typically expressed as

$$d_{\text{mreach-}k}(a, b) = \max \{ \text{core}_k(a), \text{core}_k(b), d(a, b) \}, \quad (1)$$

where a and b are two points, $d(a, b)$ is their metric distance, and core_k is the core distance at each point, which depends on a hyperparameter k , and represents the distance to the k th neighboring point. More details can be found in [20]. This step is followed by finding the minimum spanning tree of the points, based off of the connecting edges for each pair of points with weights related to the mutual reachability. The spanning tree can be converted to a hierarchy which is eventually used to extract the clusters (see Fig. 10). In the figure, the parameter λ is the inverse of the introduced distance. An-

¹In its standard implementation, k -means has as hyperparameter the number of iterations to run and the number of clusters k ; in our implementation of k -means, we first determine the seeds of the clusters and then start clustering.

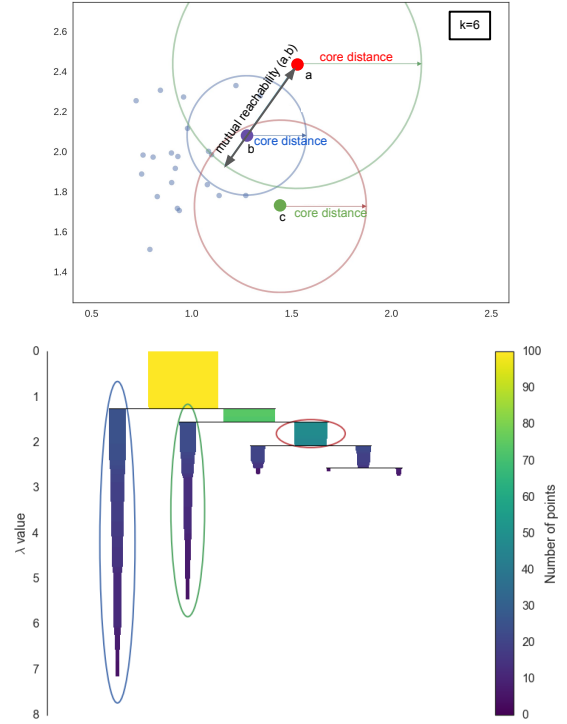


Fig. 10 Illustrative example of mutual reachability (top) for the points a and b , following Eq. (1). HDBSCAN condenses the tree upwards via density based notions, defining a maximum lambda boundary (bottom). The horizontal lines indicate the lambda value boundaries of clusters. Image taken from [20].

other hyperparameter is the minimum cluster size N_{size} , with which we can now walk through the hierarchy and, at each split, ask if one of the new clusters created by the split has fewer points than the minimum cluster size. As shown in Fig. 10, we select clusters from the condensed tree with larger persistence, i.e. with longer lifetime as measured by λ . Recent improvements in the HDBSCAN algorithm can be found in [21].

This approach is very effective in suppressing single tower noise. Hits are associated with clusters along with their probability of membership and their probability of being an outlier. These become new features that can later be used for refining the selection of physics events.

Before the SRO tests began, these AI-based approaches were initially developed and tested with well understood and characterized minimum-bias triggered uncalibrated events collected in the FT-CAL of CLAS12, as shown in Fig. 11.

Diphoton invariant mass spectra obtained with the AI-based approaches have been benchmarked against the spectrum obtained with the standard clustering algorithm to test the tuning of the hyperparameters (Fig. 11 shows only one particular case). For HDBSCAN we also extend the clustering to the entire information avail-

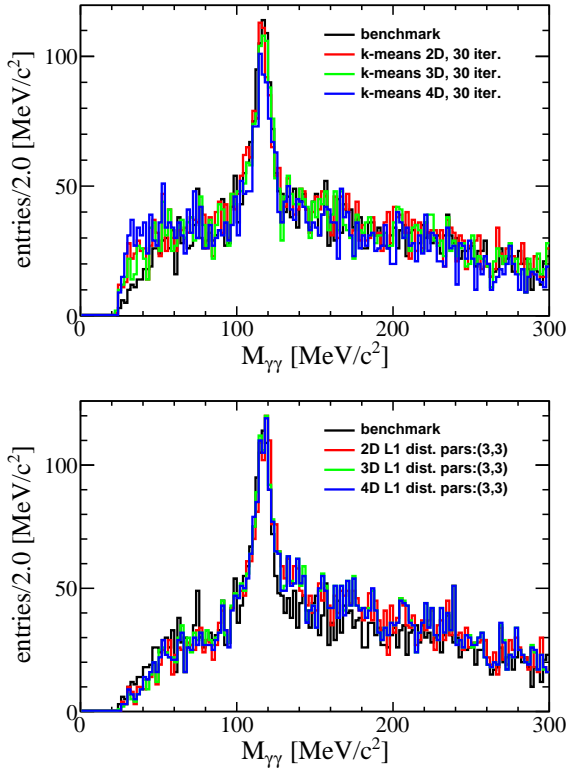


Fig. 11 Diphoton invariant mass spectra. Top: standard clustering algorithm (benchmark) and k -means, shown with different configurations of the hyperparameters. Bottom: the same with HDBSCAN.

able in the calorimeter (4D: x, y, t, E). Loose selection criteria with fiducial cuts is applied consistently in all cases in Fig. 11 to produce the corresponding diphoton invariant mass spectra. With this simple and clean dataset, the π^0 yields obtained with the different methods are comparable, but k -means retained more background at lower mass value. As expected, the runtime of k -means is comparable to the standard algorithm, while HDBSCAN is 30% slower on average due to its more complex calculations. On the other hand, HDBSCAN is a more suitable clustering strategy for more complex data, as it handles high multiplicity, noise, and complex topologies. No cuts on the membership probabilities or outlier scores of the hits have been applied in the HDBSCAN case — this is a promising opportunity that is left for future studies. In Sec. 5.2.3, we will run the AI-based clustering algorithms on data taken during the SRO tests and provide a detailed description of the accomplished analysis.

