Streaming Readout

Alexandre Camsonne

Hall A/C collaboration meeting

06/17/2022

Outline

- Triggered DAQ
- Streaming DAQ
- Example of streaming readout LHCb
- Pros and cons of streaming readout
- Streaming readout for EIC
- Streaming readout possibilities in Hall A and C
    - SBS
    - TDIS
    - SoLID
- Possible strategy for Hall A and C
- Conclusion

Triggered DAQ

- At beginning electronics or detector were very slow
- Not possible to record everything
- Need to design a trigger so event of interest is recorded and not missed because of dead time

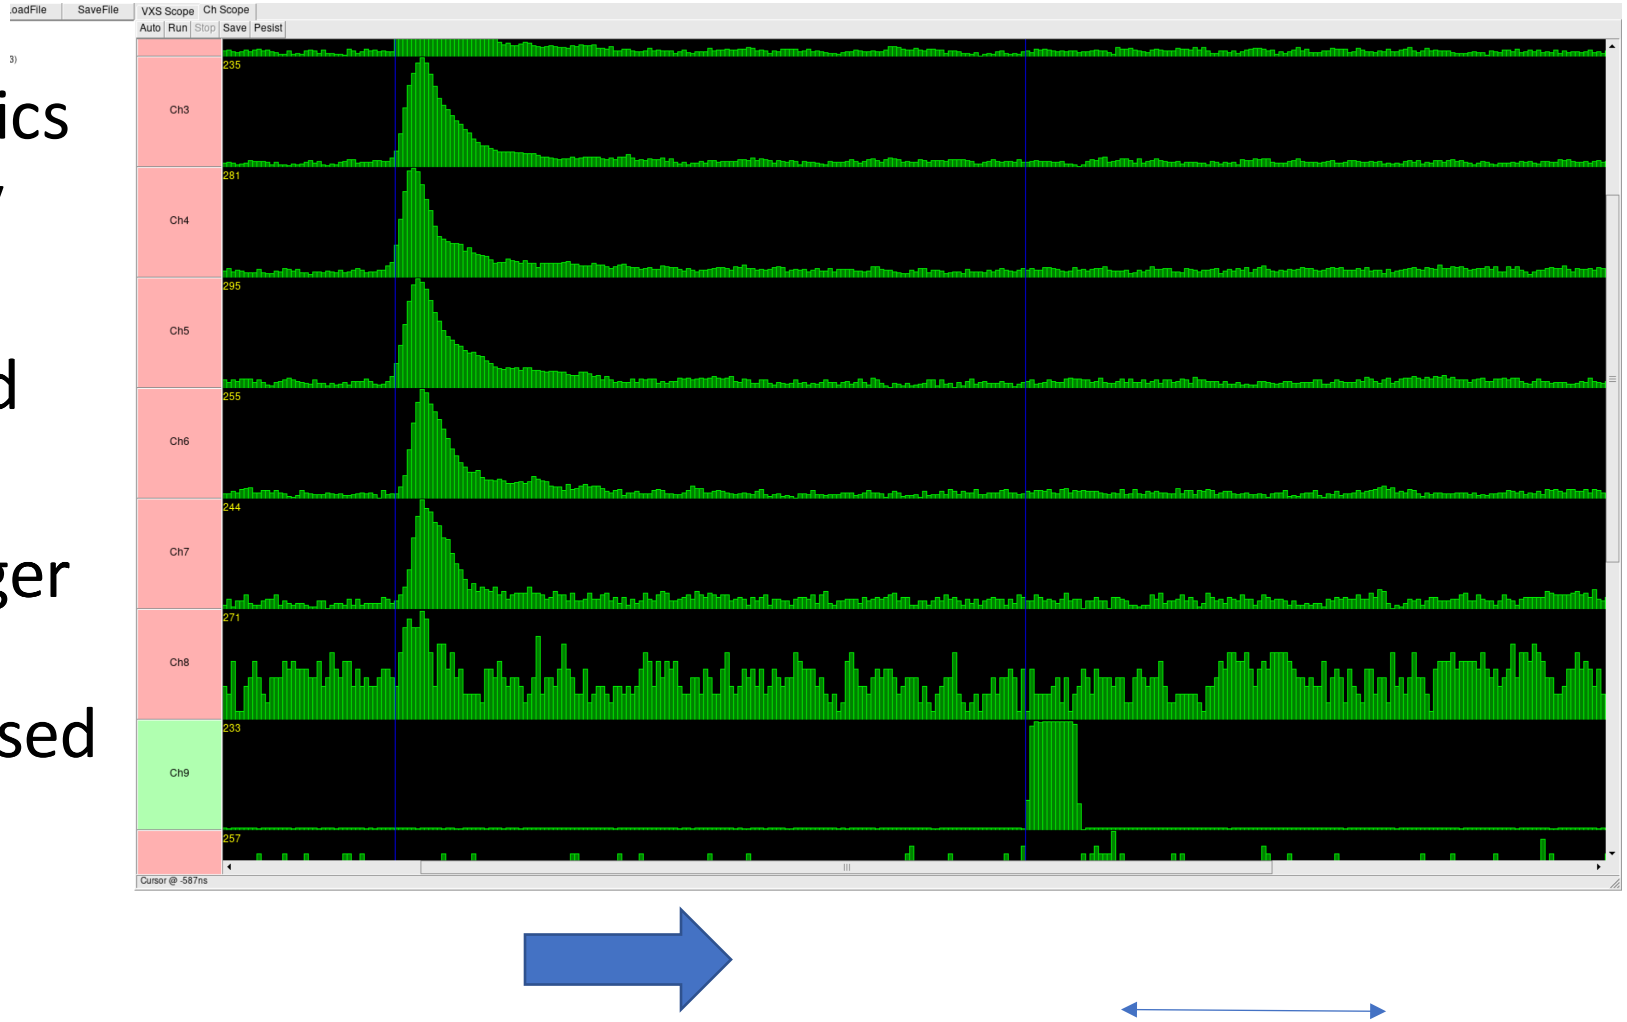

Advent of Flash ADC

- Progress in electronics (mainly wireless cell phones)
- High sampling rate, low power, continuous digitizing
- Example FADC 250

