Development of the EIC Streaming Readout – Guidelines and Specifications

Fernando Barbosa, David Abbott

2 September 2021

Streaming readout has been established as the desired mode for data processing, collection and analysis at the EIC. The consensus from the EIC user community on streaming, i.e. triggerless, readout is summarized within the Yellow Report.

1. Readout Architecture

The EIC streaming readout architecture is shown in fig. 1.

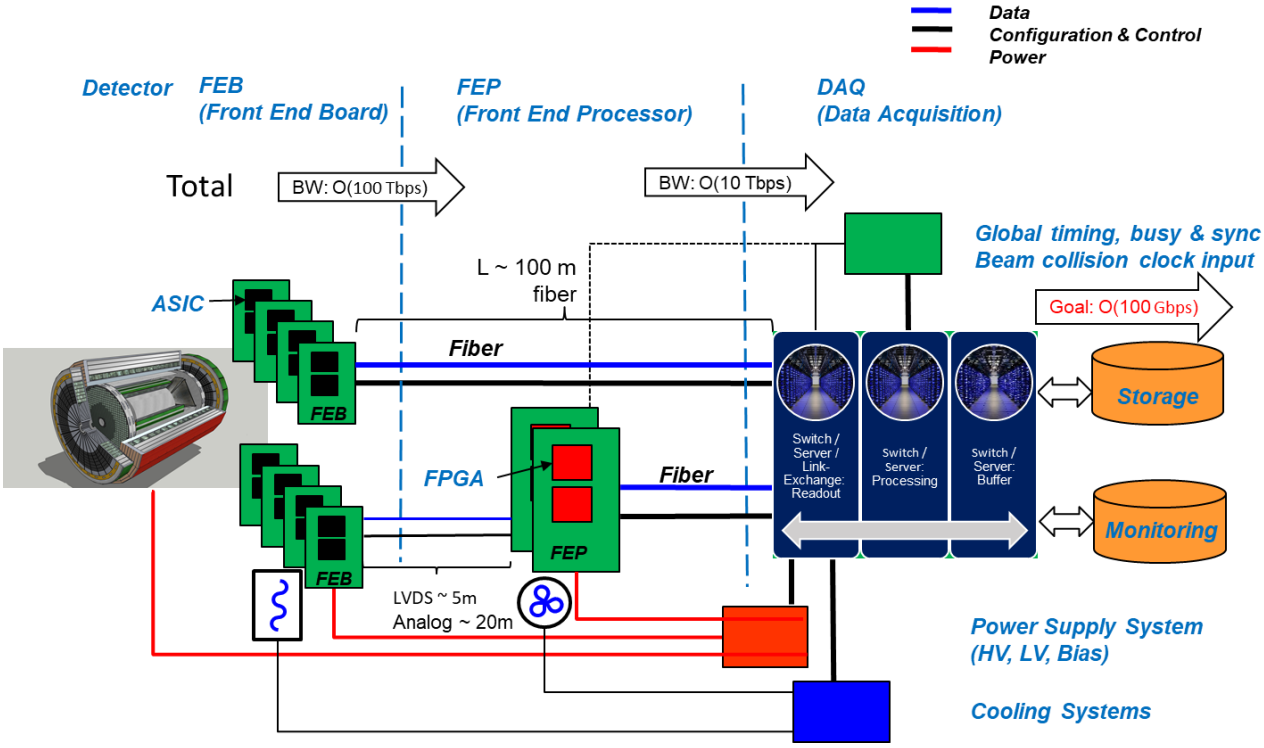

Fig. 1 – EIC Streaming Readout Architecture

The readout chain is partitioned into three distinct functional profiles: FEB, FEP and DAQ.

FEB (Front End Board)

The frontend boards are specifically designed for, mounted on or in close proximity to each of the sub-detectors. These boards conform to the geometry requirements of each of the sub-

detectors and may be further constrained by power dissipation or heat loads, cooling services, radiation levels, cabling requirements and serviceability.

