RTDP:Streaming Readout Real-Time Development and Testing Platform

Ayan Roy[1], David Lawrence[1], Jeng-Yuan Tsai[1], Marco Battaglieri[2], Markus Diefenthaler[1], Vardan Gyurjyan[1], Xinxin Mei[1]

[1]Thomas Jefferson National Accelerator Facility, Newport News, VA 23606, USA
[2]The National Institute for Nuclear Physics (INFN), Italy

E-mail: {ayan, davidl, tsai mdiefent, gurjyan, xmei}@jlab.org, battaglieri@ge.infn.it

Abstract. The Thomas Jefferson National Accelerator Facility (JLab) has created and is currently working on various tools to facilitate streaming readout (SRO) for upcoming experiments. These include reconstruction frameworks with support for Artificial Intelligence/Machine Learning, distributed High Throughput Computing (HTC), and heterogeneous computing which all contribute significantly to swift data processing and analysis. Designing SRO systems that combine such components for new experiments would benefit from a platform that would combine both simulation and execution components for simulation, testing, and validation before large investments are made. The Real-Time Development Platform (RTDP) is being developed as part of an LDRD funded project at JLab. RTDP aims to establish a seamless connection between algorithms, facilitating the seamless processing of data from SRO to analysis, as well as enabling the execution of these algorithms in various configurations on compute and data centers. Individual software components simulating specific hardware can be replaced with actual hardware when it is available.

1 Introduction

With the rapid commencement of new experiments that leads to generation of huge volumes of data, the Nuclear Physics community is actively exploring the next generation of data processing and analysis workflows to optimize scientific output. The major motivation in doing so is to achieve seamless data processing from detector readout to physics analysis and enabling the rapid turnaround of data for publications through advanced scientific computing.

Jefferson Lab (JLab) has been at the forefront of such progress and has made substantial investments in the development of the next generation of data processing and analysis workflows. Many essential components required for seamless data processing and rapid data turnaround, such as streaming readout, heterogeneous computing, and AI/ML capabilities, are already available. Furthermore, JLab is appropriately equipped with facilities to generate large amount of copious data from current running experiments, as well as the number of scheduled experiments that contributes towards developing the final system in real conditions.

In this paper, we present the Real-Time Development Platform (RTDP) framework, which leverages both existing and forthcoming infrastructure and software at Jefferson Lab, including JANA [10][11], EJFAT [7], PHASM [5] and HPDF [1]. This framework creates a seamless transition from data generated by the Streaming Read-out (SRO) system using JIRIAF [12], an example is shown in Figure 1, to offline data analysis. By enabling the execution of diverse algorithms on data collected from various

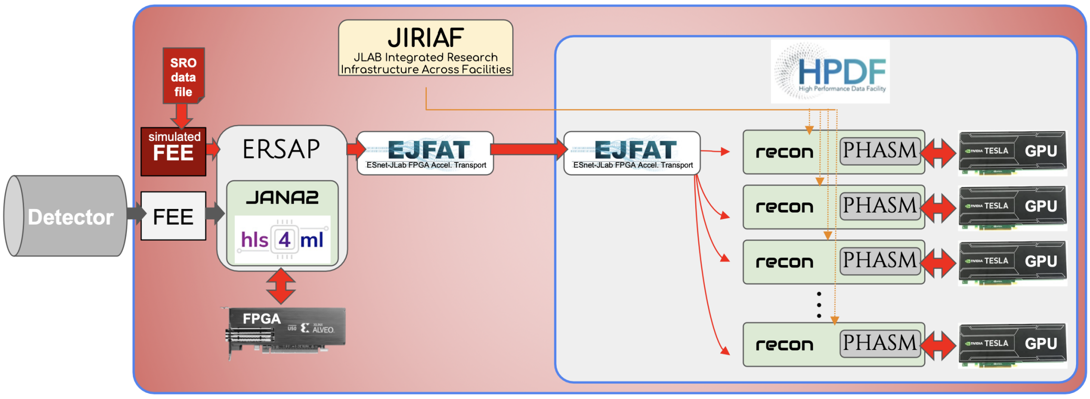

Figure 1: Diagram illustrating an example SRO system where data is streamed from the detector to numerous nodes at an High Throughput Computing (HTC) facility such as HPDF. Red items and arrows indicate pieces that will be developed as part of the SRO design and validation platform.

configurations, the RTDP framework allows for precise estimation of costs and impacts in real-world deployments. The RTDP project fills a critical gap in the development and deployment of full streaming systems for future JLab experiments like SoLID [6], CLAS12 operations at high luminosity [4] ,TDIS [8], and ePIC [9] at EIC. The other most significant benefit of an SRO system when compared to a traditional triggered system is the ability to base decisions on which data to keep and which to discard using information from the whole detector [3].

