

Streaming Readout for EIC Experiments

Douglas Hasell



Detector Workshop, EIC User Group Meeting
Catholic University of America
July 29, 2018

Goal of Streaming Readout

To readout all useful experimental data

Goal of Streaming Readout

To readout all useful experimental data

Selections based on access to all the information from all detectors

Goal of Streaming Readout

To readout all useful experimental data

Selections based on access to all the information from all detectors

Without relying on hardware triggers

Goal of Streaming Readout

To readout all useful experimental data

Selections based on access to all the information from all detectors

Without relying on hardware triggers


Controlled by adaptable software algorithms

NAS Report on EIC Science - 7/24/2018

“The continued rapid pace of technological development is starting to enable a transition from the event-oriented and triggered data-acquisitions of past and current experiments in nuclear and high-energy physics to data models where detector subsystems deliver time-stamped streams of data for processing with increasingly integrated and advanced computing resources in real time.”

-Ernst Sichtermann (LBNL)

Event Rates



Pythia ep rates $\frac{dN}{d\theta d\phi dt}$ vs η

for 10 GeV on 100 GeV

at 1 m from IP

$\mathcal{L} = 10^{33}$

Assume central tracker $-1 < \eta < 1 \rightarrow 3 \times 10^4 \times 9.97 \approx 3 \times 10^5$ tracks/s

T. Ljubicic estimated 9.4 kB / track \rightarrow 2.82 GB/s

No background noise, higher \sqrt{s} may add 30%, scales with \mathcal{L}

LHCb currently writing ~ 1 GB/s

sPHENIX planning on writing 10 GB/s

EIC Experiments Expectations

Expected EIC experimental environment


- luminosity $10^{33} - 10^{35} \text{ cm}^{-2} \cdot \text{s}^{-1}$
- event rates $100 - 1,000 \times$ rates at HERA
- detectors with $\mathcal{O}(10^6)$ channels
- can expect rates comparable to current LHC experiments $\mathcal{O}(1 \text{ TB/s})$

Such an environment will greatly benefit from streaming readout

- leverage advances in electronics, computing, storage, *etc.*
- flexible software controlled analysis and event selection
- on-line data monitoring, calibration, and alignment
- significant feature building and analysis on-line in real time
- drastically reduce event size requiring less storage capacity



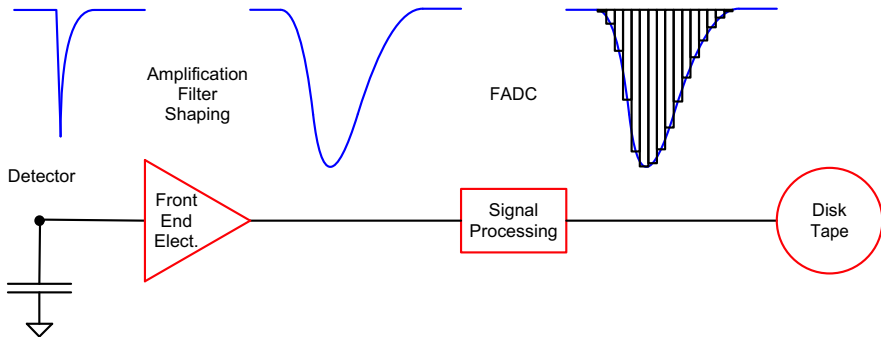
Ideal Readout Scheme - Schematic for One Channel



Consider this simple readout scheme

- detector channel experiences an event and generates a signal
- front-end electronics amplifies, filters, and shapes appropriately
- signal processing extracts time, integrated charge, amplitude, ...
- results are written to disk or tape for offline analysis

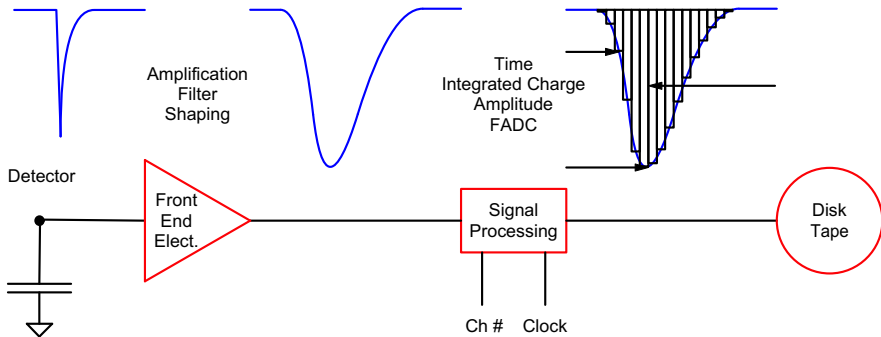
Ideal Readout Scheme - Schematic for One Channel



More modern scheme might use a flash ADC or TDC

- samples the signal at some rate and writes the results
- signal processing possible offline
- time, charge, amplitude, *etc.* parameters extracted as needed
- more information also available: pile-up, signal shape, noise, ...