5 On-beam test results

5.1 Hall-D

Tests were performed parasitically during GlueX high-luminosity runs with a 350 nA photon beam. The prototype was irradiated with a 4.7 GeV secondary electron beam centered with respect to the matrix central crystal. Figure 12 shows a sketch of the experimental setup.

Two different DAQ setups were used: triggered mode (integrated into GlueX data acquisition), and streaming readout. Tests with triggered DAQ were performed by applying the same methodology described in Ref. [22]. The signal amplitude from each PMT was recorded by an FADC whenever a lepton hit a PS hodoscope tile. For SRO tests, each PMT signal was digitized by the *WaveBoard* and streamed to TRIDAS software, where a threshold equivalent to ~ 2 GeV, defined a L1 event.

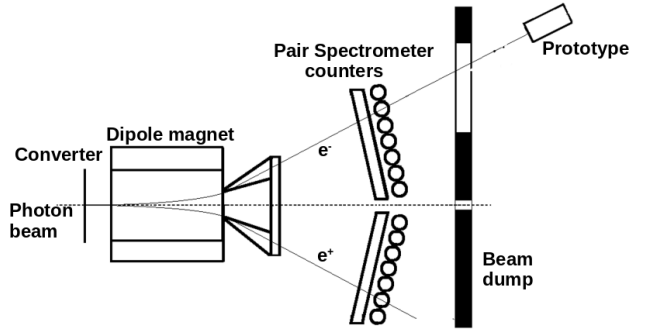


Fig. 12 Schematic of the prototype tests installed in the Hall-D beamline behind the pair spectrometer

5.1.1 Data analysis and results

To validate the performance of the SRO DAQ chain, we compared the energy resolution obtained in triggered and SRO mode. The SRO data analysis was performed within the JANA2 framework, where a dedicated clustering algorithm was implemented. Fig. 13 shows the energy spectrum of the nine channels. The effect of the L1 threshold is clearly visible for the central crystal.

The selection algorithm identified events with a large energy deposited in the central crystals (assumed to be the EM shower seed) and summed all hits in the other channels within a time window of 100 ns. A cut on the energy-weighted x - y hit position was used to exclude events hitting the side crystals after a rough inter-channel energy calibration (the procedure is the same as described later for triggered mode). The clus-

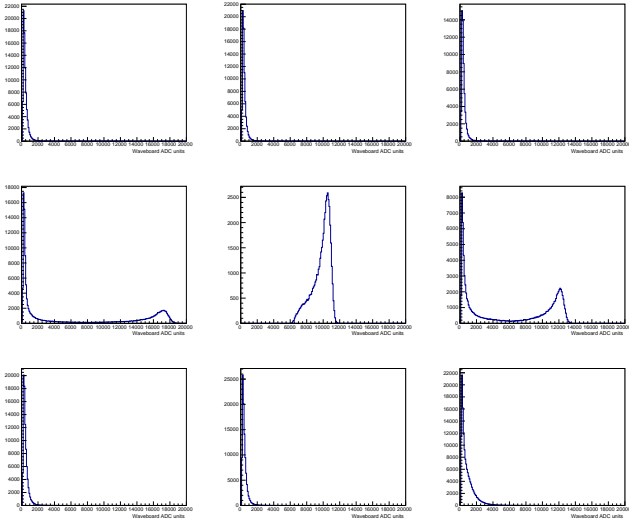


Fig. 13 Response of the nine channels of the calorimeter collected in SRO mode. The beam was centered on the central crystal.

ter energy distribution is shown in Fig. 14. The energy resolution, obtained by fitting the distribution with a Gaussian curve, is $(2.40 \pm 0.05)\%$.

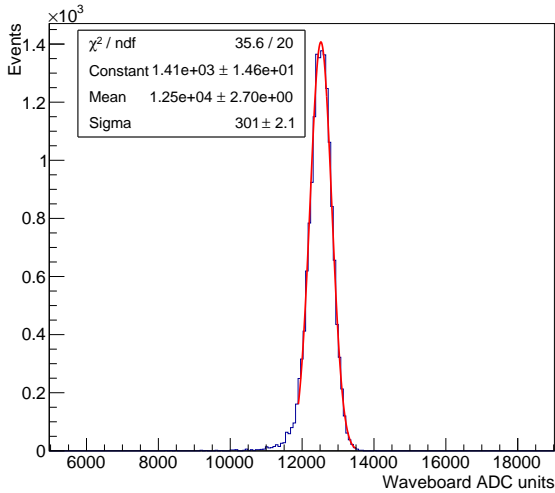


Fig. 14 Reconstructed cluster energy for the SRO configuration.

The energy resolution in triggered mode was obtained after applying the calibration procedure described in [22]. In particular, calibration coefficients were obtained by minimizing the difference between the total energy deposited in the 3×3 matrix and the electron energy measured by the Pair Spectrometer.

The cluster energy spectrum for the triggered configuration,

calculated as the sum of energy deposited in the nine channels, is shown in Fig. 15.

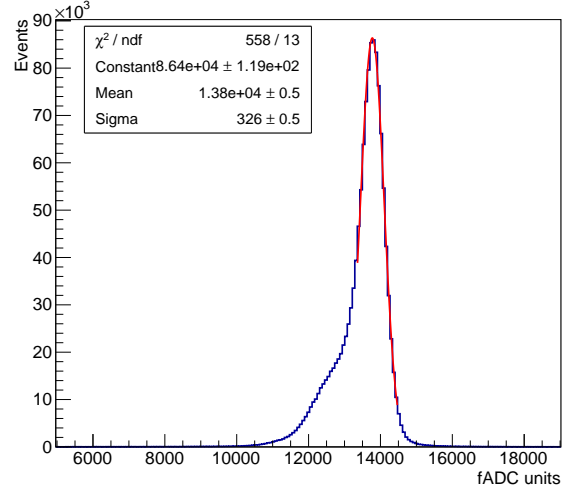


Fig. 15 Reconstructed cluster energy for the triggered configuration.

The resulting energy resolution (Gaussian fit) is $(2.37 \pm 0.05)\%$, in good agreement with results obtained for the SRO DAQ configuration.

