Progress Report: Streaming readout for EIC detectors

Jan C. Bernauer
for the eRD23 EIC Streaming Readout Consortium

EIC R&D meeting, BNL July 2019

RBRC
RIKEN BNL Research Center

Stony Brook University

Who are we: SRC members

- Catholic University of America: S. Ali, V. Berdnikov, T. Horn, M. Muhoza, I. Pegg, R. Trotta
- INFN Genova: **M. Battaglieri,** A. Celentano
- Stony Brook University / RBRC: **J. C. Bernauer**
- Massachusetts Institute of Technology: D. Hasell, R. Milner
- Thomas Jefferson National Accelerator Facility: C. Cuevas, M. Diefenthaler, R. Ent, G. Heyes, B. Raydo, R. Yoshida
- New: Brookhaven National Laboratory: M. Purschke, J. Huang

Additionally many regulars

$\longrightarrow$We welcome new members!$\longleftarrow$

What is streaming readout

» Electrically: No global electric signal ("trigger") to control data conversion / data retention

What is streaming readout

- Electrically: No global electric signal ("trigger") to control data conversion / data retention
- Logically: Hits labeled by time stamps, not event numbers. Definition of an event not part of the electronics!

Why triggered?

- Triggered DAQ was a necessity because electronics was too slow
- Trade-off: Rather have dead-time than no experiments
- Modern triggered system reduce dead-time with large efforts
  - Often by building a streaming front end and then adding a trigger module.
- *If you never heard of triggers, you wouldn't build one.*

The bottleneck moved

- First generation bottleneck: Conversion times multiple μs.
  - Fast ADCs killed it.

The bottleneck moved

- First generation bottleneck: Conversion times multiple μs.
  - Fast ADCs killed it.
- Second generation bottleneck: Bandwidth of back-planes
  - VME: 320 MByte/s per crate.
  - Today: 1 GBit/s per module "free", 10 Gbit/s per module cheap, 100 GBit/s per module reasonable.

The bottleneck moved

- First generation bottleneck: Conversion times multiple μs.
  - Fast ADCs killed it.
- Second generation bottleneck: Bandwidth of back-planes
  - VME: 320 MByte/s per crate.
  - Today: 1 GBit/s per module "free", 10 Gbit/s per module cheap, 100 GBit/s per module reasonable.
- Third generation: Long term storage.
  - This drives adoption by HEP.

What is streaming readout II

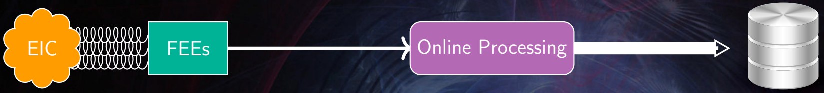

- » Convert all signals continuously
- » Almost always: Per channel zero suppression

What is streaming readout II

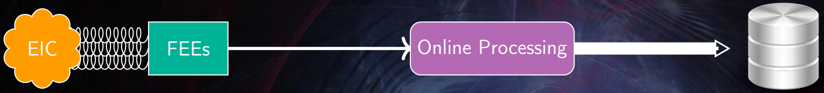

- Convert all signals continuously
- Almost always: Per channel zero suppression
- Maybe: Per detector noise suppression

What is streaming readout II

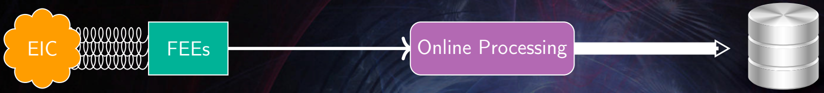

- Convert all signals continuously
- Almost always: Per channel zero suppression
- Maybe: Per detector noise suppression
- Maybe: Per channel / detector feature extraction (total energy, tracks, etc.)