The RTDP project is an ongoing work in progress that is a natural step forward following previous work to develop prototype SRO systems at JLab[2].

2 Application

The application of a fully developed RTDP can be multi fold. Figure 1 illustrates a simple configuration that could be implemented by the platform This includes:

- SRO experiments that require intricate configurations can be defined with user-friendly language.
- Individual components, such as calibration or data transport, can be represented by software simulation modules. This approach allows for easy configuration for experimentation, provides estimations of various real-world impacts, and enables faster interpretation of results.
- Provides us with the capability to incorporate simulation components along with the real world infrastructure. This allows us to make our simulation more realistics while adhering to the real world constraints for different experiments.
- This approach offers the advantage of scaling the simulation system from a fully simulated software setup on a single computer to a hardware-leveraging model in a distributed configuration.

3 Goal and Motivation

Experimental Nuclear Physics is moving towards a SRO paradigm which includes complex pipelines integrating heterogeneous hardware, and varied software which may have interference effects. Under such circumstances, simulation and testing of complex SRO systems in the offline mode is needed to assist in their design and validation before deploying it in real world. Furthermore, testing of complete, integrated SRO systems at scale for future experiments requires new tooling.

The goal of this project is to design a platform that can facilitate the seamless processing of data from an SRO system to the analysis on compute centers under various configurations. Upon completion of the project, the RTDP framework would be capable of monitoring the various components of the system in a fully developed streaming system. Furthermore, a fully developed RTDP allows us to simulate a real-time SRO data processing network from front end electronics to large compute.

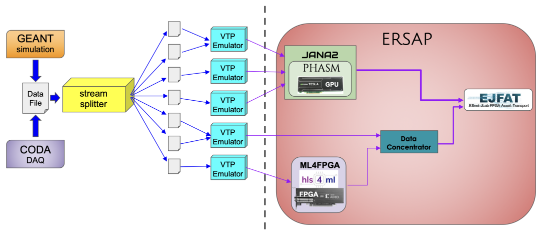

(a) Illustration of the first half of a simple configuration that SRO RTDP would support. Components on the left side of the dashed line would make up the multi-stream source that could be replaced with live data streams from a real detector. To the right of the dashed line is a multi-architectural configuration that communicates the filtered and partially processed data to a remote site via EJFAT

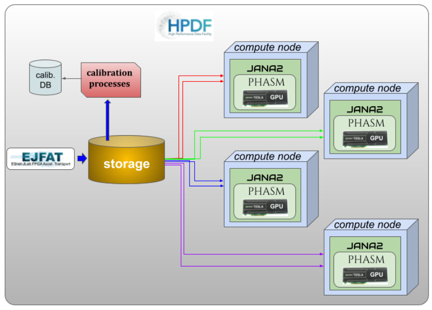

(b) An example configuration supported by RTDP for the remote site. Data received from EJFAT is sent to temporary storage while near-real-time calibrations are performed. Once the calibrations are ready, data is distributed to the compute resources which may also implement multiple compute architectures.

Figure 2: Example of SRO system

4 Objectives

The main objectives of the RTDP is as follows:

- Deployment of a (quasi) real-time SRO data processing model in a distributed setup that includes data calibration, as well as total reconstruction of data in a fully offline setup.
- Use existing infrastructure for framework optimization using the GEANT4-generated and archived beam-on data as well as optimizing the framework validation with beam-on tests.
- Assessment of the need/capability of the available network and computing resources required to run the RTDP in both offline and online settings.
- As RTDP would allow for seamless integration of simulation software with real world hardware components, we can also leverage RTDP for optimizing performance of the hardware platforms with in-depth experiments under different configurations.
- RTDP framework will be leveraged to identify potential issues relevant to a future High-Performace Data Facility (HPDF) in receiving and processing SRO data.

5 Measures of Success

The metrics for measuring a successful RTDP will include the following:

- Ability to launch synchronized processes across multiple nodes to emulate different VTP modules representing real world detectors.
- Integrated monitoring of all components in the system that helps to determine any abberations in the performance of the system.
- Ability to configure and simulate an experiment similar in size to the planned SoLID experiment at JLab.
- Test with a 400Gbps transfer speed using a configuration with at least one FPGA and one GPU component.