5.2 Hall-B

Tests were performed parasitically during two CLAS12 production runs using a 100 nA electron beam accelerated at 10 GeV by the CEBAF on a $125 \mu\text{m}$ lead target for Run-1, and a 40 cm gaseous deuterium target for Run-2.

An SRO system based on FADC250, VTP, TriDAS and JANA2 was used to read out the FT-CAL and FT-HODO.² The TriDAS time coincidence window was set to 200 ns and the TST to 50 ms. The energy threshold of the L1 trigger was set to 2 GeV. Several L2 trigger conditions were used in parallel to tag events: one event over ten, minimum bias, and one or more clusters based on both standard and AI algorithms. Minimum-bias triggered data, accumulated in Run-1, was used to identify a clean physics channel consisting of inclusive π^0 production, and corresponding results were compared with detailed Monte Carlo simulations. In the following sections we report more details about the π^0 yield estimate and data analysis.

²During Run-1, only the FT-CAL was used

5.2.1 Inclusive π^0 electroproduction model

The dominant contribution to the inclusive π^0 cross section is given by the real and virtual photoproduction of a single π^0 . Real photons can be produced by bremsstrahlung inside the lead target. The expected yield of events is given by [23]

$$Y_{\text{real,Pb}} = N_e \frac{N_A X_{\text{Pb0}}}{A} \frac{T_{\text{Pb}}^2}{2} \int_{k_{\min}}^{E_{\text{beam}}} \sigma(k)_{\gamma p \rightarrow \pi^0 p} \frac{dk}{k}, \quad (2)$$

where N_e is the number of electrons scattered, T_{Pb} the thickness of the target in units of radiation length, X_{Pb0} the radiation length times density of lead, and N_A the Avogadro number. For incoherent scattering onto nucleons, $A = 1 \text{ g/mol}$. The cross section is parametrized with the model from [24]. We consider $k_{\min} = 2 \text{ GeV}$. The contribution from virtual photons depends on an unknown form factor, function of the photon virtuality. However, it can be approximated by [23]

$$Y_{\text{virtual,Pb}} \simeq \frac{2t_{eq}}{T_{\text{Pb}}} Y_{\text{real,Pb}} \quad (3)$$

with $t_{eq} \simeq 1.7 \times 10^{-2}$. This contribution is about two times larger than the former.

The yields of pions generated by the interaction of the beam with the two exit aluminum foils can be calculated with the same formulae, replacing $\text{Pb} \rightarrow \text{Al}$. Moreover, the target can radiate additional real photons that interact with the lead target. For $80 \mu\text{m}$ of thickness, this contributes to about

$$Y_{\text{rad,Al}} = N_e \frac{N_A X_{\text{Al0}}}{A} T_{\text{Pb}} T_{\text{Al}} \int_{k_{\min}}^{E_{\text{beam}}} f(k) \sigma(k) dk \quad (4)$$

with

$$f(k) = \frac{4}{3} \left(\frac{1}{k} - \frac{1}{E_{\text{beam}}} \right) + \frac{k}{E_{\text{beam}}^2} \quad (5)$$

with

We provide the cross sections averaged over the photon energy profile, in a fiducial region $2^\circ < \theta < 6^\circ$:

$$\frac{1}{\log E_{\text{beam}}/k_{\min}} \int_{k_{\min}}^{E_{\text{beam}}} \sigma(k)_{\text{fid}} \frac{dk}{k} = 182 \text{ nb} \quad (6a)$$

$$\left[\int_{k_{\min}}^{E_{\text{beam}}} f(k) dk \right]^{-1} \int_{k_{\min}}^{E_b} f(k) \sigma(k)_{\text{fid}} dk = 177 \text{ nb} \quad (6b)$$

Subleading contributions can also be considered. The expected largest one is given by $\gamma p \rightarrow \pi^0 \pi^0 p$. As recently measured in [25], the $2\pi^0$ invariant mass is dominated by the $f_2(1270)$ resonance. Using the model from Ref. [26], we estimate that the yields of events having at least one of the pions in the FT acceptance is one order of magnitude smaller than the single π^0 events, and can be safely neglected.

5.2.2 CLAS12-FT simulations

To provide a realistic estimate of expected π^0 yield, the detector acceptance and reconstruction efficiency was evaluated by detailed Monte Carlo simulations. The geometry, materials and the detector response were simulated using GEMC, the CLAS12 Geant4 Monte Carlo package [27].

The geometry was implemented using a database of Geant4 volumes and imported from the CAD engineering model.

The target volumes implemented in the simulations, shown in Fig. 16(top) are:

- a foam scattering chamber with a $50 \mu\text{m}$ aluminum window;
- a 29 cm long cell containing liquid Helium;
- the lead target, $110 \mu\text{m}$ thick.

The FT consists of three subsystems:

- a tracker (FT-Trk), composed of 4 MicroMega layers;
- a hodoscope (FT-Hodo), with eight sectors, each containing two layers of scintillators;
- a calorimeter (FT-Cal) containing an array of 332 crystals.

The three subsystems are shown in Fig. 16-Bottom. The response of FT-Cal crystals is modeled within the simulation as follows. The time window of the calorimeter is set to 132 ns , and all Geant4 steps within the same paddles and time window are collected in one hit. The digitization routines are called at the end of each event, after the Geant4 navigation has propagated all tracks and GEMC has collected all the steps into hits. The deposited energy is converted first to the charge produced at the end of the electronics chain composed by an avalanche photodiode (APD) and preamplifier, and then to an ADC. The first conversion is based on the measured charge for cosmic rays that deposit a known energy in the crystals, while the second one is based on the FADC conversion factor. A smearing on the final ADC values is added, accounting for the Poisson distribution of photoelectrons produced by the photosensor, the Gaussian noise of the photosensor and of the preamplifier. All parameters—the number of photoelectrons per MeV of energy deposited, and the RMS width of the APD noise and of the preamplifier input noise—have been tuned to the experimental data.

The $\pi^0 \rightarrow \gamma\gamma$ events were generated in a fiducial region $2^\circ < \theta < 6^\circ$ and in a photon energy range $2 < E_\gamma < 10 \text{ GeV}$ at Pb target position accordingly with the theoretical distribution and then were passed to GEMC using the LUND data format which encodes the particle IDs, vertices and momenta for each event.