What is streaming readout II

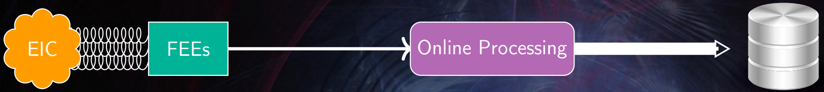

- Convert all signals continuously
- Almost always: Per channel zero suppression
- Maybe: Per detector noise suppression
- Maybe: Per channel / detector feature extraction (total energy, tracks, etc.)
- Maybe: High level feature extraction/physics extraction

What is streaming readout II

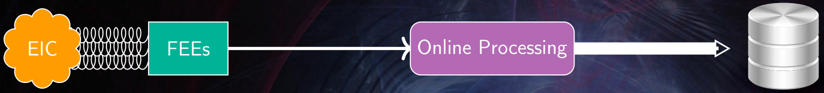

- Convert all signals continuously
- Almost always: Per channel zero suppression
- Maybe: Per detector noise suppression
- Maybe: Per channel / detector feature extraction (total energy, tracks, etc.)
- Maybe: High level feature extraction/physics extraction
- Maybe: Data selection based on extracted features

What is streaming readout II

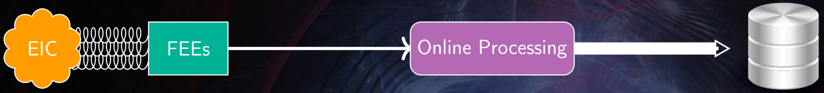

» Convert all signals continuously
» Almost always: Per channel zero suppression
» Maybe: Per detector noise suppression
» Maybe: Per channel / detector feature extraction (total energy, tracks, etc.)
» Maybe: High level feature extraction/physics extraction
» Maybe: Data selection based on extracted features
» Maybe: Remove raw data

What is streaming readout II

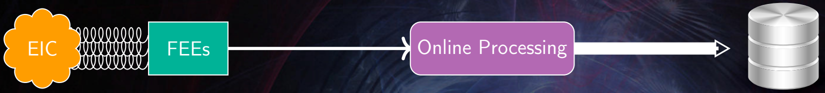

- Convert all signals continuously
- Almost always: Per channel zero suppression
- Maybe: Per detector noise suppression
- Maybe: Per channel / detector feature extraction (total energy, tracks, etc.)
- Maybe: High level feature extraction/physics extraction
- Maybe: Data selection based on extracted features
- Maybe: Remove raw data
- Save data to long term storage

General advantages of a SRO

- Reduce complexity in FEE
  - Remove "trigger" module from already streaming electronics

General advantages of a SRO

- Reduce complexity in FEE
  - Remove "trigger" module from already streaming electronics
- Relax hard timing constraints. Instead of latency, focus only on throughput.
  - Buffer memory in PCs is cheap. Could even use disk (LHCb)!

General advantages of a SRO

- Reduce complexity in FEE
  - Remove "trigger" module from already streaming electronics
- Relax hard timing constraints. Instead of latency, focus only on throughput.
  - Buffer memory in PCs is cheap. Could even use disk (LHCb)!
- Move complexity from hardware to software (and from online to offline)
  - More flexible, and more people can contribute. (Also gain do-overs)

General advantages of a SRO

- Reduce complexity in FEE
  - Remove ”trigger” module from already streaming electronics
- Relax hard timing constraints. Instead of latency, focus only on throughput.
  - Buffer memory in PCs is cheap. Could even use disk (LHCb)!
- Move complexity from hardware to software (and from online to offline)
  - More flexible, and more people can contribute. (Also gain do-overs)
- Remove bottleneck of event building.
  - Can easily scale to large channel counts.
  - Event building always ”brittle”. What happens if FEE dies?

