CLAS12 Remote Data Stream Processing

Vardan Gyurjyan[1], Michael Goodrich[1], Derek Howard[2], Graham Heyes[1], Yatish Kumar[2], David Lawrence[1], Stacy Sheldon[2], Carl Timmer[1], Jeng-Yuan Tsai[1], Nicholas Tyler[3]

[1]Thomas Jefferson National Accelerator Facility, Newport News, VA 23606, USA
[2]Energy Science Network (ESnet), Berkeley, CA 94720, USA
[3]National Energy Research Scientific Computing Center (NERSC), Berkeley, CA 94720, USA

E-mail: gurjyan@jlab.org

Abstract. Implementing a physics data processing application is relatively straightforward with the use of current containerization technologies and container image runtime services, which are prevalent in most high-performance computing (HPC) environments. However, the process is complicated by the challenges associated with data provisioning and migration, impacting the ease of workflow migration and deployment. Transitioning from traditional file-based batch processing to data-stream processing workflows is suggested as a method to streamline these workflows. This transition not only simplifies file provisioning and migration but also significantly reduces the necessity for extensive disk space. Data-stream processing is particularly effective for real-time processing during data acquisition, thereby enhancing data quality assurance. This paper introduces the integration of the JLAB CLAS12 event reconstruction application within the ERSAP data-stream processing framework that facilitates the execution of streaming event reconstruction at a remote data center and enables the return streaming of reconstructed events to JLAB while circumventing the need for temporary data storage throughout the process.

1 Introduction

At Jefferson Lab (JLAB), four experimental halls designated A, B, C, and D are at the forefront of this transformation, transitioning from triggered data acquisition to streaming readout systems. These upgrades are crucial for the Electron-Ion Collider (EIC) project, emphasizing the need for real-time data processing and analysis to harness the full potential of scientific Big Data.

The primary goal of these advancements is to enable a more comprehensive understanding of the physical phenomena under investigation. By capturing data in a continuous stream, researchers can obtain a more detailed and accurate representation of experimental outcomes. This shift from traditional data acquisition methods to streaming readout systems aligns with the broader trend in scientific research toward real-time data processing and analysis.

2 The JLAB Grand Challenge

Many current nuclear and high-energy and nuclear physics experiments are upgrading their triggered data acquisition systems to streaming readout models. However, these experiments often need an integrated system from data acquisition (DAQ) to analysis. The JLAB Grand Challenge aims to eliminate the separation between data readout and analysis by leveraging modern electronics, computing, and analytical techniques to build the next-generation computing model essential for probing femto-scale science.

It integrates detector readout and analysis to enhance scientific output, supported by a robust software and hardware stack, including the Common Data Acquisition (CODA) system, the Event-Driven Reactive Stream Analysis Processing (ERSAP) framework, and the Jefferson Lab Integrated Research Infrastructure Across Facilities (JIRIAF).

3 CODA Data Acquisition System

CODA is a hybrid data acquisition system is designed to handle both triggered and streaming data. It is a data-centric system where components communicate solely through data, allowing flexibility and scalability. The current version supports streaming readout, making it a unique tool for managing the vast amounts of data generated by JLAB's experiments.

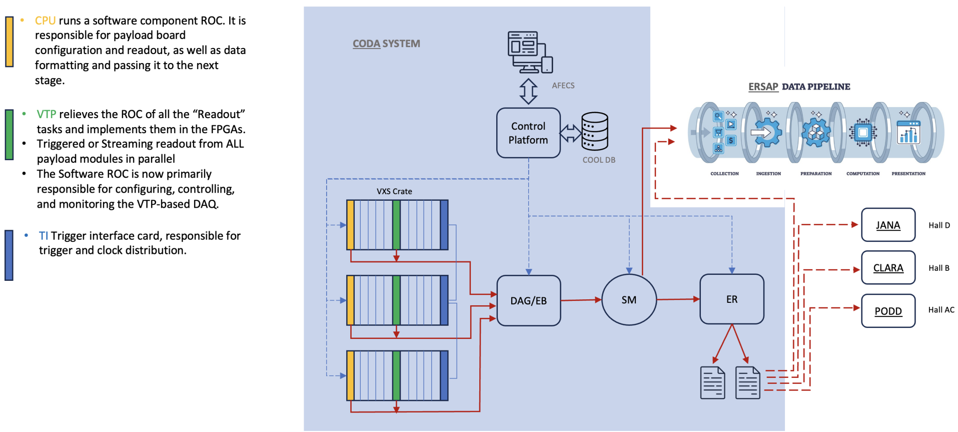

Figure 1: The CODA (CEBAF Online Data Acquisition) system consists of fundamental components including the Data Aggregator (DAG), Event Builder (EB), Event Recorder (ER), and Shared Memory (SM). These components have corresponding agents within the AFECS (multi-Agent Framework for Experiment Control Systems) used for run control, monitoring, and the creation of online/DAQ state machines. The data collected by CODA is subsequently processed offline by JLAB experimental Hall-specific data processing frameworks, such as HallAC PODD, HallB CLARA, and HallD JANA.