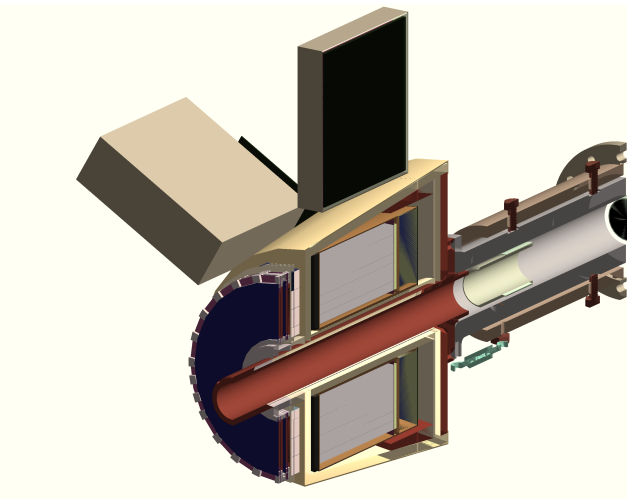
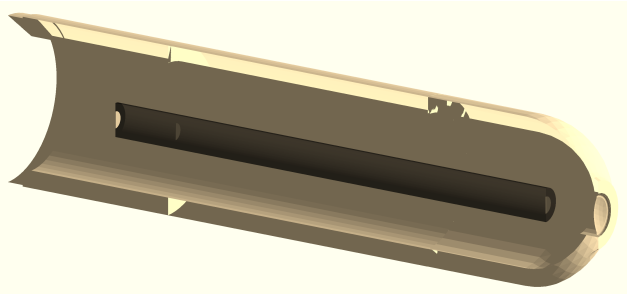


Fig. 16 Top: the lead target and scattering chamber implementation in GEMC. Bottom: details of the implementation of the three Forward Tracker subsystems. As seen by the beam (incident from the left): the disks form the FT-Trk; the FT-Hodo scintillators just behind the tracker; the FT-Cal. The hardware mount support.

The output of the GEMC simulations was analyzed by applying the same clustering algorithms used for experimental data. The simulation shows that the overall efficiency is 1.4%, providing an expected π^0 -event yield of 1800 ± 200 for $N_e = 1.8 \times 10^{14}$.³ Figure 17-Left shows the reconstructed invariant mass peak centered at the expected π^0 mass and 3.6 MeV width, corresponding to an energy resolution of 2.9%.

Pions produced by the interaction of the beam/photons with the two downstream aluminum foils were also simulated, and the corresponding expected yield was found to be 420 ± 100 . The reconstructed π^0 mass distribution, showed in Fig. 17-Right, has a lower mean value, as a result of the assumption that the vertex was located at the lead target position when calculating the invariant mass. The resulting π^0 mass peaks associated to Pb and Al production, are separated by ~ 17 MeV.

³This value corresponds to the accumulated charge during Hall-B tests.

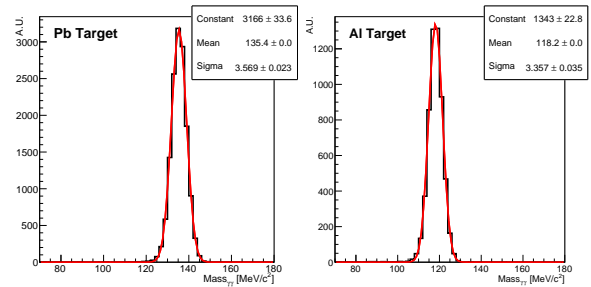


Fig. 17 Left: $\gamma\gamma$ -invariant mass spectrum of π^0 produced by the interaction of the beam with Pb target. Right: $\gamma\gamma$ -invariant mass spectrum of π^0 produced by the Al target. In both cases, the mean value and width of the distribution was determined via Gaussian fit (red line).

5.2.3 Data analysis and results

The offline data reconstruction is performed by applying the same full suite of reconstruction algorithms used in the online analysis, which are implemented in the JANA2 framework and described in Sec. 4.4. The main task is reconstructing the clusters from the raw information associated with the particles detected in the calorimeter. This raw information includes the charge and time for each channel of the detector with a signal above the front-end threshold. The first step in this reconstruction pipeline is to determine the energy and time of the individual crystal hits. The recorded charge of each hit is converted into energy by applying the proper calibration constant, which is determined from a calibration run performed in standard triggered mode. Details on the calibration procedure are reported in [12]. It is worth noting that the calibration run was performed several weeks before the SRO tests, and that significant radiation damage induced on the crystals was not accounted for.

As mentioned in Sec. 4.5, one of the nice features of an unsupervised clustering algorithm like HDBSCAN (compared to standard clustering algorithms) is the ability to cluster hits without using cuts on time, spatial or energy information, but rather looking at their correlations. Therefore, the effect of energy miscalibrations does not impact the HDBSCAN clustering performance, but only affects the reconstructed cluster energy. Also, because AI-based algorithms such as HDBSCAN are capable of performing with large hit multiplicities and in presence of substantial noise, they are particularly desirable in high luminosity experiments that will operate in SRO mode. Fig. 18 shows a comparison of the diphoton invariant mass spectrum obtained with raw data (uncalibrated), utilizing both the standard and the HDBSCAN clustering algorithms. Both spectra are obtained by applying loose fiducial cuts on

the reconstructed events, a minimum cluster size of 3 hits, and a minimum cluster energy (dominated by a threshold on hit energies). Remarkably, the unsuper-

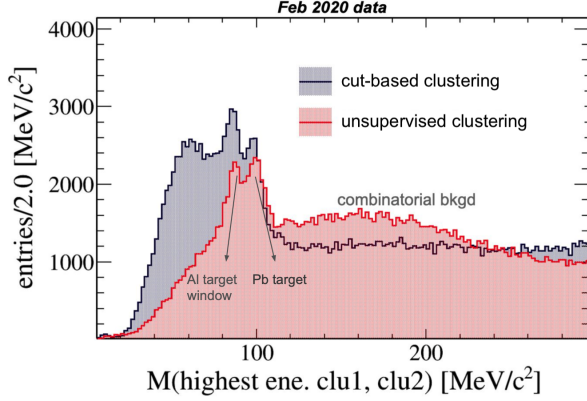


Fig. 18 Diphoton invariant mass spectrum from the two highest energy clusters obtained with the unsupervised clustering (red) as compared to the standard clustering (black) algorithms.