Spectrum of SRO

min ⟷ max

amount of online processing

- Save all data
- Lowest risk
- Maximum physics
- Highest rate

- Only keep high level data
- Highest risk
- Maximum physics/byte

- Generally: Can reach more physics than with a triggered system

Spectrum of SRO

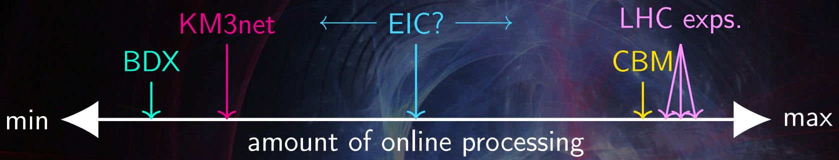

- Save all data
- Lowest risk
- Maximum physics
- Highest rate

- Only keep high level data
- Highest risk
- Maximum physics/byte

- Generally: Can reach more physics than with a triggered system

A rate study: sPHENIX based EIC detector (Jin Huang)

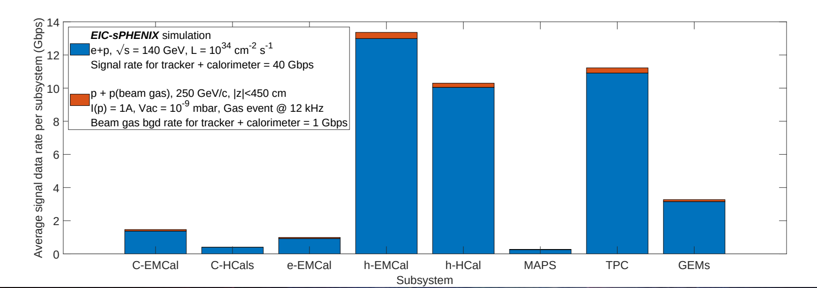

Total rate around 100 GBit/s. This is 1/2 the rate of sPHENIX!

- Rates are very well doable for sPHENIX, will be trivial for EIC time frame.
- EIC SRO is on the safe side of the spectrum!
- BeAST: similar rate
- JLEIC: First estimate: 250 GByte/s for vertex detector. Need streaming, ROI/noise suppression.

(Different detectors might give different rates, especially if channel count is drastically larger.)

Likely configuration for EIC

- Convert all signals continuously
- Per channel zero suppression and feature extraction
- Maybe: Per detector noise suppression
- For online monitoring: Per detector feature extraction (total energy, tracks, etc)
- For online monitoring: High level feature extraction/physics extraction
- Only drop data in zero suppression stage
- Save data to LTS

Possible downsides of SRO

» Reuse of existing, triggered hardware → sPHENIX

Possible downsides of SRO

- » Reuse of existing, triggered hardware → sPHENIX
- » Multiplexer chips to reduce ADC count (APV, DREAM, etc)

Possible downsides of SRO

- Reuse of existing, triggered hardware → sPHENIX
- Multiplexer chips to reduce ADC count (APV, DREAM, etc)
  - EIC minimally biased rate >100kHz,
  - 64:1 Multiplexer (APV:128), 40 MSAMP ADC
  - Single sample / trigger
    → 16% dead time. That's worth $millions per year!
  - If multiplexers are not possible, streaming is likely cheaper
- ASICs which require a trigger → maybe local trigger?
- Less experience → That's what we are here for.

For EIC, streaming readout is *insurance*

- If rates are higher than estimated, can easily compensate (by adding back "maybe steps")
- SRO is simpler, easier to reason about and debug.
- No complex trigger electronics/timing.
- Avoid critical bottle neck of event builder (can do event building over and over again offline).
- Does your trigger capture all interesting events? Also the channels you think of in the future?

Unexpected physics

- » In 2016, surprising evidence of collectivity in p+p collisions [CMS, Phys. Lett. B 765 (2017) 193]
- » Inspired search for QGP signatures in even smaller systems: e+p [ZEUS, QM18], e+e [MOD, arxiv:1906.00489]
- » HERA data limited by statistics and trigger.

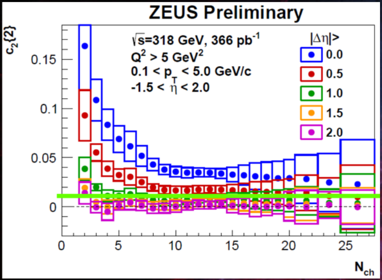

- » $c_2\{2\} < 0.01$ at $N_{ch} = 15$
- » implies $v_2 < 0.1$
- » Room for small v_2, need more statistics!

