Date: June 23, 2020

EIC Detector R&D Proposal and Progress Report

Project ID: eRD23

Project Name: Streaming readout for EIC detectors

Period Reported: from 1/1/2020 to 6/26/2020

Project Leader: M. Battaglieri and J. C. Bernauer

Contact Person: M. Battaglieri and J. C. Bernauer

Project members

J. Huang, M.L. Purschke
Brookhaven National Laboratory, Uptown, NY

S. Ali, V. Berdnikov, T. Horn, M. Muhoza, I. Pegg, R. Trotta
Catholic University of America, Washington DC

M. Battaglieri, M. Bondí, A. Celentano, L. Marsicano, P. Musico, S. Vallarino
INFN, Genova, Italy

F. Ameli
INFN, Roma La Sapienza, Italy

D. K. Hasell, C. Fanelli, I. Friščić, R. Milner
Massachusetts Institute of Technology, Cambridge, MA

J. C. Bernauer
Stony Brook University, Stony Brook, NY and Riken BNL Research Center, Upton, NY

E. Cline
Stony Brook University, Stony Brook, NY

C. Cuevas, M. Diefenthaler, R. Ent, Y. Furletova, G. Heyes, B. Raydo
Thomas Jefferson National Accelerator Facility, Newport News, VA

Abstract

The detectors foreseen for the future Electron-Ion Collider will be some of the few major collider detectors to be built from scratch in the 21$^{\text{st}}$ century. A truly modern EIC detector design must be complemented with an integrated, 21$^{\text{st}}$ century readout scheme that supports the scientific opportunities of the machine, improves time-to-analysis, and maximizes the scientific output. A fully streaming readout (SRO) design delivers on these promises, however, it can also impose limitations on the sub-detectors and electronics. The streaming readout consortium will research the design space by evaluating and quantifying the parameters for a variety of streaming readout implementations and their implications for sub-detectors by using on-going work on streaming readout, as well as by constructing a few targeted prototypes particularly suited for the EIC environment.

1 Past

What was planned for this period? What was achieved? What was not achieved, why not, and what will be done to correct? How did the COVID-19 pandemic and related closing of labs and facilities affect progress of your project?

We had a very successful workshop in May, organized by Jefferson Lab. Originally planned to be hosted at Christopher Newport University, the COVID pandemic forced us to meet online instead. Despite this, we had a record attendance, many interesting talks, and very fruitful discussions. The topics discussed included:

- Beam tests: We discussed tests of SRO systems in test beams for CLAS12 and TPEX. Preliminary results were very encouraging, and a lot of interesting feedback about the real-world usage of SRO was gained. Data analysis is still progressing, and future beam times are planned.

- Electronics: We discussed existing SRO-capable readout ASICs developed for CERN and BESIII experiments like the TIGER and SAMPA designs. Designs like these can already fulfill many requirements and will likely be the basis for additional ASIC design. The community feels that the timeline for new ASIC designs becomes more and more tight. It seems critical that the detector requirements are fixed as soon as possible to allow adequate time for a green-field design if required. We also discussed the WaveBoard digitizer design. It makes use of COTS ADC and FPGAs, but can be easily adapted for the specific readout requirements.

- Timing: We discussed timing distribution strategies and requirements.

- Software: We discussed adaptations of existing analysis frameworks like JANA and SRO oriented frameworks like CLARA or TRIDAS. It is clear that the field would like to have an integrated framework from online DAQ to analysis to facilitate the convergence of these two. The approach needs to be scalable, must be able to make use of general purpose CPU clusters but also be able to integrate special-purpose accelerator nodes. Meanwhile, the sPHENIX collaboration has demonstrated the analysis of streaming-mode data from a TPC prototype in a test beam in both their online monitoring package `pmonitor`, as well as in the offline `Fun4All` package.

- Transport protocols: There are several possibly suitable network transport protocol libraries and technologies available on the market, and the community would like to avoid locking into one of these. ADIOS promises to be a meta-library with plugable underlying transport engines. However, performance tests are underwhelming, and the abstraction is leaky, with many library features' availability depending on the chosen protocol. While ADIOS does not seem to be a good solution, we feel that the goal of some sort of transport abstraction should be followed further.