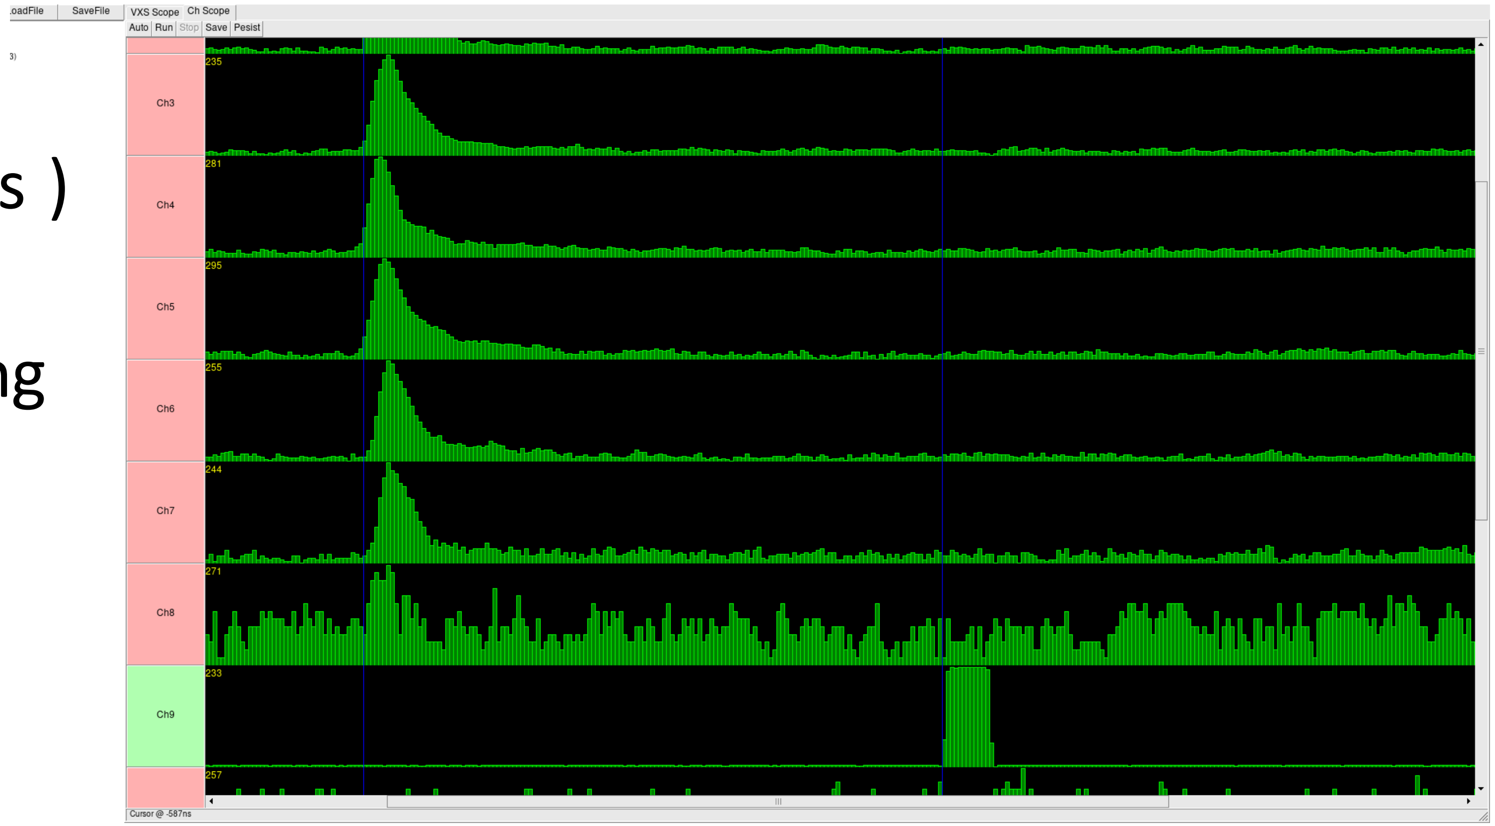

What is streaming readout DAQ

- Streaming is familiar word to everybody now : listen to musics or watch videos on internet. Record streams and playback