CODA components, including the Readout Controller (ROC), Data Aggregator (DAG), Event Builder (EB), and Shared Memory (SM), operate independently yet cohesively. This modular design allows for easy integration and customization to meet the specific needs of different experiments [1]. The intelligent agent-based run-control system orchestrates the entire process, ensuring efficient data acquisition and processing [2].

4 ERSAP: A Framework for Real-time Data Stream Processing

The ERSAP framework leverages the Flow-Based Programming (FBP) paradigm to facilitate the design and deployment of distributed data processing applications. It is built to support multilingual and heterogeneous hardware integration, thereby simplifying the development of complex data processing workflows [3]. By representing user algorithms as event-reactive actors, ERSAP enables both local and remote data processing environments, enhancing scalability and performance.

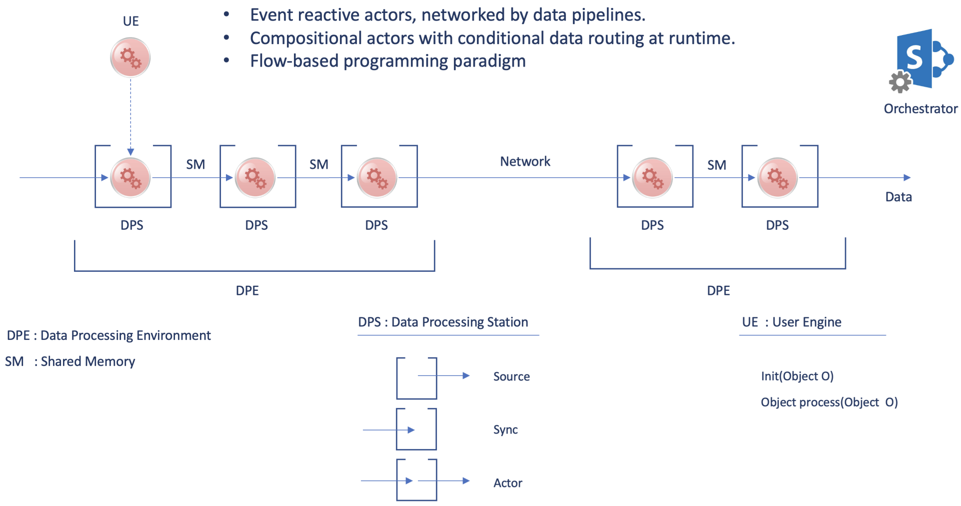

Figure 2: Core components of the ERSAP framework.

ERSAP applications are structured as directed graphs where nodes represent actors and arcs represent data links. This architecture allows for the seamless integration of diverse technologies and hardware accelerators within the same application. The framework's primary goal is to provide flexibility for the creation and evolution of data processing applications, fostering the implementation of new ideas and technologies while preserving the integrity of existing data pipelines.

The framework's core components include reactive actors, data-stream pipes, communication channels between actors, and an orchestrator or workflow manager (see Fig. 2). User-provided data-event processing algorithms, referred to as User Engines (UE), adhere to a simple interface. The Data Processing Station (DPS) uses this interface to represent the UE as an actor within ERSAP. DPSs are deployed within Data Processing Environments (DPE), which provide a runtime environment for actors. For optimization, data transport between actors within a DPE utilizes shared memory, thereby avoiding data copies. ERSAP defines three types of actors based on input and output assignments, allowing for multiple inputs and outputs.

ERSAP applications are described and composed using YAML declarative syntax, supporting conditional routing at runtime. The graphical designer software used to build these applications provides an intuitive interface for constructing and managing data processing pipelines, further simplifying the development process.

5 Concept Validation and Experiments

The CLAS12 (CEBAF Large Acceptance Spectrometer for operation at 12 GeV beam energy) detector at Jefferson Lab represents a pivotal instrument in experimental nuclear physics, meticulously engineered to deepen our understanding of hadronic physics and nuclear interactions [4]. This sophisticated apparatus is integral to investigating the intricate structure of protons, neutrons, and atomic nuclei. Positioned within the Continuous Electron Beam Accelerator Facility (CEBAF), CLAS12 is comprehensively designed to explore the dynamic properties of quarks and gluons—the fundamental building blocks of matter—via high-precision electron scattering experiments.