The FEBs are characterized by the use of low noise, low power circuitry with analog frontends and digitization, processing and drive capabilities. This high density mixed-mode circuitry is normally implemented with Application Specific Integrated Circuits (ASIC) when a very large number of readout channels justifies development efforts; alternatively, discrete implementations with commercial-off-the-shelf (COTS) components may be developed, if permitted by the established design boundaries.

The FEBs will likely have configuration, control and timing interfaces, in addition to the data readout drive. It is desirable to implement data transport interfaces via optical fibers to minimize cabling, especially from high granularity sub-detectors, improve throughput and extend range. Although multi-mode fibers are generally acceptable, high precision timing interfaces will likely require the use of single-mode optical fibers.

Heat management for each of the sub-detectors is constrained by a variety of factors such as detector thermal requirements, granularity or density of circuitry, space availability and power dissipation. In some cases, a cooling system may already be necessary and available for temperature stability of the sensor elements, such as SiPMs; in other cases, a cooling system may be necessary for strictly managing the heat dissipated by the readout circuitry.

FEP (Front End Processor)

The FEPs are located outside of the proper sub-detector volume, or in close proximity, and interface to FEBs. The FEPs aggregate the data output streams from multiple FEBs and are designed to handle different types of FEBs, i.e., a few FEP designs with generic interfaces will handle multiple sub-detectors. Some sub-detectors may not necessitate the use of FEPs, such as MAPS, as these may interface directly to switches, servers or link exchange modules (e.g., FELIX).

The FEPs make extensive use of FPGAs for data processing, providing an opportunity to decrease the available output data bandwidth requirements by a factor of ten (10). This can be accomplished by various methods or algorithms, such as zero suppression or via ML and AI filtering. High performance FPGAs, available as COTS, provide the processing power, speed and flexibility of use. The choice of FPGAs and their location within the experimental area must be carefully considered with regards to radiation levels.

As with the design considerations for the FEBs, the FEPs will play a critical role in configuration, control and timing of the various sections of the readout chain via optical fibers. For their location and serviceability, heat management will likely be accomplished by forced convection.

Switch/Server/Link Exchange

The Front-End Link eXchange (FELIX) boards, servers and network switches will be located far from the detector, facilitated by the extensive use of optical fibers. Servers and switches are available as COTS.

The FELIX board, originally developed for the ATLAS experiment, interfaces the DAQ and the detector front-end and functions as a router between custom serial links and a commodity switching network using standard technologies. FELIX is designed to be detector independent, supports the CERN GBT protocol to connect to front-end units, supports distribution of timing, trigger and control and supports calibration operations. It is expected that an updated design will be necessitated to fully benefit the EIC streaming readout. Optical fibers from the FEB and FEP are employed for data transport.

Rates and Timing

With the implementation of a fully streaming readout model it is important to provide enough available bandwidth at the front-end (FEB) to handle potentially highly variable rates (signal/background/noise) from all detector systems. The total anticipated bandwidth limits for the whole detector are shown to be on the order 100 Tbps at the front-end. With the implementation of the FEP stage we assume limits can be decreased by a factor of ten to 10 Tbps at the server/switch/link exchange level. Given current estimates on the physics interaction rates, following back-end processing and buffering, the expected data output rates to more permanent storage is anticipated to be around 100 Gbps.

Timing, generally consists of global timing, busy, synchronization and clock distribution within the experimental area. Precision timing will require precision timing referencing the beam or beam crossings, as well as multiple control loops addressing stability and drift. A single accelerator clock reference may exhibit timing jitter of a few hundred femtoseconds, implemented with single-mode fibers and periodic phase correction.

Clock distribution jitter via multi-mode optical fibers, clock de-skewing and distribution via backplanes and employing dPECL technology and using COTS components can result in clock jitter of about 4 ps, which is sufficient for detectors requiring timing resolution in the 100 ps range. As some EIC timing detectors are specified for 20 ps timing resolution, it is expected that clock jitter will need to be better than 1 ps. The use of single-mode optical fibers, careful attention to design and component and technology selection, as well as consideration for feedback control loops and calibration should result in clock timing jitter better than 1 ps.