6 Progress

In December 2023, a data capture exercise was performed where we were able to capture CLAS12 data streamed across the JLab campus using a 100 Gb/s high speed Network Interface Card (NIC) featuring hardware timestamps. The data was captured using synchronized streams from multiple network sources. During the data capture exercise the Cebaf Online Data Acquisisition system (CODA) had a configuration

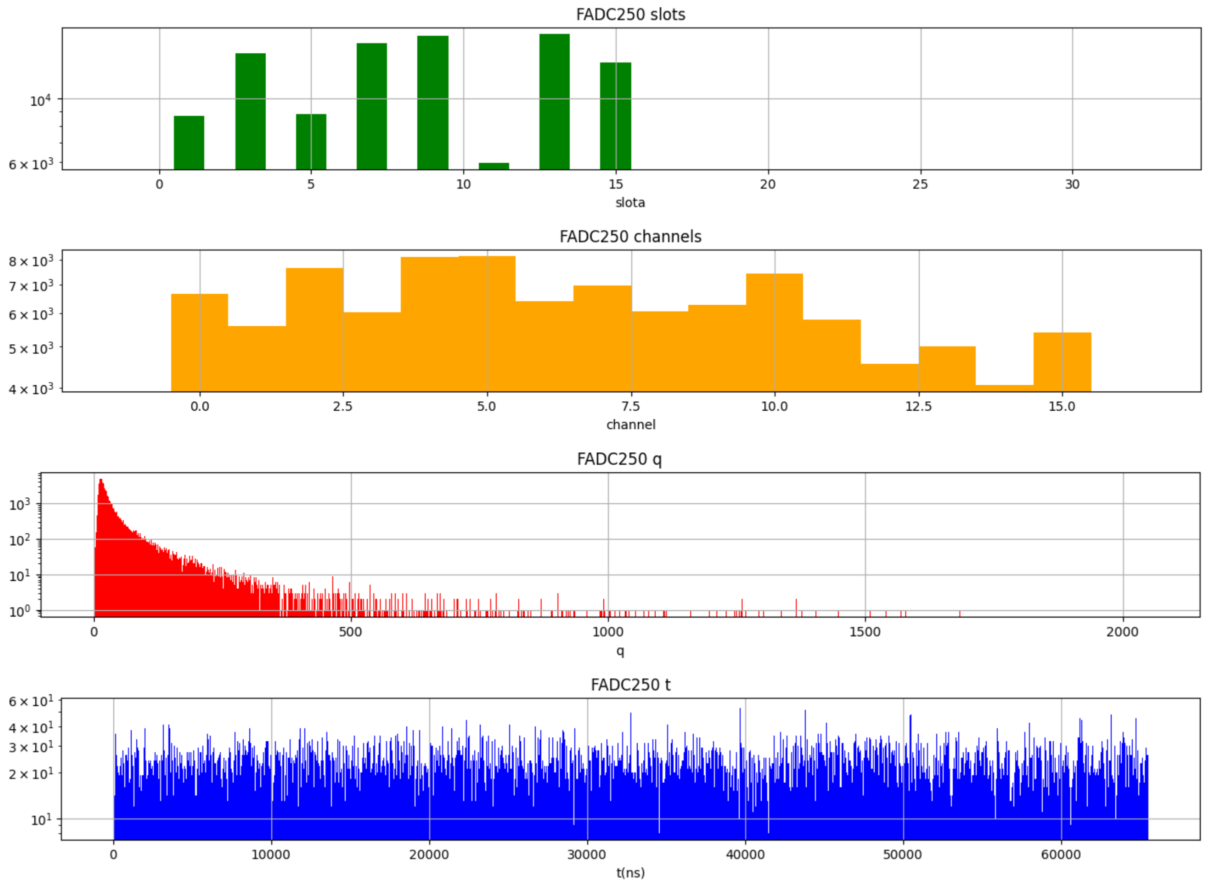

Figure 3: Flash ADC values extracted from the payloads of the data captured in Dec. 2023. **top:** (green) slot number of module in crate (0-23) **second:** (orange) channel number (0-15) **third:** (red) ADC pulse integral **bottom:** (blue) time relative to frame start

that was limited to 2 crates in a single sector. These corresponded to two different calorimeter detectors (ECAL and PCAL) being readout. The data was captured at multiple beam currents: $10, 25, 50, 75, 90$, and 100 nAs respectively. Data captured during this exercise is shown in Figure 3 and 4. In the data capture exercise, packets were captured using *tcpdump* from 4 different TCP ports, but the plots represent data from a single, example port. In Figure 4, the red points represent the beam current during the same period, plotted against the right-hand axis. The gradual decline in data rate is attributed to the liquid hydrogen target depleting over this time. The exact cause of the spikes in the data rate is unknown. This shows that even during periods of steady beam operation, the network traffic may exhibit structure.