- Machine Learning: A ML approach can outperform classical filtering algorithms. We discussed the possibility to implement such algorithms in FPGAs, as well as the challenges in verifying such algorithms.

- Mid-range opportunities: planned experiments in the mid-term such as Moeller and SOLID at JLab and SPHENIX at BNL, could provide the ideal test bench for SRO solutions. The complexity and the size of these experiments are close to the EIC providing a valuable test in realistic conditions of different SRO schemes.

We had a strong engagement with the Yellow Report progress, with many of us regularly attending and contributing to the discussions inside the YR DAQ meetings. The YR DAQ sub-meetings invited the broad community to participate and was not restricted to streaming readout. Nevertheless, nobody favored, or considered as useful, a triggered readout design for EIC. We view this as a strong endorsement for our approach.

The advisory committee asked us: "..as the Yellow report comes into focus later this year the group should be using one or more of the emerging sub-detectors as a test case to carry through the detailed cost comparison." We could not fulfill the charge at this time. The YR progress has yet to settle on sub-detectors. While we are prepared to work on such a comparison when such a selection becomes more clear, we fear that such a cost estimate will be perceived as biased, just from the fact that the EIC community-wide call of the YR committee did not identify anybody who doesn't prefer an SRO design.

To facilitate the exchange of information with the YR committees, but also as a central repository for our collected results, a overleaf repo was created: `https://www.overleaf.com/7341518614pkrtqxwvgjzg` . We hope that this can become the central repository for the collective work on streaming readout.

From the discussion at the workshop and YR meetings, some principal guidelines can be established:

- The DAQ system should be as physics-agnostic as possible to maximize the physics output.
- For risk-mitigation, data-mining, etc., it is beneficial to keep as much of the raw data as possible. Current estimates indicate that the complete physics rate can indeed be archived, even with technologies available now. However, the final physics rate depends strongly on the detectors selected. It is therefore possible that some reduction is necessary, for example via cross-detector noise reduction, range-of-interest or time-of-interest data selection.
- It is however likely that a full reduction to physics objects online is not required.
- It is likely that some detector types will need running calibration which cannot be done online.
- At the same time, high-level online or near-online analysis will allow an unprecedented quality assurance during data taking. Minimizing the work required to find, debug and repair problems in the data will considerably improve time-to-publication.
- The actual noise and background levels cannot be predicted with 100% reliability, and might be considerably worse in the bring-up of the machine.
- The system needs to flexible and allow for later upgrades and additions of detectors.

- It is beneficial to minimize the required processing power in the counting room, as its availability would be crucial and costly. The system should allow to make use of on-site computing infrastructure as much as possible.

Based on this, it is clear that the DAQ system must be flexible enough to allow for, but not require, online data reduction via filter/data-selection elements. Compared to other large-scale SRO implementations, the community will have only limited experience with the final detectors, so it is critical that the DAQ system can support the initial detector start-up. We have developed a possible start-up scenario for the readout system, to minimize overall risk:

- In the beginning, the detector performance and noise levels might be sub-optimal. To allow for a simple start-up, we envision a system mimicking a traditional triggered system: A data selection signal is generated from some subset of fast detectors, *e.g.* from a calorimeter. This low-latency signal can then be used to drop "uninteresting" data at an early stage in the chain. If built with current technology, such a signal could be used, for example, on the Felix readout boards to mark time segments to forward to the host computer, reducing the bandwidth requirement downstream. Depending on overall detector rate and technology advances in the next few years, such a filter could also be implemented in software. Such a system is comparatively easy to implement—the requirements are somewhat relaxed compared to the equivalent triggered readout, and would require minimal extra hardware.

- At least in the initial phase, as much raw data as possible will be saved, and no higher-order filtering will be performed until enough experience with the detector system could be established. It is likely that the achieved luminosity will be less than the design luminosity in any case, allowing for such a reduced filtering even if a more sophisticated filtering is required at the design luminosity.