Ideal Readout Scheme - for More than One Channel




For more than one channel need to add channel ID and time stamp

Scheme provides trigger free, independent readout of all channels

- BUT ! \Rightarrow breaks down with increasing event size \times event rate
- \Rightarrow limited by computer resources, time, manpower, and \$\$\$
- need to discard noise, background, and unwanted events

Traditional Triggered Readout Scheme



Traditional solution is a triggered readout scheme

- signal or signal parameters stored in pipeline / buffer or delayed
- trigger decision based on fast signals from subset of detector channels
- often special detectors, hardware, and electronics involved
- sometimes multiple levels of buffers and triggers required

Trigger logic can veto or select events but ...

- decision based on a subset of the data
- logic based on past experience and expectations
- finds what is expected → might miss the unexpected

Streaming Readout is Possible

Leverage advances and falling costs of electronics, computers, storage, . . .

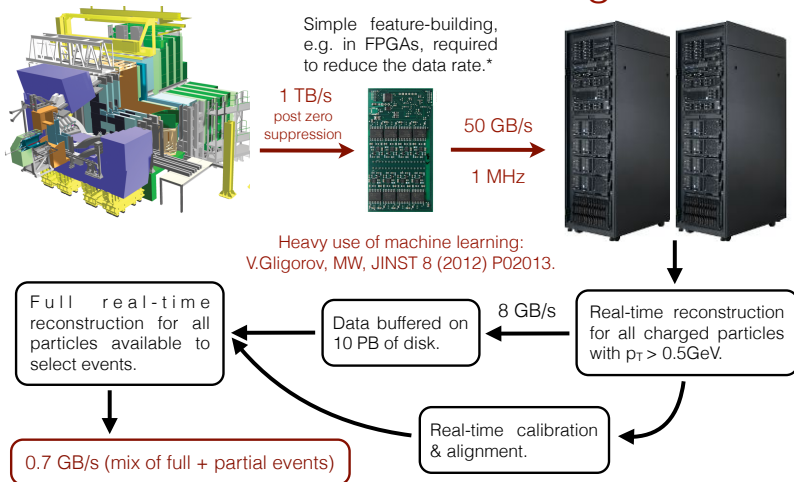
- ASICs → multiplexed ADC/TDC chips, rad. hard, low power, . . .
- FPGAs → affordable, multi-channel, digital signal processing
 - now with UNIX OS to simplified programming
- high bandwidth copper and optical fibre networking solutions
- affordable, multi-core CPU clusters to analyse data in real time
- reconstruction algorithms: neural networks, machine learning, . . .
- TPU chips - artificial intelligence accelerator ASIC

Many experiments already moving toward streaming readout !

Readout Scheme at LHCb

JINST 8 (2013) P04022

Real-Time Processing




*LHCb will move to a **triggerless-readout** system for LHC Run 3 (2021-2023), and process 5 TB/s in real time on the CPU farm.

Slide provided by Mike Williams (MIT)




Streaming Readout Scheme



Goal: Do as much as possible on-line and save everything that is useful

- extract relevant parameters for all channels
- extract time, position, momenta, angles, energy, ...
- discard noise and background hits when possible
- save high-level information
- combine data in sub-detector or sectors to form event segments
- option to organised into events or not !
- all detector information available for making decisions
- on-line calibration, data monitoring, and alignment


ASIC Board on Detector



Specifically designed boards mounted on the detector

- FEE must be matched to detector requirements
- multi-channel ASIC chips available with FADCs and multiplexing
 - e.g. 64 ch, 12 bit, 1 GSPS, < 20 mW/ch, rad. hard, < \$10/ch
- instead of FADC chip TDC chip could be used if appropriate
- zero-suppression, only viable signals passed on
- copper (supply power) or optical fibre (electrical isolation) to DSP


FPGA Signal Processing



Digital signal processing (DSP) with inexpensive FPGA boards

- can be situated at some distance (accessible) from the detector
- multiple input channels, de-multiplex FADC information
- analyse input channels in parallel for: time, charge, amplitude, ...
 - compress data from high frequency samples to a few parameters
 - flexibility to optimise software as needed
- high bandwidth output to CPU / GPU / TPU reconstruction

FPGA Signal Processing




Data monitoring and calibration

- option to write some fraction of the data to local storage
- cross check of signal analysis
- calibration, pedestals, gain, *etc.* for reconstruction

Add channel ID and time stamp for each channel's data stream

Time synchronisation over all channels important !

Reconstruction at Sub-Detector Level




CPU / GPU / TPU analyses sub-detector information

- analysis / reconstruction of localised data
- form clusters, track segments, PID, ...
- high-level parameters: position, time, momenta, angle, energy, ...
- save this high-level information in time slots

Perhaps this is sufficient ? Output at this stage ?