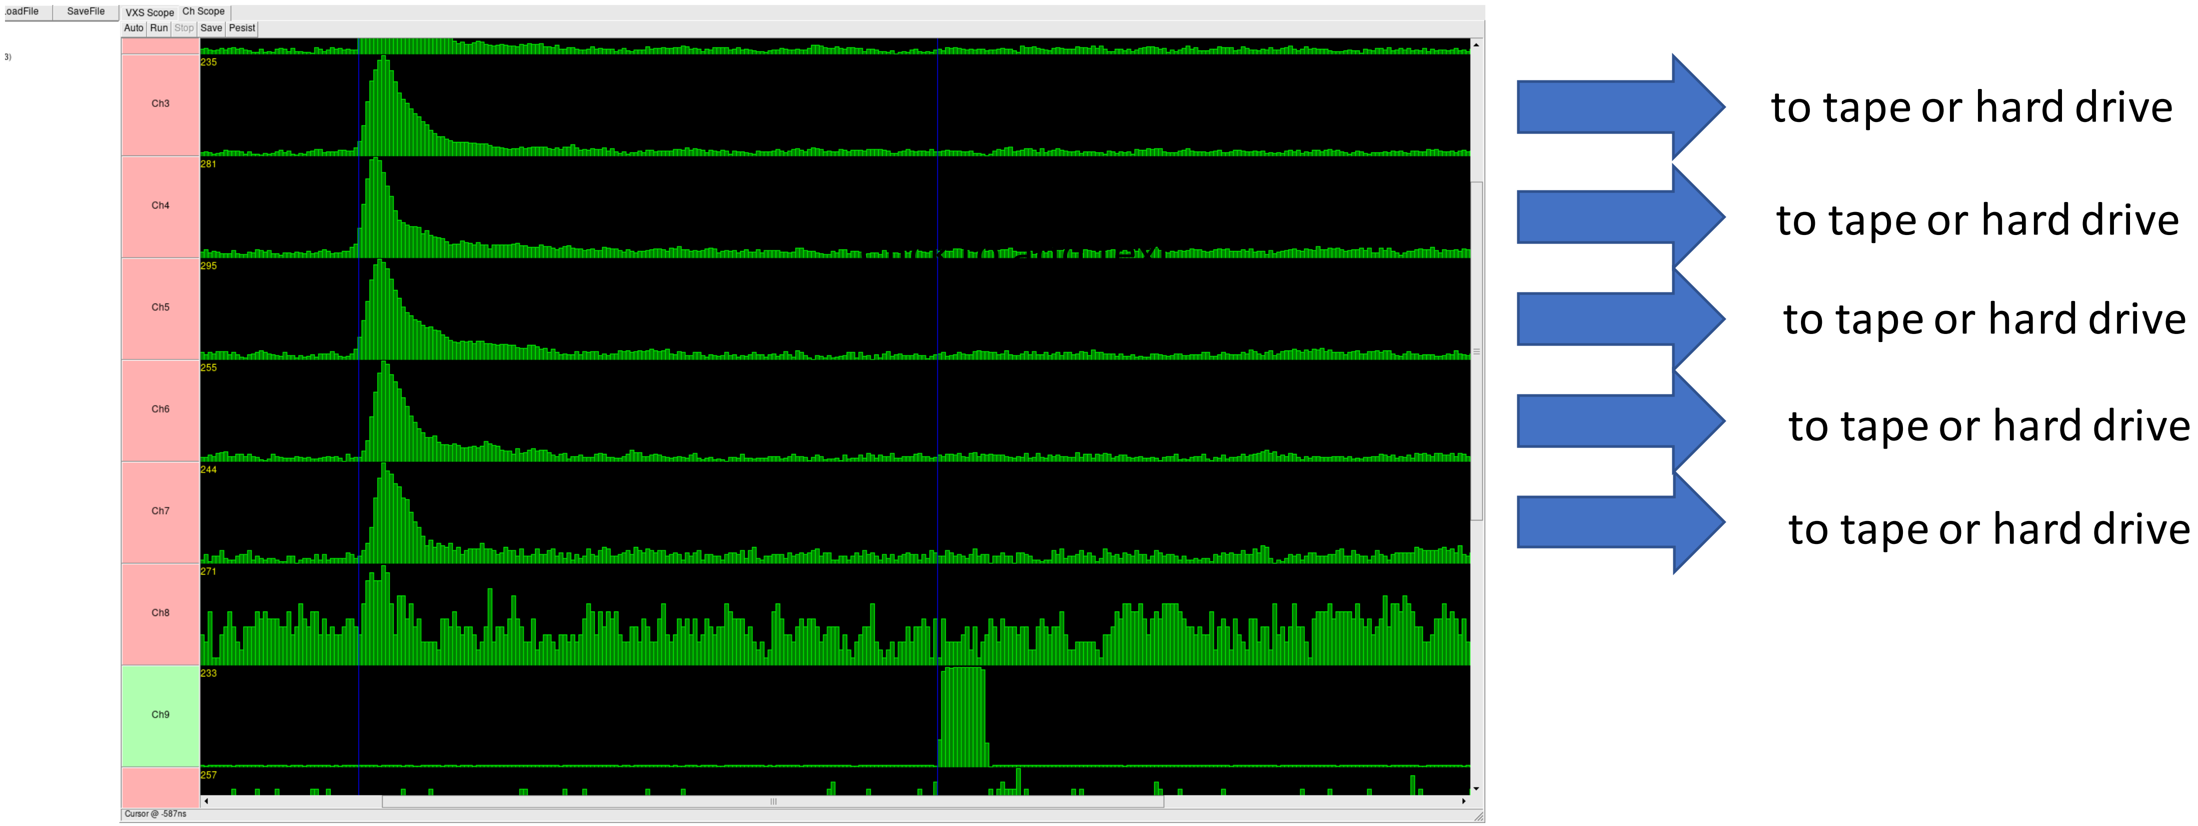

Semi streaming

- Can self trigger on data to reduce data – put a threshold

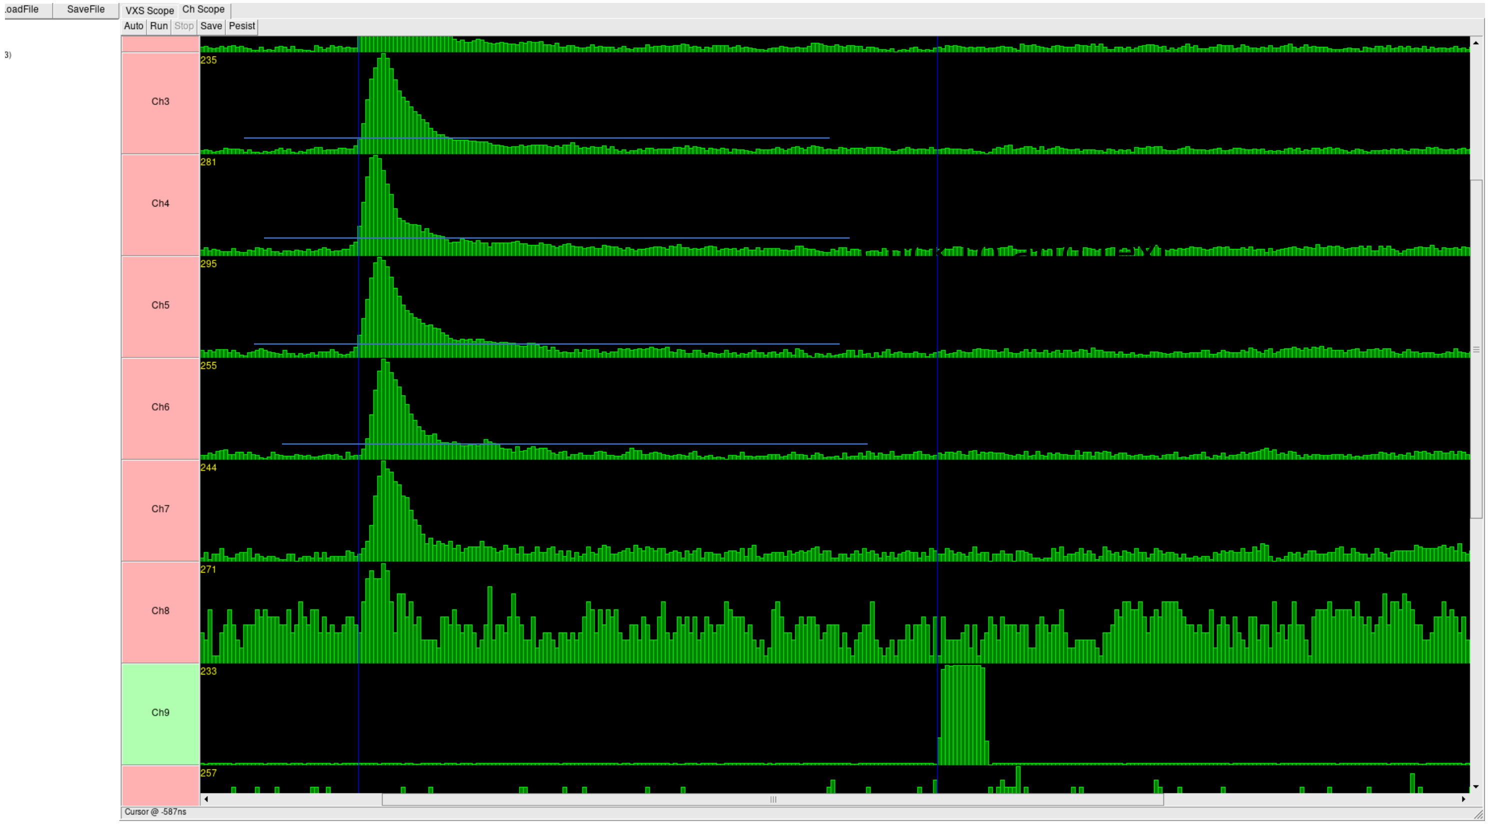

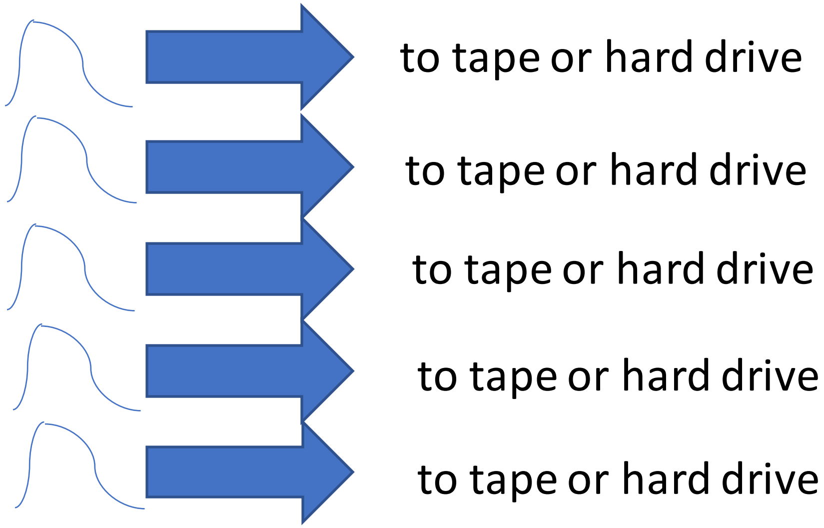

Reduce data rate from sampling rate 250 MHz to single detector rate ~ MHz

Ingredients for a streaming DAQ

- No dead-time
  - Fast front end electronics
    - FADC ideal (FADC ASICs or CHIPs : Alphacore, TI, SAMPA)
    - Electronics self triggered
  - Large data pipes
    - Ideally bandwidth should match electronics
    - 1 FADC channel : 2 bytes x 250 MHz = 500 MB/s = half of 10 gigE link
  - High speed data recording
    - Hard drive
    - Tape drive : LTO8 = 360 MB/s
- High resolution timestamp in each stream to be able to put back channels together
- Highly desired : background reduction and compression to reduce data footprint

Pros of streaming readout

- Most unbiased way to take data
- Can record data and go back to old data
  - Data always good, no way to mess trigger
  - Can look at additionnal physics channel that a trigger would have excluded
- Simple electronics design
- Simple scaling
  - Adding a detector is trivial
- Push data to regular farm computers
- No custom electronics
  - High level programming available