- If required, step by step, more and more filters can be activated to reduce the data rate to long-term storage, and more and more high-level online monitoring can be established.

In an ideal scenario, the transition from step to step can happen quickly. However, such a design mitigates the risk that sub-optimal machine or detector performance inhibits productive data taking and eases the initial start-up by building up the complexity over time.

Beside the workshop and YR activity, the following group activities, not directly supported by eRD23 funds, were planned. We give here a progress report for completeness:

- sPHENIX (BNL, SBU): sPHENIX is producing the fifth version of the SAMPA ASIC for the streaming readout of the time projection chamber. This version supports 80 ns shaping time and is optimized for 20 MHz digitization, which is also better suited for EIC application comparing to the early version of SAMPAs. sPHENIX has recently completed the SAMPAv5 production readiness review (`https://indico.bnl.gov/event/8571/`). The initial engineering-run chip has been delivered to BNL and 25 wafer production run has started on May 5, 2020. The team is aiming for testing the new chip with the FELIX readout after the CoVID lockdown. In larger scope, as the sPHENIX construction has started, we will continue to gain experience on constructing a large detector system with three streaming readout detectors that constitutes the full tracking system.

- TPEX/DESY test beam (CUA, INFN/JLAB, MIT, SBU): The test beam at DESY, originally scheduled for April 27 - May 10, 2020, was postponed because of the pandemic. DESY has begun operation in June with restrictions but so far we have not be granted any test beam slots. It is of course unclear how soon we would be willing to travel to Germany in any case. In the meantime, we have focused on the analysis of data taken in the September, 2019, test beam period to compare performance of triggered and streaming readout. Since international travel is likely affected for a longer time, we are investigating the possibility of additional test beam time at JLAB, ideally with three parallel readouts (commercial triggered electrons, commercial SRO capable electronics and INFN waveboard 2.0) and the scintillating glass prototypes of eRD1. It would be simpler to drive the 5×5 lead tungstate calorimeter from MIT together with the necessary manpower to JLab than shipping equipment and people to Germany in these uncertain times.

- VHDL TDC (SBU): The undergraduate project of Marisa Petrusky to build a streaming-capable TDC using FPGAs from the ZYNQ family made considerable progress, and Marisa is likely to continue the work in the next semester.

- Calorimeter (FT-Cal) test (JLab): SRO of two almost full crates of fADC250 modules in a beam experiment. The test proved the transition of our VXS modules from VME2eSST readout, to a true serial stream of data through the VXS Trigger-Processor. The SRO software was a modified version of TRIDAS, the performance was good enough to be a proof-of-concept. Data analysis is still in progress off-line. Selection and reconstruction algorithms will be used in a new test planned for August. Development of the new CODA revision to handle the streaming mode is a work in progress, but there is strong effort on this with the restructuring of personnel in the newly formed EPSCI group. Furthermore, procurement of an Arista network switch, based on a large capacity XILINX UltraScale FPGA, is progressing despite the COVID shutdown.

- FADC250 in streaming mode (JLAB): The FADC250 and VXS Trigger Processor switch module firmware has been updated to allow the 4ns timing and pulse charge information to be transferred to the VTP module in streaming mode. This new firmware has been tested with 16 FADC250 modules (full VXS crate) and further studies are a work in progress to stream raw ADC data in a zero suppressed fashion. The zero suppression scheme will only send the samples within a predetermined window if a signal passes over a programmed threshold. This zero suppression mode will be refined and tested to match the back-plane bandwidth of 20 Gbps between the FADC250 modules and the VTP Refinements to the VXS Trigger Processor firmware that manages the TCI/IP stack has improved the full utilization of the 10 Gbps output links where four links make up the 40 GbE output path from the VTP.