vised clustering provided a π^0 yield consistent with the cut-based standard clustering, but with significantly improved signal-to-background ratio in the region sensitive to the π^0 peak. The standard clustering algorithm utilizes cuts to aggregate hits around the seed of a cluster, and suffers when dealing with data that is not properly calibrated. In the near future we plan to further characterize the performance of this unsupervised approach: preliminary results during our tests have shown a longer runtime of $\sim 30\%$ relative to the standard clustering algorithm when run in a single thread. This clustering technique seems a promising alternative to traditional cut-based approaches, particularly when deployed online on data that is not fully calibrated, and is capable of rejecting noise hits and dealing with complicated topologies of clusters and higher multiplicities. This comes at the cost of introducing the 2 hyperparameters described in Sec. 4.5, which can be tuned with dedicated studies. These results show that the AI-based algorithms are robust and provide an efficient event selection on raw data where miscalibrations are expected. In the following, we describe the data analysis performed using the standard clustering algorithm in more detail. After correcting for the *time walk* effect, the reconstructed hits were sorted by energy. A cluster was identified starting from the seed and adding adjacent crystals with a signal above 50 MeV and within a time window of 10 ns. Once a cluster was identified, the energy, time and position were computed according to the procedure described in Ref. [28]. Fig. 19 shows the distribution of the number of clusters per event: the largest

bucket corresponds to one cluster, and only about 8% of the events have two or more. The analysis aimed

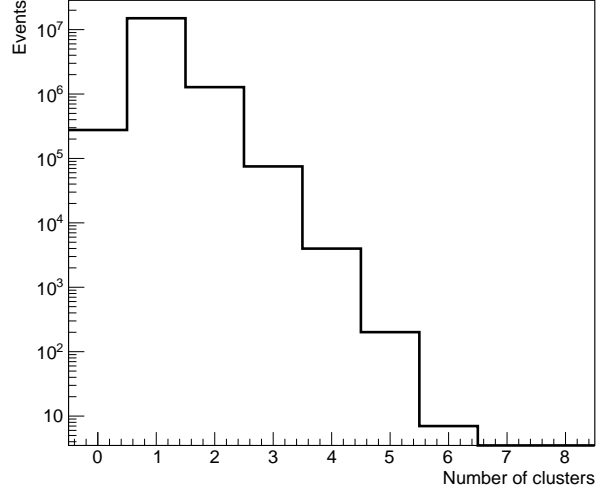


Fig. 19 Distribution of number of clusters per event.

at identifying the $\pi^0 \rightarrow \gamma\gamma$ decays only considers two-cluster events. It uses following selection cuts:

- two clusters appear within a time window of 10 ns from each other;
- both clusters have energy above 500 MeV;
- the sum of the two clusters' energy is less than 8 GeV;
- the number of crystals involved in each cluster is greater than 3;
- the opening angle between the two clusters is greater than 2° ;
- the polar angle of both clusters falls within $2.5^\circ < \theta < 4.5^\circ$.

Fig. 20 shows the reconstructed $\gamma\gamma$ -invariant mass spectrum. It is characterized by two peaks centered at 114.1 ± 2.7 MeV and 129.6 ± 2.1 MeV respectively. The presence of the two peaks and the relative distance are in agreement with the Monte Carlo simulation results reported in Sec. 5.2.2. We interpreted the higher-mass peak as being due to π^0 production from the lead target, and the other peak as being due to π^0 -production from downstream materials (e.g. aluminum windows). The two peaks are centered at lower mass values, relative to the expected values, due to the aforementioned miscalibration. To determine the π^0 production yield, the $\gamma\gamma$ -invariant mass spectrum was fit with two Gaussian functions plus a quadratic polynomial function for the background. The number of events in the first and second peak, determined by integrating the respective

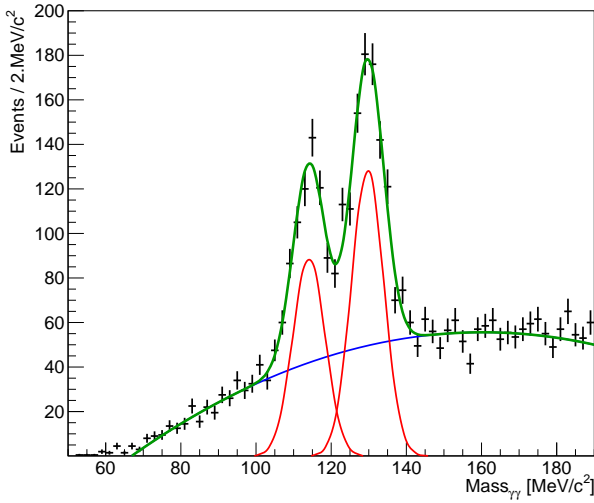


Fig. 20 Distribution of $\gamma\gamma$ invariant mass. The two peaks were fit with Gaussian functions (red dashed lines) plus a quadratic polynomial function for the background (blue line). The green line represents the overall fit. As discussed in Fig. 17, the lower mass peak corresponds to the Al window, the larger mass peak to the Pb target.

Gaussian function from -3 to 3σ , were found to be 966 ± 164 and 1378 ± 275 , respectively. The latter is in agreement within 30% of the theoretical expectation for generated π^0 by the interaction of the beam with the lead target via real and virtual photoproduction mechanisms. The former exceeds the expected yield related to the production from the two Al windows by a factor ~ 4 . This discrepancy could be due to the presence of other materials placed near the two Al windows (e.g. glue, mechanical support) contributing to the π^0 production and consequently increasing the measured yield.