  - Complex algorithm can be used
  - Heterogenous computing friendly
  - Ideal for AIML

Cons of streaming

- Large amount of data
- Sensitive to background
  - Detector noise : exemple SiPMT dark rate
  - Electronics background : if thresholds to low in electronics noise
  - Physics background
    - Photons conversion
    - Unwanted high rate physics process
- FADC : use more power than triggered electronics, can make issue for integration

The Network Bandwidth will grow!

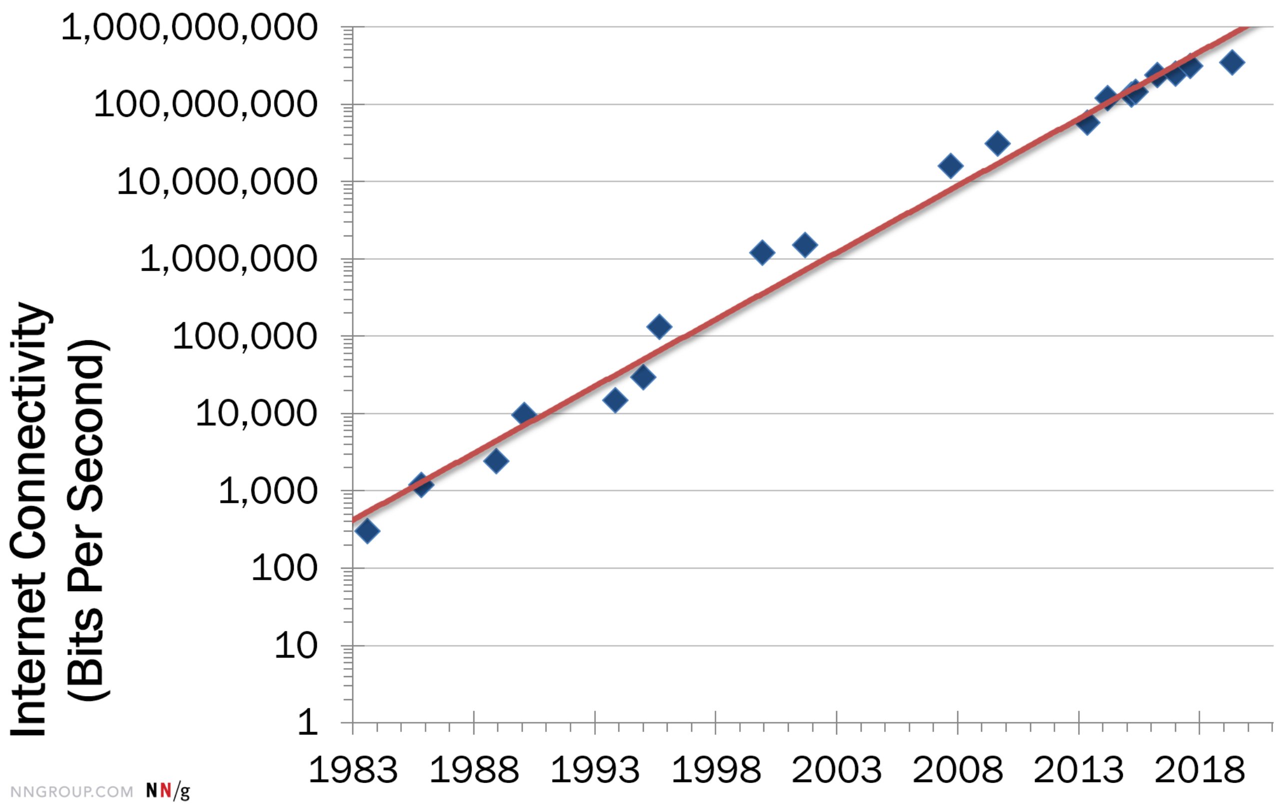

Eli Dart (leader of Science Engagement group of ESnet) e-mail:

"As far as network backbone capacity in the 2030s goes, it's too early to say. ESnet's optical system is highly capable (we just deployed it as Jason said), but we don't yet know what networking technologies will be available in 2032 - ten years is a long time. That said, we can make some guesses. It is likely that 1.6Tbps Ethernet will be available by then."

Jason Zurawski (Esnet, in group of above) e-mail:

"I think BNL plans to support multipole 400G connections, JLab is limited due to a lack of fiber recourses in the area, but will at a minimum upgrade some of the 10G connections that exist today to 100G soon (and move to 400G sometime in the future I would imagine).'