- no need to form complete events, leave for off-line analysis
- next stage, organisation into events, can be optional


Reconstruction at Sub-Detector Level



Cross check of reconstruction

- write some fraction of event to local storage
- check reconstruction algorithms and calibration
- as confidence in reconstruction / training grows
 - discard most of “raw” data → event size becomes manageable

Combine Sub-Detector Information for Event Selection



Combine sub-detector data to form complete tracks / events

- connect track segments, associate with calorimeter clusters, *etc.*
- categorise events
- write to disk for off-line physics analysis
- directly produce DSTs → reduce event size

Events immediately ready for physics analysis

- reduce time to achieve results / publication

Streaming Readout Issues

Readout scheme and detector, software, and analysis strongly coupled

- FEE needs to be tailored to detector and its requirements
- FPGA signal processing software specific to detector requirements
- possible to identify basic designs that can be modified to suit
 - change shaping time, sampling rate, number of bits, *etc.* as needed
 - code libraries to extract timing, amplitude, charge, analysis, *etc.*
- standardise components to be adaptable to many experiments
 - plug-n-play for streaming readout like old CAMAC and NIM modules

Detector must also be designed with streaming readout in mind

- detector can not require a trigger to initiate readout
- prefer fast detector response, avoid long response times
- plan ahead to simplify analysis
- detector, readout, and analysis must be designed together

Streaming Readout Development

MIT workshops to review status and plan for future

- first workshop January 27, 2017, second January 29–30, 2018
- now holding monthly video meetings
- electronics groups from JLab, BNL, ANL, SLAC involved
- TOPSIDE, sPHENIX, SOLid, LHCb, GlueX, DarkLight experimenters
- representatives from industry CAEN, AlphaCore



Lots of groups working towards streaming readout

- national labs have great resources but “doing their own thing”
- ASIC chip developers very interested but need direction
- need to coordinate efforts

Streaming Readout for sPHENIX

sPHENIX as a foundation of EIC streaming DAQ


- ▶ Versatility of EIC event topology calls for trigger-less streaming DAQ
 - 0.5 MHz interaction at top luminosity
- ▶ Start using sPHENIX-EIC full detector simulation to estimate trigger-less DAQ
 - Total streaming signal data rate on order of 100 Gbps
- ▶ Matched well with sPHENIX FELIX-based DAQ through-put rate
 - Designed to record data up to 200 Gbps
 - Similar architecture with ATLAS/LHCb/ALICE DAQ upgrade in 2020+
- ▶ EIC-connected calorimeter and tracker prototypes for EIC streaming testing



Streaming Readout for sPHENIX

A streaming DAQ architecture: FELIX

- ▶ Using PCIe FPGA card bridging stream-readout FEE on detector and commodity online computing
 - Similar approach taken at ATLAS, LHCb, ALICE upgrades and sPHENIX
- ▶ Implementation: FrontEnd Link eXchange (FELIX) PCIe FPGA card
 - 2x 0.5-Tbps optical link to FEE: 48x bi-directional 10-Gbps optical links via MniPODs and 48-core MTP fiber
 - 100 Gbps to host server: PCIe Gen3 x16
 - Large FPGA: Xilinx Kintex-7 Ultra-scale (XCKU115)
 - Interface to multiple timing protocols (SPF+, White Rabbit, TTC)
 - Developed at BNL for ATLAS Phase-1 upgrade, considering to use for streaming FEE readout in sPHENIX, proto-DUNE, CBM
 - Continued development to upgrade to 25-Gbps optical links, Vertex7 FPGA and PCIe-Gen4
- ▶ sPHENIX running 30x FELIX card production later calendar year 2018.
 - Interests of extra cards for EIC stream readout R&D welcomed



FELIX Card – v2.0/BNL712



FELIX timing interface mezzanine



FELIX-server test stands at BNL



Jin Huang <jhuang@bnl.gov>



Slide from Jin Huang (BNL)



Imaging Calorimeter for TOPSide

A) Imaging Calorimetry

Replace

Tower structure with very fine granularity (lateral and longitudinally)
Few 1,000 channels \rightarrow few 10,000,000 channels
Option to reduce resolution on single channels to low-bit depth



Technologies developed in past decade

Silicon sensors with $1 \times 1 \text{ cm}^2$, $0.5 \times 0.5 \text{ cm}^2$ and 0.16 cm^2 pixels
Scintillator strips ($4.5 \times 0.5 \text{ cm}^2$) or scintillator pads ($3 \times 3 \text{ cm}^2$)
Resistive Plate Chambers with $1 \times 1 \text{ cm}^2$ pads
Micromegas and GEMs with $1 \times 1 \text{ cm}^2$ pads

These technologies have been (mostly) validated



J. Repond: TOPSIDE



8

Slide from Jose Repond (ANL)

Conclusion

Streaming readout is possible !

- advances in technology → ASIC, FPGA, muti-core CPU, ...
- falling costs in electronics, computing, storage, and networking
- advances in software → neural networks, machine learning, TPUs, ...
- can expect further improvements over the next decade

Future EIC experiments will benefit from streaming readout

Implications for detector, electronics, software, and analysis

Important that streaming readout approach be endorsed now

So all groups can include this in their designs !



Thank You