5.2.4 JLAB SRO-DAQ performance

During the Run-2 tests, a study of SRO DAQ performance was conducted. From the front-end, a data rate of about 800 MB/s per uplink was measured with no data frame dropping (100% livetime). Since the setup consisted of 3 VXS crates with 6 fiber uplinks, the total data rate reached up to 4 GB/s.

To study the performance of the back-end, the front-end thresholds and TriDAS parameters (i.e. the number of instances of HMs and TCPUs) were varied. During tests, the memory occupancy and the CPU load per TriDAS process were checked against the data throughput. An uneven distribution of data sources was found to have a significant impact on TriDAS performance. This is not a surprise, since the system was originally designed for a neutrino telescope, where all detection

elements produce almost the same data throughput, providing a well distributed and balanced load to the HM. The best performance was achieved with a single memory assignment to fulfill the requirements of every instantiated HM. However, throughput homogeneity is not guaranteed in CLAS12 streams. The topology of the physics events created sizeable gradients in the throughput across different sectors of the FT-CAL and FT-HODO detectors. The first version of the TriDAS implementation, which is not yet optimized, handles this problem by dimensioning all memory buffers according to the maximum size necessary to accommodate the largest data stream. This of course biases the measured memory occupancy.

The front-end thresholds were varied to provide a data throughput ranging from a few tens of Mbit/s up to almost 100 Mbit/s. The HM processes were instantiated on one Linux server with 48 cores, 1 GHz each and 64 GB RAM. The number of HM instances per run were raised from 5 HMs, 10 HMs and 20 HMs. The detector was subdivided in 5, 10 and 20 sectors, accordingly. The CPU load increased almost linearly with the number of HM instances, 500%, 850% and 1600%, respectively. This is implicit in the multi-threaded design of TriDAS. Meanwhile, the HM memory occupancy remained almost constant at about 12–1 GB per run. This is consistent with the 500 kB/channel/timeslice buffer size, and the fact that the number of HMs is inversely proportional to the number of served channels per HM, which is the total number of FT+Hodo channels, i.e. a constant on the order of ~ 500 .

Ten instances of TCPUs, each capable of handling 5 timeslices at time, run on two CPU servers. As mentioned in Sec. 4.3, the TCPU implements different trigger-level algorithms. The Level 1 performance was found to be strongly affected by hit sorting in the considered timeslice. The profiling of this nonlinear performance was reported in [14]. The Level 2 trigger was not always used, in order to determine the impact of running TriDAS with or without the JANA algorithms. The CPU load per TCPU instance ranged from 400% without any JANA trigger, to 800% including the *standard clustering*, the 1 : 10 scaler and the minimum bias selection algorithms, and, up to 1600% when processing the AI clustering. Generally the memory usage remained within 20–24 GB. However, it doubled when running the AI algorithm, indicating the need for optimization.

Although several positive results, a properly quantification of TRIDAS performance was not possible due to some weakness revealed by the on-beam tests. The main issues were due to the VTP to native TriDAS data format conversion and some system limitations related to a not homogenous load distribution. Both issues were

fixed in the new version of the code. A thorough profiling of TriDAS performance will be performed in the next on-bem tests and results reported in a follow-up publication.

6 Future work

6.1 The ERSAP SRO framework

The Environment for Real-time Streaming, Acquisition and Processing (ERSAP) framework is an effort at JLab to develop streaming readout and data processing systems that will satisfy future experiments at the laboratory which are currently in various stages of planning. The goal is to develop a common framework for building both streaming data acquisition and data stream processing pipelines. The new experiments are eager to take advantage of SRO technologies, due to the well-known intrinsic limitations of the triggered readout systems. The ERSAP framework organizes a data stream processing application into a network of interconnected actors which communicate via message-passing and execute user-defined algorithms referred to as ‘engines’. The framework abstracts these engines as micro-services and provides them with a run-time environment. This includes handling the network programming, data serialization, and general IO. ERSAP also provides tools for engine configuration and scaling, potentially freeing the user from having to write multi-threaded code. The only requirement ERSAP imposes is that the user-defined engines must adhere to a simple data-in/data-out interface.

A key goal for ERSAP is to be flexible enough to allow data processing applications to evolve. It should encourage the implementation of new ideas and technologies while preserving the integrity of existing data pipelines. To this end, it provides three basic components: a reactive Actor micro-service, a communication channel between Actors (which serves as a data-stream pipe or conveyor belt), and finally an application orchestrator. During operation, a stream of data-quanta will flow through the directed graph of reactive microservices, and the network itself will define the higher-level application logic.

One basic difference between ERSAP and other frameworks is that, rather than moving instructions between Actors, the data is moved instead. An incoming data-quantum triggers the execution of an actor (in other words, the actor *reacts* to the incoming message). Communications between actors is restricted to message passing, and the channels between actors are specified externally to the actors themselves. A key consequence of

this design is that actors are programmatically independent, i.e. they can be built and deployed as stand-alone processes.

ERSAP is a framework that uses an independent, reactive actor model and a flow-based programming paradigm. It encourages a functional decomposition of the overall data processing application into small single-function artifacts called micro-services. These artifacts should be easy to understand, develop, deploy and debug. Because they are programmatically independent, they can be readily scaled and individually optimized, unlike the components of a monolithic application. Another important advantage is fault tolerance: individual actors are allowed to crash and be automatically restarted without bringing down the entire pipeline.

6.2 GEMC SRO

The GEMC simulation framework has been updated to collect hits from multiple events based on the Time at the Readout Electronics (TRE) of each Geant4 step, as provided by the digitization routine. The data is organized and collected in the same data structures used by the VTP, called Frames.

During digitization, the TRE is calculated for each Geant4 step by taking into account the track navigation in the various material, the response of the detector, and the propagation delays due to the hardware and cables. A dedicated GRunAction class handles:

- creating the VTP frame buffers
- linking the distribution of hits from various events to frame buffer ids using the TRE
- flushing the buffers to disk

The digitized data identifier is converted into hardware addresses (organized by crate/slot/channel) using the detector translation tables. Each VTP frame refers to the corresponding hardware crate, uniquely identified by absolute time t and the frame duration Δt_f using the formula $id = \text{floor}(t)/\Delta t_f$. The frame data includes:

- the Frame Payload, which contains the collection of hits, represented as a vector of unsigned int, and also includes the hardware address.
- the Frame Header, which contains the payload properties and time of frame.