7 Future Work and Conclusion

Thus to conclude, in this paper we highlight the initial data capturing exercise that was successfully performed in December 2023. We have presented some of the initial results of the analysis that was performed on the data captured during the exercise. For future we plan to create different simulated VTP emulators and a synchronizer program that will be leveraged for the offline simulations and replay the data captured during the data capture exercise. We will also be working towards developing appropriate schema for monitoring all aspects of monitoring system, establishing databases for monitoring system using existing JLab servers as well as integrating existing JLab monitoring component to SRO-RTDP.

8 Acknowledgements

Jefferson Science Associates, LLC operated Thomas Jefferson National Accelerator Facility for the United States Department of Energy under U.S. DOE Contract No. DE-AC05-06OR23177.

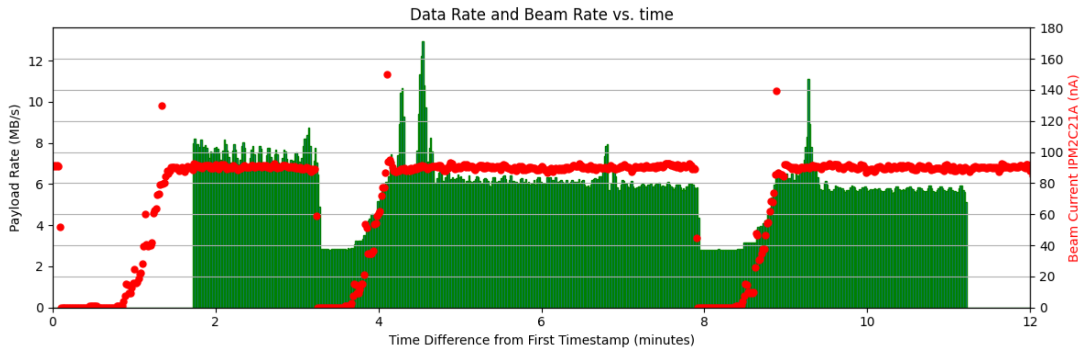

Figure 4: Data rate of captured packets (payload only) vs. time in minutes relative to start run time.

References

[1] https://www.hpdf.science/, 2024.

[2] Fabrizio Ameli, Marco Battaglieri, Vladimir V Berdnikov, Mariangela Bondí, Sergey Boyarinov, Nathan Brei, Andrea Celentano, Laura Cappelli, Tommaso Chiarusi, Raffaella De Vita, et al. Streaming readout for next generation electron scattering experiments. *The European Physical Journal Plus*, 137(8):958, 2022.

[3] JC Bernauer, Cameron Thomas Dean, C Fanelli, J Huang, K Kauder, D Lawrence, JD Osborn, C Paus, JK Adkins, Y Akiba, et al. Scientific computing plan for the ecce detector at the electron ion collider. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 1047:167859, 2023.

[4] S Boyarinov, B Raydo, C Cuevas, C Dickover, H Dong, G Heyes, D Abbott, W Gu, V Gyurjyan, E Jastrzembski, et al. The clas12 data acquisition system. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 966:163698, 2020.

[5] Nathan Brei, Xinxin Mei, Daniel Lersch, Kishansingh Rajput, and David Lawrence. Phasm: A toolkit for creating ai surrogate models within legacy codebases. Technical report, Thomas Jefferson National Accelerator Facility (TJNAF), Newport News, VA . . . , 2024.

[6] SoLID Collaboration et al. Solid (solenoidal large intensity device) updated preliminary conceptual design report. *Jefferson Lab Hall A SoLID Experiment, URL https://solid. jlab. org*, 2019.

[7] Michael Goodrich, Carl Timmer, Vardan Gyurjyan, David Lawrence, Graham Heyes, Yatish Kumar, and Stacey Sheldon. Esnet/jlab fpga accelerated transport. *IEEE Transactions on Nuclear Science*, 70(6):1096–1101, 2023.

[8] A Hall. Measurement of tagged deep inelastic. 2015.

[9] R Abdul Khalek, A Accardi, J Adam, D Adamiak, W Akers, M Albaladejo, A Al-Bataineh, MG Alexeev, F Ameli, P Antonioli, et al. Science requirements and detector concepts for the electron-ion collider: Eic yellow report. *Nuclear Physics A*, 1026:122447, 2022.

[10] David Lawrence. Multi-threaded event reconstruction with jana. In *Journal of Physics: Conference Series*, volume 119, page 042018. IOP Publishing, 2008.

[11] David Lawrence, Amber Boehnlein, and Nathan Brei. Jana2 framework for event based and triggerless data processing. In *EPJ Web of Conferences*, volume 245, page 01022. EDP Sciences, 2020.

[12] Gyurjyan Vardan, Larrieu Christopher, Heyes Graham, and Lawrence David. Jiriaf: Jlab integrated research infrastructure across facilities. In *EPJ Web of Conferences*, volume 295, page 04027. EDP Sciences, 2024.