Consideration for calibration, testing and timing of the various sections and components of the EIC streaming readout requires that every part of the readout chain must be designed for the intended purpose of streaming readout and also for a triggered implementation. That is, the readout system should be able to operate with and without a trigger.

Cooling System

Cooling systems will be required to manage the additional heat from the power dissipation of the electronics, to cool a particular set of sensors for lower noise operation, for temperature stability or in combination. Most hardware is expected to be COTS; heat exchangers may be custom made to suit a sub-detector due to geometry, for example; controls are generally offered by the original manufacturer or can be implemented with PLCs. Liquid cooling may be better applied to FEBs, also due to geometry constraints of the sub-detectors; forced convection will suffice for cooling of the FEPs and servers.

Power Supplies

The power supplies supporting the experimental equipment will be COTS. These will be modular, multi-channel units with chassis, and as normally offered within this field. Various modules will supply HV, LV or Bias: HV: 1 kV ~ 8 kV, a few mA and used with PMTs and MPGDs; LV: up to 15 V and up to 10 A per channel; Bias: ~ 50 V, a few mA for SiPMs.

These power supplies are floating supplies, which will be referenced to the detector clean ground grid in the experimental area. The LV and Bias supplies can be configured for either polarity. These supplies include global and single channel controls and forced convection cooling.

Cabling

Optical fibers offer considerable advantages over copper wiring and are, therefore, the preferred choice for data transport. Optical fibers are available with multiple fiber links in single and multi-mode varieties. Copper wiring may be necessitated, however, in a few applications where signal drive capabilities are appropriate for the transmission length.

2. ASICs for Front-End Readout

The number of channels specified for each sub-detector function and type is shown in table 1, which was obtained from the *EIC Readout Specifications Matrix* spreadsheet. This information was collected from the original EIC proposal and from recent assessments and it is subject to change.

Table 1 – EIC channel count

Detector	Sub-system	Type	Sub-type	Channels
Tracking				
	Silicon maps	Si MAPS	Pixel	200M
	TPC	MPGD	Pads	160K
	GEM	MPGD	Strips	217K
	uRWELL	MPGD	Strips/Pads	
	Cylindrical Micromegas	MPGD	Strips	60K - 80K
	sTGC	MPGD	Pad, Strip, Wire	
	TRD	MPGD	Strips/Pads	
Calorimetry				
	n-EMCal PWO	Cal	SiPM	1628
	n-EMCal SC Glass	Cal	SiPM	1168
	n-Hcal (KLM type) (10 layers)	Cal	SiPM	10k
	p-EMCal	Cal	SiPM	31k
	p-Hcal	Cal	SiPM	3/6/9k
	b-Hcal (KLM Type) (5 layers)	Cal	SiPM	24k
	b-Ecal (ScFI part)	Cal	SiPM	4k
	b-Ecal (Si layers)	Cal	Si Sensor	480M
PID				
	mRICH @ e-endcap	RICH	MCP-PMT (LAPPD) or SiPM	350k
	dRICH @ h-endcap	RICH	MaPMT	330k
	GEM RICH	MPGD	Strips	220k
	hpDIRC @barrel	DIRC	MCP-PMT	100k
	psTOF @barrel	TOF	LAPPD or LGAS	
	LGAD TOF	TOF	PMT/SiPM	
	LAPPD/MCP-PMT TOF	TOF	PMT/SiPM	
Far Forward Detectors				
	ZDC		PMT/SiPM	225(EMCAL) +36 (HCAL)
	Low Q2 tagger		PMT/SiPM	288 (2 EMCALs)
	Low Q2 tagger	Si strip	Si strip	6.4k

	Luminositymonitors		**PMT/SiPM**	**288 (2 EMCALs)**
	Roman Pots/OffM/B0	**Si**	**Si**	**500k/750k/32M+320k**
	Proton Spectrometer		**PMT/SiPM**	
	Lepton Polarimeter	**Calorimeter**	**PMT/SiPM**	**110**
		Strip/position	**Diamond strip**	**1000**
	Hadron Polarimeter		**PMT/SiPM**	

The Si detector readout (in red) is included with the development of the MAPS sensors and it is, therefore, not included with the development of the ASICs here described. Also, the need for other items (in orange) is not fully understood at this time in the context of the EIC needs and are also not included.