- EMCAL (CUA, INFN/JLAB): In collaboration with JLAB team, the INFN group is instrumenting a PbWO4 electromagnetic calorimeter prototype in streaming mode and test it with a secondary electron beam available in Hall-D downstream of the Pair Spectrometer. The TRIDAS software has already been installed on INDRA Lab and experimental counting house's machines. The WB v2.0 will be used as bench test for the whole system. When the DAQ will be fully working, we will instrument a crate containing JLAB fADC250. Via the VTP board, data will be streamed to the TRIDAS software, filtered by a CPU-based L2 trigger and recorded on tapes. Unbiased data will be also collected for further checks. We are planning to test some of the exclusive features offered by SRO in an on-beam test: on-line calibration, real-time channel selection, on-line data analysis, etc. Results will be reported at the fall meeting. When the system will be fully tested and understood, the lead tungstate crystals will be replaced by CUA scintillating glasses read by SiPMs.

- Timing module (BNL): The extension of the 2020 RHIC operations through at least August has provided us with a new opportunity to test a timing system prototype under realistic conditions with the actual RHIC RF system clock source. One purpose is to test if the clocks can stay locked during the RF frequency changes that occur during injection and acceleration. It will also allow us to measure the jitter, and find the best point for an experiment timing system to tap into the accelerator RF system. As BNL re-opens, those tests are planned in the July/August time frame. So far, we have agreed on the way to access the RF system close to the main control room, and firmware has been developed to accommodate the RF signals.

- RFSoC (BNL): Since the last report, the team has continued the work on the RF System-on-Chip (RFSoC) single-chip Gsps-ADC+FPGA+ARM platforms (e.g. Xilinx ZCU111, ZCU216), in joint R&D with Packed Ultra-wideband Mapping Array (PUMA)[1]. We continue to test the possibility of using single Gsps-ADC to readout multiple detector channels via analog-multiplexer/digital-demultiplexer in the frequency domain, aiming the application of high channel-density readout and cable reduction in the EIC settings.

- Collaboration with Alphacore: Again we plan to collaborate with the ASIC design firm Alphacore to test their latest ADC and pre-amplifier boards with a view to providing a range of useful specifications for test and development front-end electronics. However, the delay in receiving funding from DOE and the lack of electronic expertise at MIT currently has put this on hold. We hope that when things open up again that we can hire technical personnel to work in this area.

How much of your FY20 funding could not be spent due to pandemic related closing of facilities?
Since the allocated funds were supposed to fund travel, we did not spend any money. We expect that we will spend some for the upcoming workshop (technically FY 21, but funds of

[1]arXiv:1907.12559 [astro-ph.IM]

FY 21 are often not available by the time travel has to be arranged), as well as postponed travel for beam times etc.

Do you have running costs that are needed even if R&D efforts have paused?
No.

2 Future

What is planned for the next funding cycle and beyond? How, if at all, is this planning different from the original plan? What are critical issues?

We plan to continue our participation in the Yellow Report progress. We believe that the YR will strongly support a streaming readout based solution for EIC. From the discussions and presentations of the SRO solutions of the different labs, it has become obvious that the overall strategies are very similar. Detail choices like data structure layout etc. are often arbitrary, but can hinder interoperability to a large degree. A standardization of these details is possible in a rather straightforward manner, and the next goal of our group in the coming month. This standardization process will be the focus of our next meeting, tentatively scheduled for early November.

To test the real-world usability of these standards, it is important to test them in real-world settings. For the next couple of years, there will be a series of test-experiments in which small-scale streaming readout solutions, based on these protocols, will be tested. Crucial at-scale demonstrations of SRO will be realized with sPHENIX and the SRO conversion of CLAS12, for which development is already progressing. The experience from the test-beams and the design-process for sPHENIX and CLAS12 will be crucial to revise and improve the standards so that they are optimal for EIC and beyond.