		Annualized Growth Rate	Compound Growth Over 10 Years
Nielsen's law	Internet bandwidth	**50%**	57×
Moore's law	Computer power	**60%**	100×

Tape roadmap

	Generation 5	Generation 6	Generation 7	Generation 8	Generation 9	Generation 10	Generation 11	Generation 12
Compressed Capacity	3TB	6.25TB	15TB	30TB	60TB	120TB	240TB	480TB
Compressed Transfer Rate	280MB/s	280MB/s	750MB/s	900MB/s (Note1)	TBD	TBD	TBD	TBD

Breakthrough on storage would help : quantum storage ? 1 bit per atom ?

Golden example of streaming readout LHCb

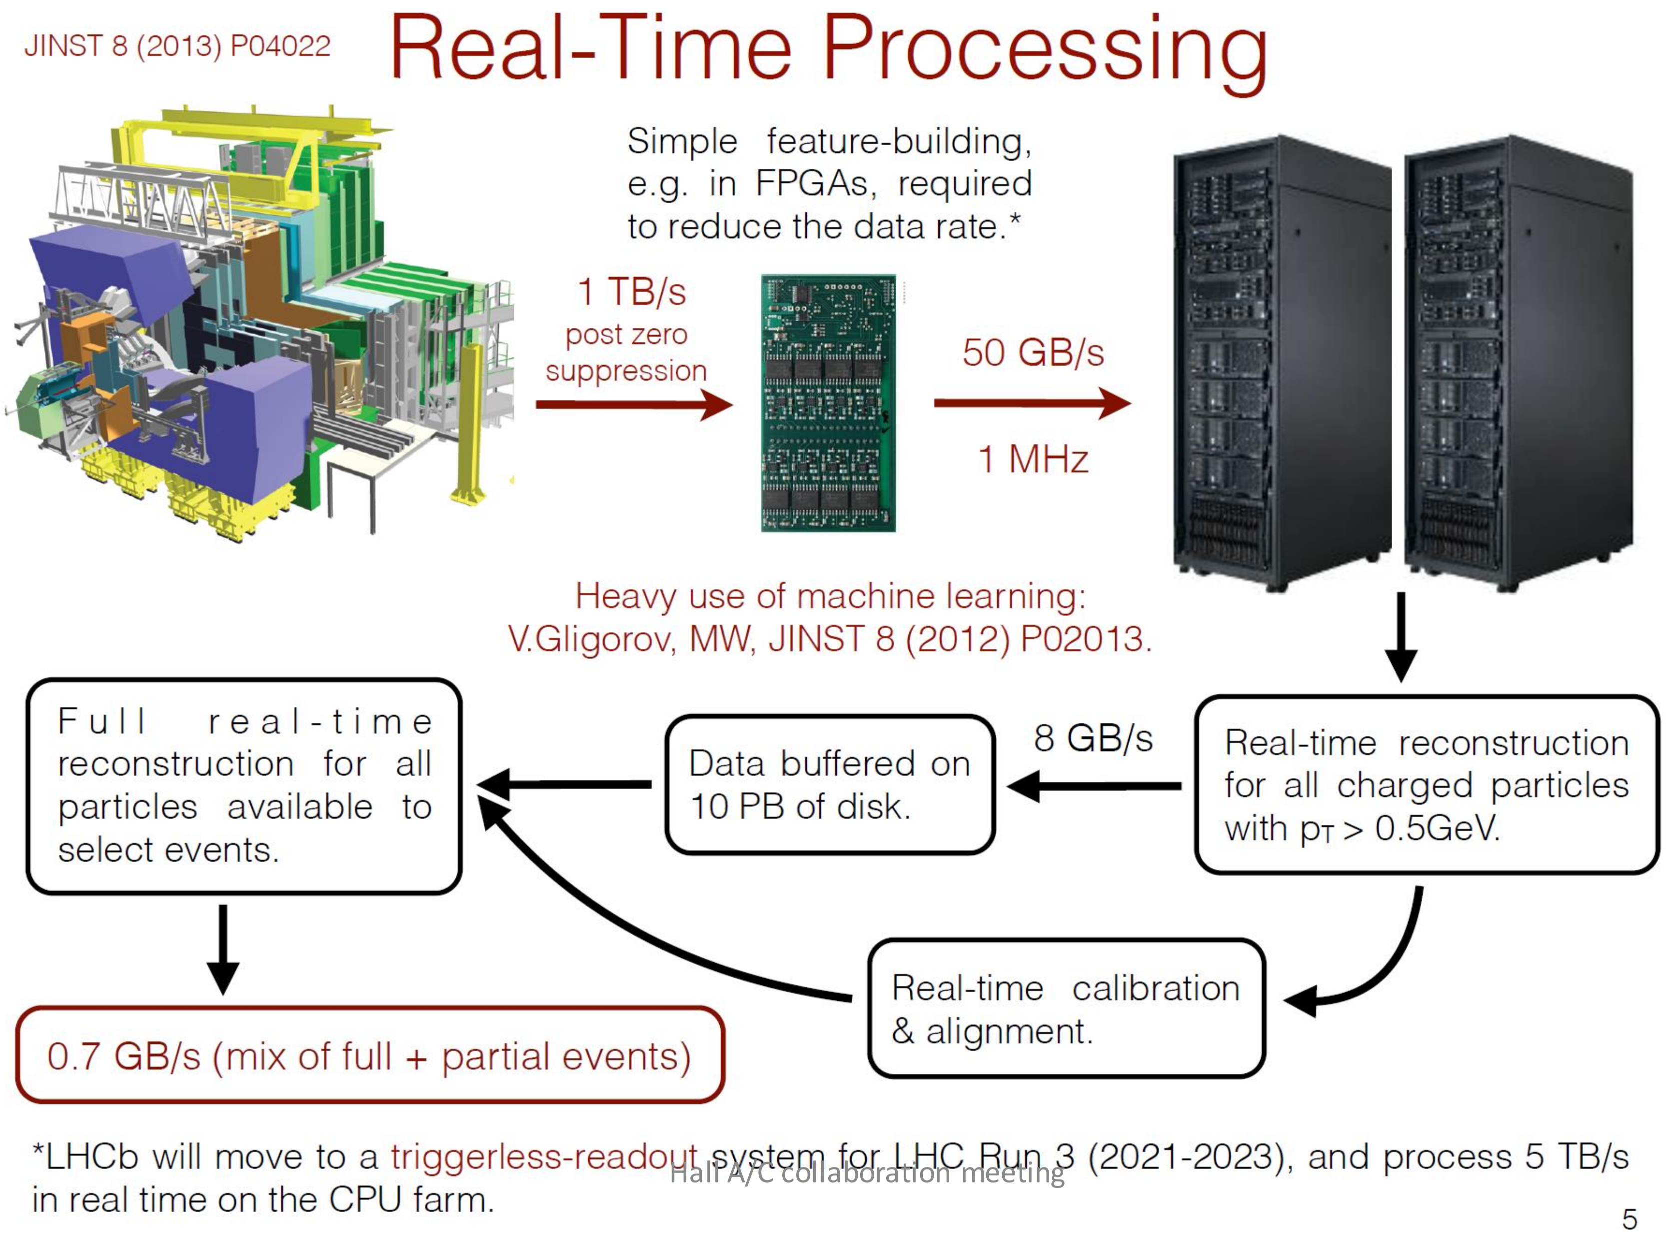

Golden example of streaming readout LHCb