The very large number of channels of MPGDs and photon sensors justify the development of two (2) types of ASICs. Although SiPMs and MCP-PMT/MaPMT may be further segmented into two distinct developments, it is reasonable to conceive of an ASIC that is possibly applicable to both. Furthermore, a non-ASIC solution for the photon sensors may also be considered if allowed by careful consideration of the specifications and sub-detector constraints. The following specifications are, therefore, preliminary and these will need to be further refined based on an accurate accounting of the number of channels, which may justify further developments, and final sub-detector requirements.

2.A Front-End ASIC Specifications for MPGDs - Preliminary

Detector

Capacitance	<200 pF nominal (500 pF maximum).
Noise	<3000 e^- @ 100 pF
Charge	25 fC – 100 fC (1 pC maximum).
Gain	$5x10^3 – 2x10^4$
Signal Time	100 ns – 500 ns (10 us ion drift time maximum), multiple hits per channel.
Signal Range	$<10^6$ e^-
Rates	<2 kHz per channel.

Readout

Attributes	Amplification, digitization and buffering.
Features	Amplitude and time per hit; waveform samples for testing and calibration functions. Zero suppression; triggerless and triggered operation.
# Channels	64
Input Impedance	<70 Ohm
Gain	2 mV/fC – 30 mV/fC, configurable.

Peaking Time	40 ns – 250 ns shaping, configurable.
Crosstalk	<1 %
ADC Resolution	12 bit (>10 bit ENOB)
TDC Resolution	<20 ns
Sampling Rate	>80 MSPS
Optional	Discriminators and scalers are desirable.
Triggering	Streaming (triggerless) readout is the default mode. Triggered operation is required for testing and calibration functions.
Pulsing	Channel group pulsing desirable for testing function.
Output	TBD. Data format to be determined and to be consistent with optical fiber data transport between FEBs and FEPs.
Control Interface	TBD. Slow controls and configuration interface to be consistent with optical fiber data transport between FEBs and FEPs.
Technology Node	65 nm CMOS or higher.
Packaging	BGA or other SMT industry standard packages.
Power	1 W ± 0.25 W or < 20 mW per channel.
Supply	<+3 V DC

2.B Front-End ASIC Specifications for SiPM/MCP-PMT/MaPMT - Preliminary

Detector

Capacitance	60 pF – 5 nF (depending of cell size and grouping), <30 pF for PMTs.
Noise	1 p.e. @ <100 kHz, nominal (extends to 3 p.e.), lower for PMTs.
Gain	$<10^6$
Signal Time	3 ns – 80 ns
Rise Time	1 ns – 3 ns
Signal Range	<1 V into 50 Ohm or $<10^5$ pixels.
Rates	<50 kHz per channel.
Bias	Vop ~ 50 V DC for SiPMs.

Readout

Attributes	Waveform sampling. Amplification, signal conditioning, digitization and buffering.
Features	Amplitude and time per hit. Zero suppression; triggerless and triggered operation. Input offset voltage adjustment.
# Channels	64
Input Impedance	<50 Ohm, depends on configuration.
Gain	1 - 10, configurable.
Peaking Time	<40 ns for SiPMs, ~1 ns for PMTs.
Crosstalk	<1 %

ADC Resolution	10 - 14 bit
TDC Resolution	1 ns for SiPMs, <100 ps for PMTs.
Sampling Rate	>80 MSPS for SiPMs, >1 GSPS for PMTs.
Optional	Discriminators and scalers are desirable.
Triggering	Streaming (triggerless) readout is the default mode. Triggered operation is required for testing and calibration functions.
Pulsing	Channel group pulsing desirable for testing function.
Output	TBD. Data format to be determined and to be consistent with optical fiber data transport between FEBs and FEPs.
Control Interface	TBD. Slow controls and configuration interface to be consistent with optical fiber data transport between FEBs and FEPs.
Technology Node	65 nm CMOS or higher.
Packaging	BGA or other SMT industry standard packages.
Power	1 W ± 0.25 W or < 20 mW per channel.
Supply	<+3 V DC

A data formatting convention is desirable for development and consistency of interfacing and serviceability of the various readout sections, including the design and development of ASICs, FEBs, FEPs and the DAQ. Such a convention, however desirable, may not be unique due to various constraints placed on the design of the front-end so that different FEBs may have different output configurations. The standardization of output data configuration from FEPs, however, is enabled by the flexibility afforded by the use of FPGAs. It is expected that a single convention will be developed for the interfacing to the DAQ (Switch/Server/Link exchange).