Further projects in SRO are listed below. They are not directly funded by eRD23, but some funds might be used for travel in support of these activities:

- TPEX test-beams at JLAB and DESY.
- CLAS12 SRO test-beams: we have plan to extend tests performed on FT-Cal to the whole Forward Tagger that include a hodoscope, for particle Id, and a tracker for precise determination of the scattered electron angles. The CLAS12 Forward Tagger represents the perfect bench test for the SRO. This will be the first step toward the SRO implementation for the full CLAS12 detector. Replacement of front end components not compatible with a SRO operation already started (old CAEN TDC are being replaced by VETROCs TDC) and has been planned for the next couple of years. Data collected on-beam will be used to test the full readout chain and the data selection.
- Further development of FPGA-based TDC (SBU undergraduate project).
- Other opportunities for streaming-readout tests as they come up.

3 Funding request

We anticipate that in-person meetings and test beam efforts will be able to resume and ramp up in frequency. We request funds to support travel for: informal meetings with various subdetector groups, our bi-annual workshops, and to travel to test-beams. These funds will not only support travel of the project members, but also of graduate and undergraduate students[2] to introduce them to the field of DAQ and help them develop crucial electronics skills. Since the money from FY19 could not be spent, we request a total of $20,000 for FY21. An 80% or 50% scenario would scale down this sum accordingly.

4 Additional information:

4.1 Manpower

Include a list of the existing manpower and what approximate fraction each has spent on the project. If students and/or postdocs were funded through the R&D, please state where they were located, what fraction of their time they spend on EIC R&D, and who supervised their work.

All personnel is currently funded by external sources. We report here time spend on SRO related activities, whether directly EIC-related or not.

- SBU/RBRC: Jan C. Bernauer (Assistant Professor) and Ethan Cline (PostDoc) spend about 30% of their time on SRO-related projects. Further, the undergraduate student Marisa Petrusky will likely return to the group for the Fall 2020 semester.

- INFN/CUA: the personnel involved in the aforementioned activities at INFN is: Marco Battaglieri (senior staff scientist), Andrea Celentano (staff scientist), Luca Marsicano (post-doc), Mariangela Bondi (post-doc), Simone Vallarino (master thesis student) and Paolo Musico (senior staff engineer) in Genova, F.Ameli (senior staff scientist) in Rome and T. Chiarusi (senior staff scientist) in Bologna. Each of us has spent approximately 30% of the time on this activity, partially shared on synergistic activities, in particular the BDX experiment at Jefferson Laboratory.

- BNL: J. Huang, M. Purschke will commit 10-20% time developing the SRO system for EIC, which will be supported under BNL LDRD 19-028 and in synergy with on-going work on sPHENIX SRO tracking system. This work will be supported by an experienced engineering team at BNL including J. Kuczewski, J. Mead, and A. Dellapenna and in collaboration with the ATLAS DAQ team at BNL.

- JLAB: Streaming readout work is performed by staff scientists and engineers from the JLAB DAQ and Electronics groups. G. Hayes and C. Cuevas are the respective group leaders with B. Raydo, E. Jastrzembski and J. Gu working 10% of their time on various streaming readout activities. Eric Pooser, a postdoc in Hall-A worked with us in the last year but left for a position in industry. Hall-B and Hall-D DAQ experts S. Boiarinov and A. Somov joined the team. The streaming readout development activities are an ongoing extension of the hardware developed and implemented for the 12GeV experimental programs at JLAB. These activities are collaborative with the SRO consortium groups and complementary work to implement real-time calibration and analysis is also supported by JLAB LDRD-2014.

[2] E.g. the SBU group will work with four undergraduate students in the Fall semester.

- MIT: Main effort will be through D. Hasell, R. Milner, I. Friščić, S. Lee, P. Moran, and B. Johnston working on the DESY test beam. I. Friščić will assume a shared role at JLAB and MIT working on streaming readout. C. Fanelli will work together INFN and JLAB teams to implement high level JANA-based algorithms within the TRIDAS framework.

4.2 External Funding

Our meeting SRO V was funded and organized by the RBRC. SRO VI was online and organized by JLAB. For SRO VII, we submitted a proposal to organize it as part of the CFNS Stony Brook workshop series. 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.

4.3 Publications

Please provide a list of publications coming out of the R&D effort.

N/A