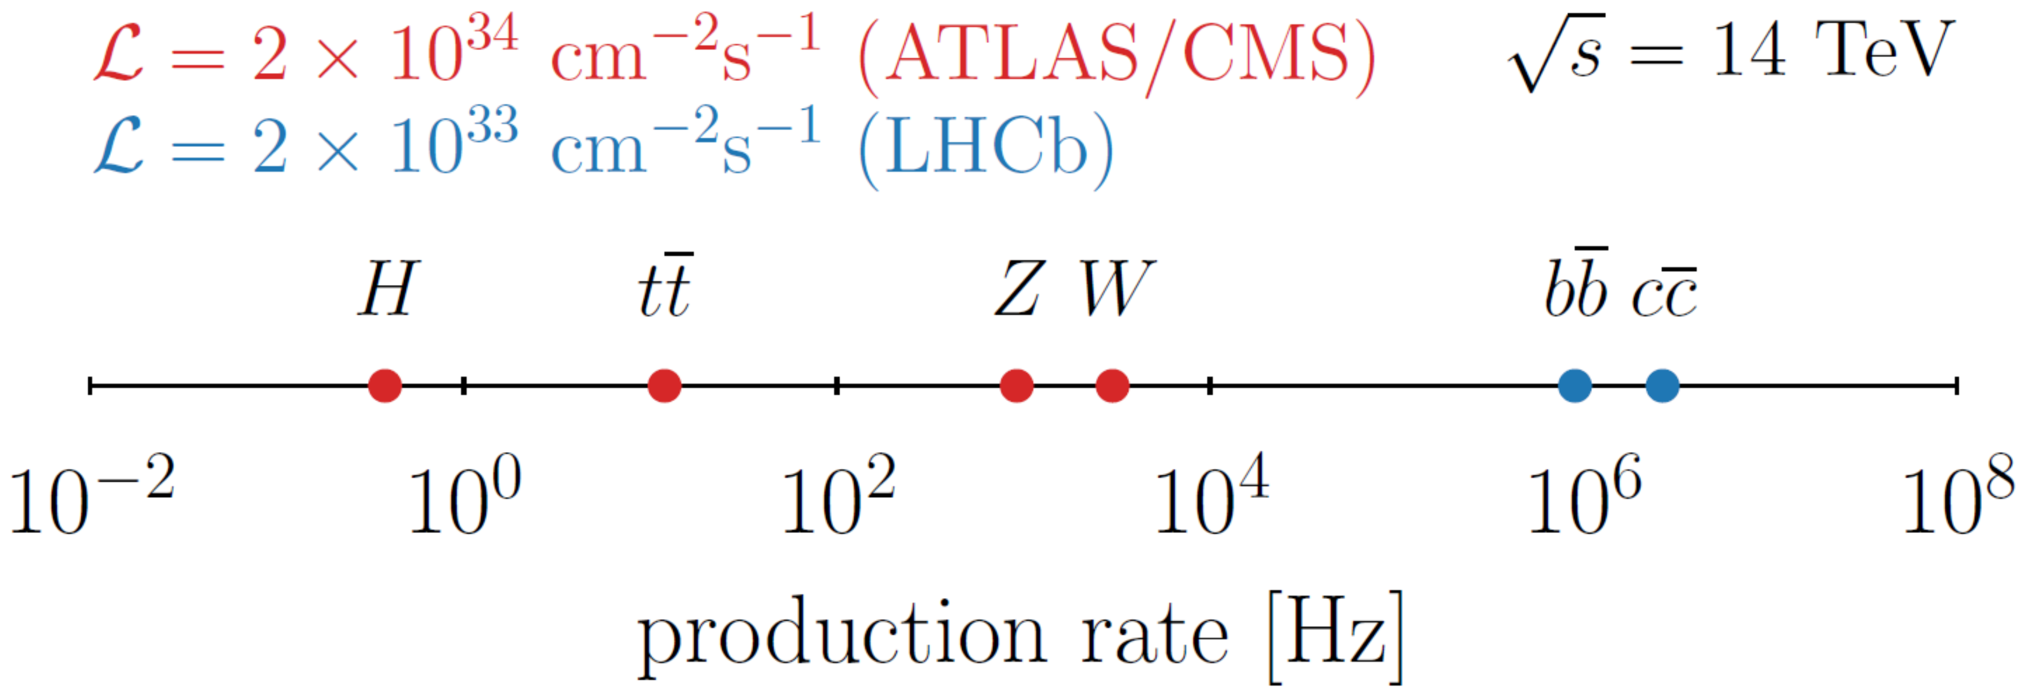

CERN experiments

	LHCb, **Run 3**	ALICE, **Run 3**	ATLAS/CMS, Run 4
Hardware trigger	**No**	**No**	Yes
Readout rate	40 MHz pp	50 kHz PbPb	$\sim$ 1 MHz pp
Data rate	5 TB/s	3.5 TB/s	$\sim$ 5 TB/s

ALICE

- Large event sizes
- Reconstruction time dominated by TPC tracking with $\sim 80\times$ speedup on GPUs
- Clear use case for GPUs

LHCb

- Extremely high rate of small events
- No step dominates reconstruction time
- Advantages of GPUs compete with challenges

LHCb

Allen is LHCb's GPU-based first level software trigger (HLT1)