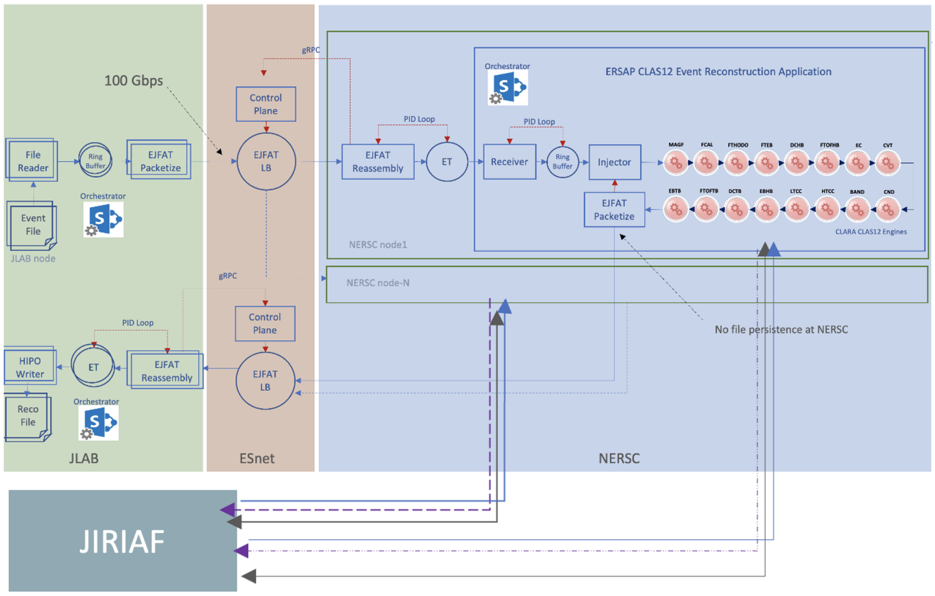

Figure 3: CLAS12 Remote Data Stream Processing Data-Flow Diagram. The EJFAT packetizer segments data events into UDP packets, facilitating their transport. The EJFAT reassembly engine reconstructs these UDP packets into a cohesive data event for backend processing. The ET and RingBuffer components at the backend function as in-memory FIFO buffers to mitigate network latency and spikes. Specifically, the RingBuffer is an integral component of the ERSAP receiver actor. The nodes in the ERSAP data pipeline represent CLAS12 reconstruction actors that process data corresponding to specific CLAS12 detector components. Blue lines connecting JIRIAF [6] to the backend illustrate remote resource allocation requests, while the black line indicates that these resources are part of the JIRIAF distributed Kubernetes cluster, which communicates with the Kubernetes control plane at JLAB. The dotted line represents the Prometheus monitoring data path from the ERSAP and JIRIAF JRM monitoring server.

In a demonstration of the EJ-FAT system's capabilities, a dataset from a CLAS12 experiment was replayed as if it were streaming directly from the data acquisition system during live experiments. This transition from traditional file/batch processing to a streaming processing paradigm was undertaken to rigorously test the EJ-FAT system's proficiency in real-time data stream processing applications. The operational CLAS12 reconstruction application, initially designed and executed within the CLARA framework [5], was modified and re-implemented within the ERSAP framework. This modification was crucial in developing a streaming, real-time event reconstruction application that leveraged existing CLAS12 subdetector reconstruction algorithms. Each component's reconstruction algorithms within the CLAS12 detector were encapsulated as ERSAP data-processing reactive actors. These actors were interconnected to form comprehensive applications capable of accurately identifying particles and determining their momenta, trajectories, and energies before and after interactions. The precision and accuracy of this reconstructed data were subsequently verified by a physics validation actor, ensuring the reliability of the reconstruction application. In a proof-of-concept data stream processing experiment, raw experimental data from CLAS12 was successfully streamed to the National Energy Research Scientific Computing Center (NERSC) through the EJ-FAT load balancer at the Energy Sciences Network (ESnet) at impressive rates exceeding 100 Gbps. This data underwent massive parallel processing on 40 Perlmutter nodes, utilizing over 10,000 cores, and was then streamed back to Jefferson Lab in real time for validation, persistence, and final physics analysis (see Figure 5).Figure 6 presents sample plots obtained during remote data processing scaled at NERSC, while Figure 6b illustrates a 100 Gbps data flow through the EJFAT LB, as monitored by ESnet.