3. DAQ

While the general consensus amongst the nuclear physics community is that streaming (triggerless) readout should be the de-facto standard for the EIC detector(s) data acquisition, there are a number of questions that need to be answered, and standards defined, before we can begin to implement a practical, flexible, efficient and scalable streaming DAQ system.

One of the first questions to get a handle on is defining the boundary between what we are requiring of the “online” DAQ system and the “offline” processing. Historically in the traditional triggered DAQ, events were defined by a trigger, distributed to all the front-end detector electronics. The data are collected, and events built from data fragments. These events, may be filtered in some way by an online processing farm, but ultimately what is left are written to files. This represents the end of the Online DAQ responsibilities.

In the EIC streaming model, there will be many independent streams of data coming off the detector electronics (FEB). These streams will get aggregated at some level either by the FEPs or in the FELIX Boards (or both). This new set of aggregated streams is what will be made available to “back-end” processing. The expectation is that all the stream processing in the back

end will be done on COTS based networks, servers, and other high performance computing hardware (GPUs etc...). The scale of this infrastructure is very much dependent on both the aggregate bandwidth of the streams and the level of processing required to get the aggregate data set to some pseudo-permanent storage.

Another consideration is that all the back-end processing does not necessarily need to be resident in the counting house. However, it provides a reasonable compromise to require the resources necessary to complete the online processing to be local. Certainly the FELIX Servers will need to be resident in the counting house. Providing enough network, processing and storage infrastructure locally to complete a first tier analysis would allow a level of autonomy and efficiency for the DAQ, particularly if it is focused on data processing related to items like event identification, data compression or detector related background and noise filtering to reduce the overall size of the data set. Second tier processing would probably best be handled in a centralized data center where resources could be reserved based on running experiments.

Perhaps the best place to draw from for guidance on developing a streaming DAQ system for EIC is to look at the current efforts ongoing at both BNL and JLab. Both labs have active programs for evolving their systems to support streaming readout. At BNL they will use a hybrid DAQ utilizing both streaming and triggered readout for sPHENIX. They will be using current generation FELIX cards as a key element of the DAQ architecture. At JLab the CODA data acquisition system is being updated to support both triggered and streaming readout from a custom FPGA-based board called the VXS Trigger Processor (VTP). This board is currently being used in all the experimental halls. A revision of this type of board could be envisioned as a template for the EIC FEP boards.

High Resolution Clock Distribution

The backbone for any streaming DAQ is a global high resolution clock distribution system. Both BNL and JLab use existing custom hardware implementations for their respective DAQ systems. Without getting into details on one particular design, it is more important to consider the requirements of such a system. The system clock should be delivered to all front-end components including the FEBs, FEPs and the FELIX boards. As mentioned in an earlier section, in order to realize anticipated time resolution measurements down to 20ps, the clock will have to be delivered to the FEB electronics with a jitter on the order of 1ps. In addition, the system should be able to generate a synchronized reset to all the endpoints so that all the local timestamp counters will be aligned to facilitate the management and aggregation of the resultant streams.

While the system clock source can be internally generated and controlled, the expectation for EIC is that it will be based on the collider's bunch crossing frequency (~ 100MHz). Hence the system clock should be able to run on an external source as well.

Design of the FEB electronics must allow for the implementation of an external system clock such that the data that are generated by the ASICs or digitizers can be timestamped accordingly. The data must also be able to be "framed" in blocks corresponding to a defined time slice window. This window will represent a specific number of clock ticks and would be configurable by the DAQ system (nominally in JLab streaming systems now this frame size corresponds to around 65µs).