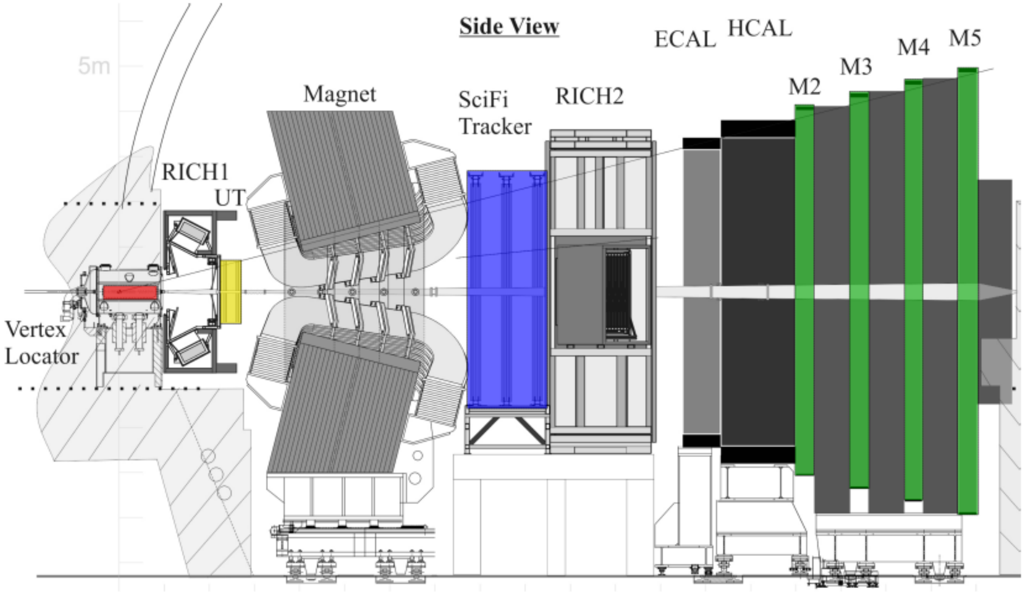

- Decode data from the **VELO**, **UT**, **SciFi**, and **Muon** systems
- Cluster detector data into "hits"
- Build tracks (**VELO**, **UT**, and **SciFi**)
- Find primary vertices (PVs) (**VELO**)
- Match tracks to **Muon** hits