The distribution of Geant4 events into frames is summarized in Fig. 21.

The GEMC output is available in various formats, which are identical in content: text (ASCII), EVIO (the Jefferson Lab data acquisition format) [29], ROOT and

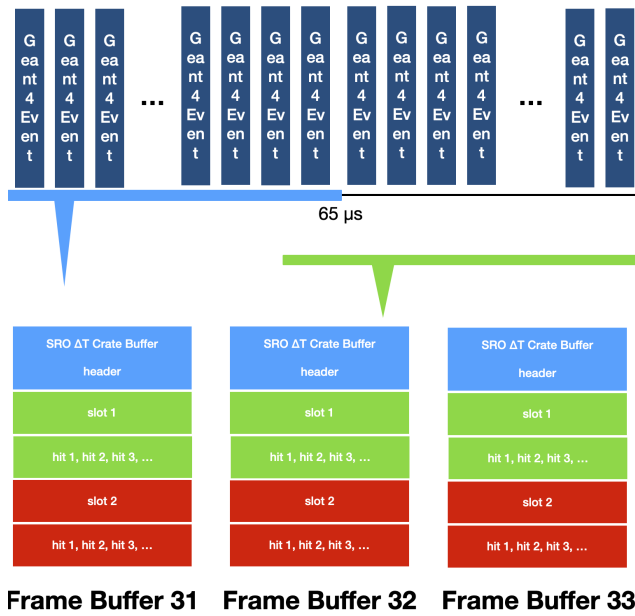


Fig. 21 Hits from a Geant4 event can end on different frames depending on the digitization time, propagation time, etc. Many Geant4 events can have hits in the same Frame. When no more Geant4 event can fill a frame, that frame is written out to disk.

VTPSRO, the binary format used for the CLAS12 SRO. GEMC was used to produce both event-based output and a frame-based streaming readout output that emulates the JLab SRO front-end.

7 Summary

Streaming readout is a powerful and flexible option in data acquisition adopted by many current and future experiments, which would be particularly beneficial for electron machines at the intensity frontier (e.g. JLab and the EIC). Despite its many advantages and the simplicity (in principle) of implementation, the superiority of streaming readout in real applications still needs to be proven. In this paper we reported results of on-beam tests of the first Jefferson Lab implementation of a full SRO DAQ system, including front-end electronics (JLab-FADC250 and INFN-*WaveBoard* digitizers), back-end software (TriDAS) and high-level analysis framework (JANA2). Tests were performed at the lab with two different experimental setups. In the first, we exposed the EIC PbWO₄ crystal EM calorimeter prototype to the Hall-D Pair Spectrometer test beam for a direct comparison of triggered and SRO performance. Results showed that the SRO performed as expected, providing a calorimeter energy resolution compatible with data collected using a traditional triggered DAQ. The second setup included the Hall-B CLAS12

Forward Tagger calorimeter and hodoscope to measure a physics channel, inclusive π^0 hadroproduction, during a standard high-intensity electron-beam production run. Results were compared to the expected yield calculated with a realistic theoretical model of the reaction making use of a sophisticated Geant4 simulation to determine the detector efficiency. The good agreement between prediction and measurement provided a significant validation of the SRO DAQ performance during a realistic electron scattering experiment. Furthermore, the implementation in the framework of AI-supported real-time tagging and selection algorithms, demonstrating how the SRO DAQ provides new capabilities well beyond the standard triggered DAQ. Based on these positive results, the current framework is being upgraded to a microservice architecture (ERSAP) to build a common framework for both streaming readout and offline data processing.

Acknowledgments

We would like to acknowledge the CLAS12 and GlueX collaborations as well as the JLab technical staff for their accommodation and support of this effort.

The INFN Group has been supported by Italian Ministry of Foreign Affairs (MAECI) as Projects of Great Relevance within Italy/US Scientific and Technological Cooperation under grant n. MAE0065689-PGR00799.

This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Nuclear Physics under contract DE-AC05-06OR23177. The work of CF is supported by the U.S. Department of Energy, Office of Science, Office of Nuclear Physics under grant No. DE-SC0019999.

Part of the work was supported by the Jefferson Lab LDRD project INDRA-ASTRA (2020-LDRD-LD2014)

References

1. R. Abdul Khalek et al. Science Requirements and Detector Concepts for the Electron-Ion Collider: EIC Yellow Report, 2021. [arXiv:2103.05419](https://arxiv.org/abs/2103.05419).
2. SoLID Collaboration. SoLID (Solenoidal Large Intensity Device) Updated Preliminary Conceptual Design Report, 2019. URL: <https://hallaweb.jlab.org/12GeV/SoLID/files/solid-precdr-Nov2019.pdf>.
3. MOLLER Collaboration. The MOLLER Experiment: An Ultra-Precise Measurement of the Weak Mixing Angle Using Møller Scattering, 2014. [arXiv:1411.4088](https://arxiv.org/abs/1411.4088).
4. A. Adare et al. An Upgrade Proposal from the PHENIX Collaboration, 2015. [arXiv:1501.06197](https://arxiv.org/abs/1501.06197).
5. V.D. Burkert and others. The CLAS12 Spectrometer at Jefferson Laboratory. *Nucl. Instrum. Meth. A*, 959:163419, 2020. [doi:10.1016/j.nima.2020.163419](https://doi.org/10.1016/j.nima.2020.163419).