In order to support a hybrid DAQ system where both streaming and triggered readout can be implemented, the clock distribution should also be the system to allow for the input of external triggers that can be synchronized with the system clock and distributed to all the front-end electronics in a programmable but deterministic way. This will facilitate certain DAQ operational modes (e.g. pulser-based calibration runs) for any subset of detectors.

FPGA-Based Stream Management

Data Streams being generated on the FEB need to be driven in a deterministic way, and they must be synchronized to the global clock. Depending on the specific capabilities of the ASICs it may be necessary to provide some complimentary processing resources at the front-end to support the data framing as well as some initial zero-suppression or threshold filtering of the data. To the degree that one can minimize potential radiation effects, use of FPGA-based processing is the preferred solution here. This gives the User more flexibility in stream management.

For the FEP boards the most flexible solution would be to implement some System on Chip (SoC) option where a full Linux-based OS can be implemented along with the FPGA resources. Depending on the experiment's requirements this would allow Users to combine software-based, network-centric applications for configuration and control of the hardware and firmware resources to best suit operational conditions.

The JLab VTP module supports a SoC ARM-based Linux OS to configure and monitor a Vertex 7 FPGA for stream processing. At BNL, the PCIe accessible FELIX board can be configured and controlled via server based applications. The EIC equivalent systems will be the primary aggregation points for the "raw" detector data streams. Because these are the main aggregation points for the front-end DAQ, there will need to be some well-defined but configurable algorithms for merging streams and managing potential congestion and data loss both for the incoming streams and the outgoing aggregated streams being queued up for back-end processing.

It may be noted that while both JLab and BNL currently have custom hardware solutions for these Aggregation & ReadOut Control points (A-ROC), there are recent COTS-based solutions that could potentially support both the FPGA resources and flexibility necessary to support the EIC DAQ. One example are the FPGA-based programmable switches from Arista (e.g. model 7130).

Configuration and Control of the Front-End

The schematic shown in Fig. 2 is a generalization of the proposed connection between the detector front-end electronics and the aggregation and readout control points (FEP and FELIX). While the individual FEB designs may have several options for configuration and control as well as for data output (e.g. JTAG, USB, ethernet, etc.), it would be most helpful for the production DAQ to standardize on one communications protocol between the FEB and the A-ROC; Specifically one that can utilize a common optical link (possibly multi-fiber) for both configuration and control communications, as well as for the high speed serial data.

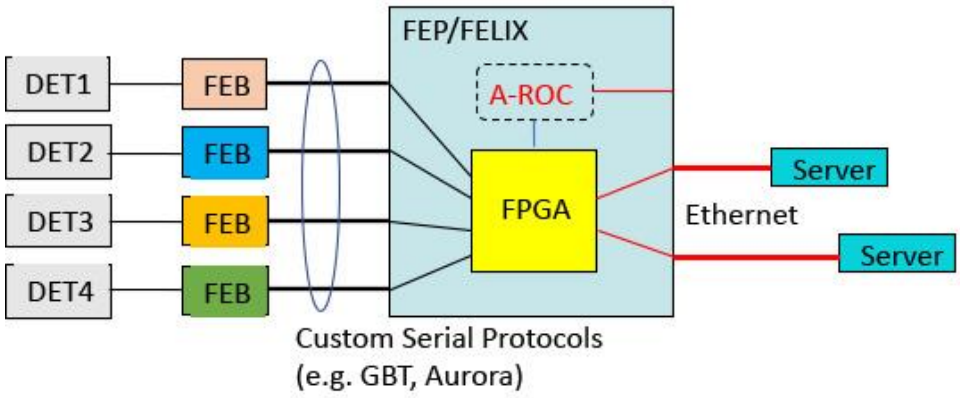

Figure 2 – General schematic of communication with FEB electronics.

The configuration and control communications with FEBs would be handled through local software applications running on the FEP or FELIX server and using libraries designed to use the custom serial protocol. In principle the global clock could be sent to an FEB on the same link. Data frames generated by the FEB could be tagged to identify control communications coming back to the A-ROC.