(Figure / discussion courtesy Austin Batsy, IS2019, ZEUS QM18)

- » Need to record large amounts of high multiplicity events with minimal trigger bias, including low Q^2 and diffractive events.
- » At LHC: streaming/locally triggered front ends, HLT on N_{ch}
- » At EIC: Natural fit for streaming readout

Report from our last workshop (selected topics)

Hardware

- ASIC development at INFN Turin
- BDX digitizer boards
- JLAB: FADC250 streaming conversion via ethernet, clock/sync recovery,...
- BNL: FELIX cards, SAMPA,...

Software

- Real time analysis at LHC
- Design considerations for Mainz new analysis software
- TriDAS software for streaming readout
- Prototype of streaming protocol

Plus many talks about experiments, validation, etc.

Collaboration with the software consortium

Identified overlaps on both ends of the software stack:

- Aim to find general data format for exchange
- MC frameworks organized around time instead of events $\Leftrightarrow$ i.e. what is required for pileup simulations.
- Since analysis and DAQ converges: Problem of orchestration of DAQ/Analysis cluster similar

Proposal for next funding period

We realize the R&D program has very limited funds. We request 20k of travel funds.

- We want to organize mini workshops / visits with detector R&D groups
- Or offset their travel costs to our regular workshops

Our goals are:

- Educate about opportunities / possible pain points of SRO.
- Encourage detector groups to think about detector, readout, feature extraction, and analysis as an integrated system
- Develop concepts for streaming capable readouts for the individual detectors as part of a global EIC solution.
- Connect detector groups and DAQ groups to facilitate use of existing solutions for prototypes / tests.

Next workshop

RBRC
RIKEN BNL Research Center

- Next workshop will be at BNL – hosted by RBRC
- Three days: November 13 to 15
- Website: https://www.bnl.gov/srv2019/

Thank you!

Glossary

Front-end electronics (FEE)
The electronics which interfaces with the detector, typically converting the analog signal from the detector via an analog-to-digital (ADC), charge-to-digital (QDC), or time-to-digital (TDC) converter into the digital domain.

Glossary

Triggered readout

A data acquisition system in which hardware produces an electrical signal according to a trigger criteria based on a subset of detector information available quickly. The signal is used to control the conversion of detector signals into the digital domain, or to trigger the read-out of a data-window from a continuously filled buffer.

Glossary

Second-level / high-level trigger

In triggered systems, higher-level triggers are often used to reduce deadtime (via a fast clear) or data amount (by dropping the so-far recorded data for that event). Each level in such a system typically has different time constraints and complexity limits. For example: a certain time frame could not be forwarded to the tracker if certain conditions are not met. In certain, complex, triggered setups, the later stages can resemble a streaming system, where a stream of events flows through a network of analysis nodes, and data selection criteria either accept or drop the event. The main remaining difference for this part is then that the data is organized and tagged by an event number instead of time stamps.

Glossary

Pipelined/buffered readout

A triggered readout system where event data is stored on the front ends and read out asynchronously by the backend.

Streaming readout

A data acquisition system without an element producing electrical signals to control the conversion into the digital domain or readout of a buffer. Each channel, independently, record data over a certain threshold and stream them to a CPU farm for further elaboration.

Glossary

Zero suppression
Removal of data if close to the no-signal level of the detector. For example, in ADC data, removal of the pedestal.

Noise suppression
Removal of data produced by intrinsic or extrinsic detector noise, for example by correlation with neighboring channels or shape analysis.

Feature extraction
Calculation of higher-level information. E.g. calculation of hit time and energy from ADC samples, or calculation of track information from hits. Often, but not necessarily, accompanied with the removal of the underlying lower-level data.

Glossary

Online Physics analysis
Analysis of the high-level information provided by the feature extraction steps to produce physics-relevant information (e.g. missing mass).

Data selection
In a SRO system, data can be algorithmically selected for further processing and long-term storage. Not selected data is dropped and not further processed. This is equivalent to the function of first and higher-level triggers in triggered systems, but can make use of all detector information and results from further analysis steps including feature extraction and online physics analysis.