6. R. Aaij, J. Albrecht, M. Belous, P. Billoir, T. Boettcher, A. Brea Rodriguez, D. vom Bruch, D. H. Cámpora Pírez, A. Casais Vidal, D. C. Craik, et al. Allen: A High-Level Trigger on GPUs for LHCb. *Computing and Software for Big Science*, 4(1), Apr 2020. URL: <http://dx.doi.org/10.1007/s41781-020-00039-7>, doi:10.1007/s41781-020-00039-7.
7. P. Buncic, M. Krzewicki, and P. Vande Vyvre. Technical Design Report for the Upgrade of the Online-Offline Computing System. 4 2015.
8. The LHCb Collaboration. Upgrade Software and Computing, 2018. URL: <http://cds.cern.ch/record/2310827>.
9. David Rohr. Usage of GPUs in ALICE Online and Offline processing during LHC Run 3. *EPJ Web Conf.*, 251:04026, 2021. arXiv:2106.03636, doi:10.1051/epjconf/202125104026.
10. F. Barbosa, C. Hutton, A. Sitnikov, A. Somov, S. Somov, and I. Tolstukhin. Pair spectrometer hodoscope for Hall D at Jefferson Lab. *Nucl. Instrum. Meth.*, A795:376–380, 2015. doi:10.1016/j.nima.2015.06.012.
11. A. Somov, I. Tolstukhin, S.V. Somov, and V.V. Berdnikov. Commissioning of the Pair Spectrometer of the GlueX experiment. *J. Phys. Conf. Ser.*, 798(1):012175, 2017. doi:10.1088/1742-6596/798/1/012175.
12. A. Acker et al. The CLAS12 Forward Tagger. *Nucl. Instrum. Meth. A*, 959:163475, 2020. doi:10.1016/j.nima.2020.163475.
13. S. Boyarinov, B. Raydo, C. Cuevas, C. Dickover, H. Dong, G. Heyes, D. Abbott, W. Gu, V. Gyurjyan, E. Jastrzembski, B. Moffit, C. Timmer, I. Mandjavidze, D. Heddle, C. Smith, A. Celentano, R. De Vita, N. Baltzell, K. Livingston, B. McKinnon, and W. Moore. The CLAS12 Data Acquisition System. *Nucl. Instrum. Meth. A*, 966:163698, 2020. doi:10.1016/j.nima.2020.163698.
14. T. Chiarusi, M. Favaro, F. Giacomini, M. Manzali, A. Margiotta, and C. Pellegrino. The Trigger and Data Acquisition System for the KM3NeT-Italy neutrino telescope. *Journal of Physics: Conference Series*, 898:032042, oct 2017. doi:10.1088/1742-6596/898/3/032042.
15. F. Ameli, M. Battaglieri, M. Bondi, A. Celentano, S. Boyarinov, N. Brei, T. Chiarusi, R. De Vita, C. Fanelli, V. Gyurjyan, D. Lawrence, P. Musico, C. Pellegrino, B. Raydo, and S. Vallarino. Streaming Readout of the CLAS12 Forward Tagger Using TriDAS and JANA2. *EPJ Web Conf.*, 251:04011, 2021. doi:10.1051/epjconf/202125104011.
16. S. Aiello et al. Measurement of the atmospheric muon depth intensity relation with the NEMO Phase-2 tower. *Astroparticle Physics*, 66:1–7, 2015. URL: <https://www.sciencedirect.com/science/article/pii/S0927650514001960>, doi:https://doi.org/10.1016/j.astropartphys.2014.12.010.
17. AI4EIC-exp: I Workshop on Artificial Intelligence for the Electron Ion Collider — experimental applications. 2021. <https://indico.bnl.gov/event/10699/>. URL: <https://indico.bnl.gov/event/10699/>.
18. R. J. G. B. Campello, D. Moulavi, A. Zimek, and J. Sander. Hierarchical density estimates for data clustering, visualization, and outlier detection. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 10(1):1–51, 2015. doi:10.1145/2733381.
19. J. A. Hartigan and M. A. Wong. Algorithm AS 136: A k -means clustering algorithm. *Journal of the royal statistical society. series c (applied statistics)*, 28(1):100–108, 1979. doi:10.2307/2346830.
20. How HDBSCAN works. hdbscan 0.8.1 documentation. URL: https://hdbscan.readthedocs.io/en/latest/how_hdbscan_works.html.
21. L. McInnes and J. Healy. Accelerated Hierarchical Density Based Clustering. In *Data Mining Workshops (ICDMW), 2017 IEEE International Conference on*, pages 33–42. IEEE, 2017. arXiv:1705.07321, doi:10.1109/ICDMW.2017.12.
22. T. Horn et al. Scintillating crystals for the Neutral Particle Spectrometer in Hall C at JLab. *Nucl. Instrum. Meth. A*, 956:163375, 2020. arXiv:1911.11577, doi:10.1016/j.nima.2019.163375.
23. Y.-S. Tsai. Pair Production and Bremsstrahlung of Charged Leptons. *Rev. Mod. Phys.*, 46:815, 1974. [Erratum: *Rev. Mod. Phys.* 49, 521(1977)]. doi:10.1103/RevModPhys.46.815, 10.1103/RevModPhys.49.421.
24. V. Mathieu, G. Fox, and A. P. Szczepaniak. Neutral Pion Photoproduction in a Regge Model. *Phys. Rev. D*, 92(7):074013, 2015. arXiv:1505.02321, doi:10.1103/PhysRevD.92.074013.
25. M. Carver et al. Photoproduction of the $f_2(1270)$ meson using the CLAS detector. *Phys. Rev. Lett.*, 126(8):082002, 2021. arXiv:2010.16006, doi:10.1103/PhysRevLett.126.082002.
26. V. Mathieu, A. Pilloni, M. Albaladejo, Ł. Bibrzycki, A. Celentano, C. Fernández-Ramírez, and A. P. Szczepaniak. Exclusive tensor meson photoproduction. *Phys. Rev. D*, 102(1):014003, 2020. arXiv:2005.01617, doi:10.1103/PhysRevD.102.014003.
27. M. Ungaro et al. The CLAS12 Geant4 simulation. *Nucl. Instrum. Meth. A*, 959:163422, 2020. doi:10.1016/j.nima.2020.163422.
28. V. Ziegler et al. The CLAS12 software framework and event reconstruction. *Nucl. Instrum. Meth. A*, 959:163472, 2020. doi:10.1016/j.nima.2020.163472.
29. Jefferson Lab Data Acquisition, The EVIO data format. URL <https://codajlab.org/drupal/content/event-io-evio>.