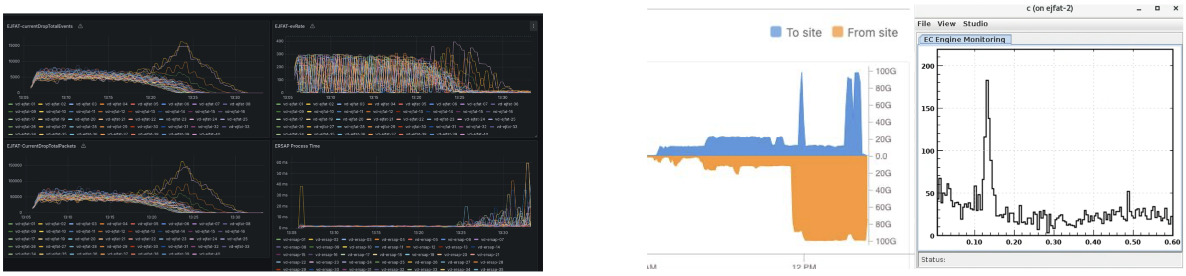

(a) Forty horizontally scaled data processing pipelines are being activated.

(b) These plots illustrate data throughput into and out of ESnet during the concept validation experiment, and a plot from the data processing validation highlighting the pi0 missing mass reconstruction.

Figure 4: Example of commissioning plots for remote data stream processing.

This innovative data processing pipeline is significant as it eliminates intermediate data storage along the data-stream path, with data originating and being stored exclusively at Jefferson Lab. This approach effectively reduces file I/O latencies, enhancing overall processing efficiency.

6 Conclusion

In conclusion, the integration of the JLAB CLAS12 event reconstruction application within the ERSAP framework represents a significant advancement in real-time data processing for high-energy and nuclear physics experiments. By transitioning from traditional file-based batch processing to a more efficient data-stream processing paradigm, the CLAS12 remote data stream processing initiative has demonstrated the capability to handle large volumes of experimental data with enhanced efficiency and reduced latency. This innovative approach not only simplifies data provisioning and migration but also minimizes the need for extensive disk space, enabling real-time data processing and quality assurance during data acquisition.

The successful implementation and validation of the EJ-FAT system and the ERSAP framework in real-time, remote data stream processing experiments highlight the potential for scalable, distributed data processing solutions in scientific research. The ability to stream data at rates exceeding 100 Gbps and perform massive parallel processing at remote data centers, such as NERSC, underscores the robustness and scalability of the developed system.

This work lays the groundwork for future developments in the field of high-energy and nuclear physics, providing a flexible and efficient framework for real-time data processing. The methodologies and technologies introduced here can be adapted and expanded to meet the evolving needs of scientific research, ultimately contributing to a deeper understanding of fundamental physical phenomena. The collaboration between JLAB, ESnet, and NERSC serves as a model for future endeavors, demonstrating the power of integrated, interdisciplinary efforts in advancing scientific discovery.

References

1. W. A. Watson III, J. Chen, and G. Heyes, "CODA: A scalable, distributed data acquisition system," in IEEE Transactions on Nuclear Science, 1994. OSTI ID:54023.

2. V. Gyurjyan, D. Abbott, G. Heyes, E. Jastrzembski, C. Timmer, and E. Wolin, "AFECS. multi-agent framework for experiment control systems," Journal of Physics: Conference Series, vol. 119, no. 2, p. 022025, 2008.

3. V. Gyurjyan, D. Abbott, N. Brei, M. Goodrich, G. Heyes, E. Jastrzembski, D. Lawrence, B. Raydo, and C. Timmer, "ERSAP: Toward Better NP Data-Stream Analytics With Flow-Based Programming," IEEE Transactions on Nuclear Science, vol. 70, no. 6, Jun. 2023.

4. M. Goodrich, C. Timmer, V. Gyurjyan, D. Lawrence, G. Heyes, Y. Kumar, and S. Sheldon, "ESnet/JLab FPGA Accelerated Transport," IEEE Transactions on Nuclear Science, vol. 70, no. 6, Feb. 2023. DOI: 10.1109/TNS.2023.3243871.

5. V. Gyurjyan, et al., "CLARA: A Contemporary Approach to Physics Data Processing," J. Phys.: Conf. Ser., vol. 331, p. 032013, 2011. DOI: 10.1088/1742-6596/331/3/032013.

6. V. Gyurjyan, C. Larrieu, G. Heyes, and D. Lawrence, "Jiriaf: Jlab integrated research infrastructure across facilities," in EPJ Web of Conferences, vol. 295, p. 04027, 2024.