Work as a standalone application or as part of LHCb's software stack

Can be compiled for CPU and GPU with CUDA or HIP

LHCb data processing

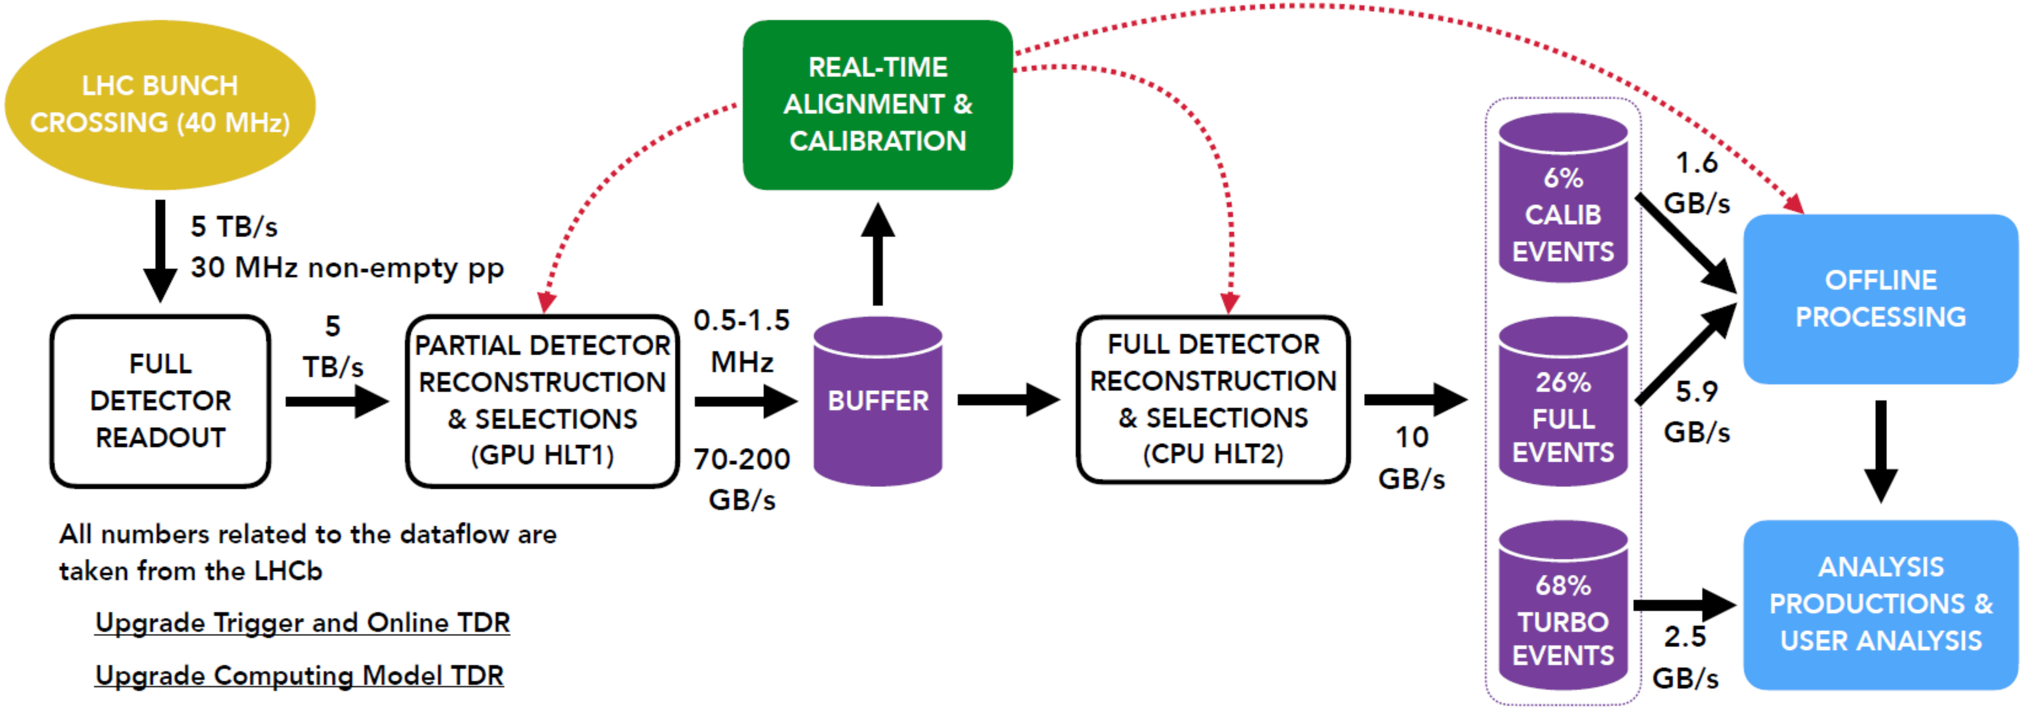

All numbers related to the dataflow are taken from the LHCb

- Upgrade Trigger and Online TDR
- Upgrade Computing Model TDR

LHCb data processing

Hardware trigger performance LHCb-TDR-016

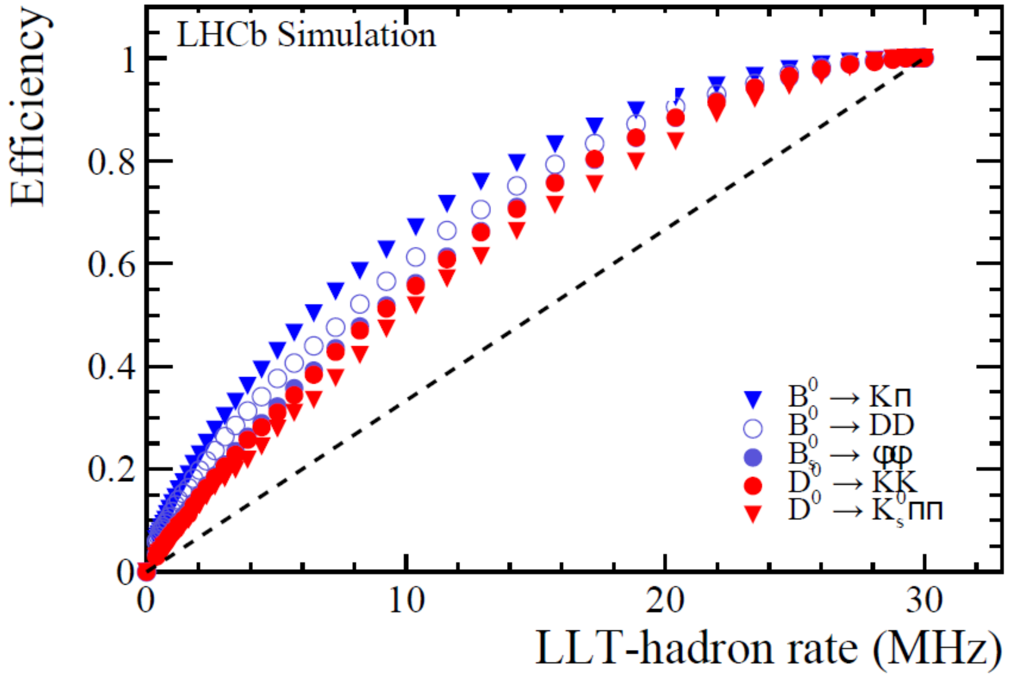

- Heavy flavor hadrons are long-lived and decay into low-momentum particles
- Can't effectively trigger on heavy flavor decays using hardware signatures
- Solution: process every event (30 MHz, 5 TB/s) in software

~~Triggering~~ $\to$ Real-Time Analysis

LHCb data processing

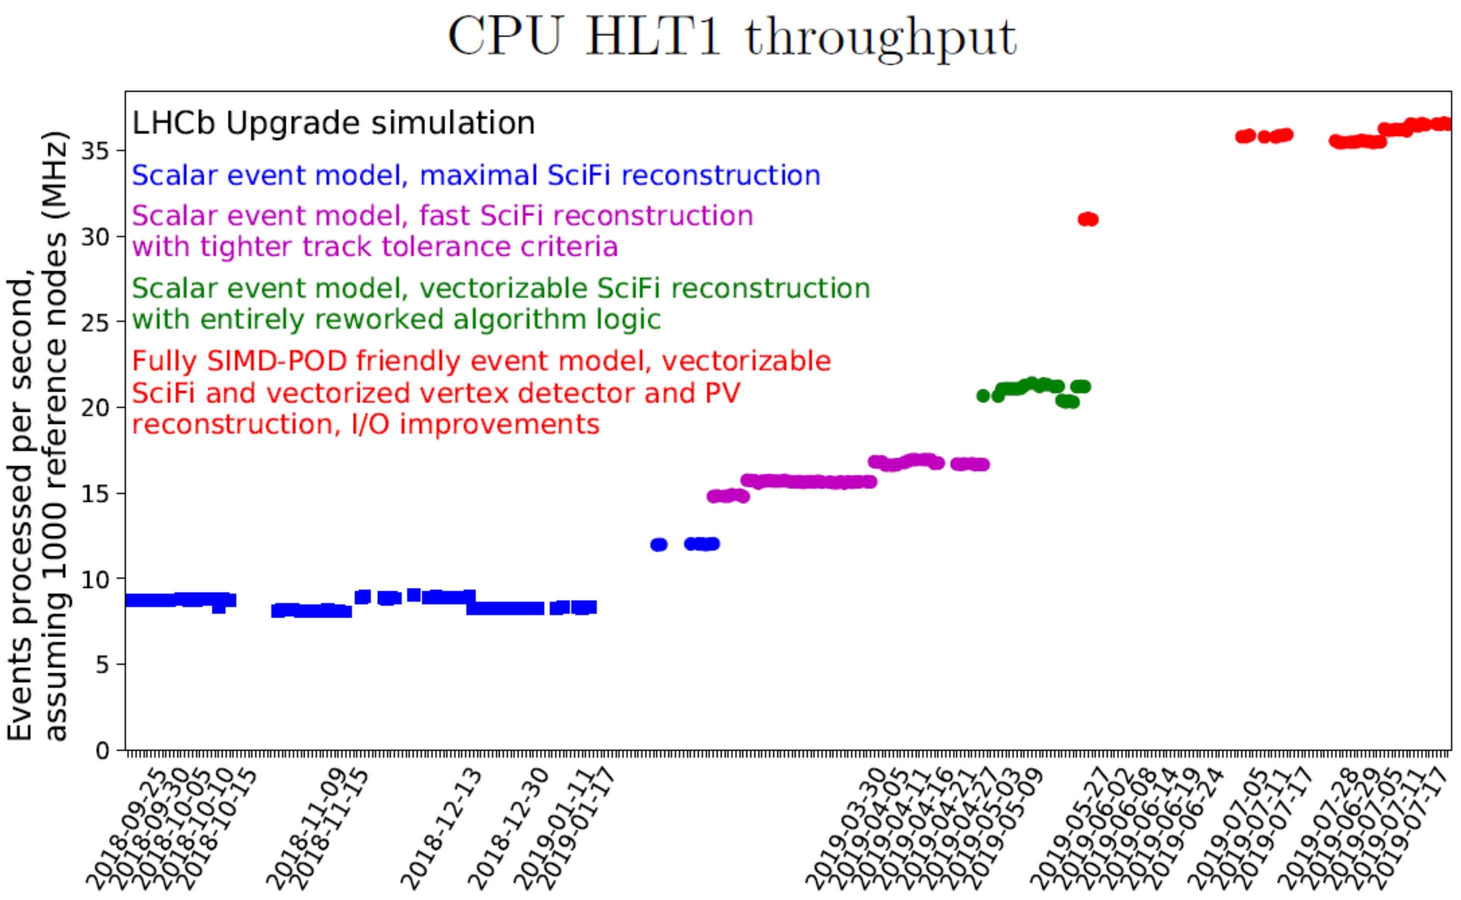

Streaming for EIC detector

- Total crossection

σ_{tot}(μb)		E_e [GeV]		
		5	10	18
E_p [GeV]	41	25.9	30.1	35.0
	100	32.1	37.1	41.6
	275	39.4	44.6	49.3

EIC Multiplicities from YR

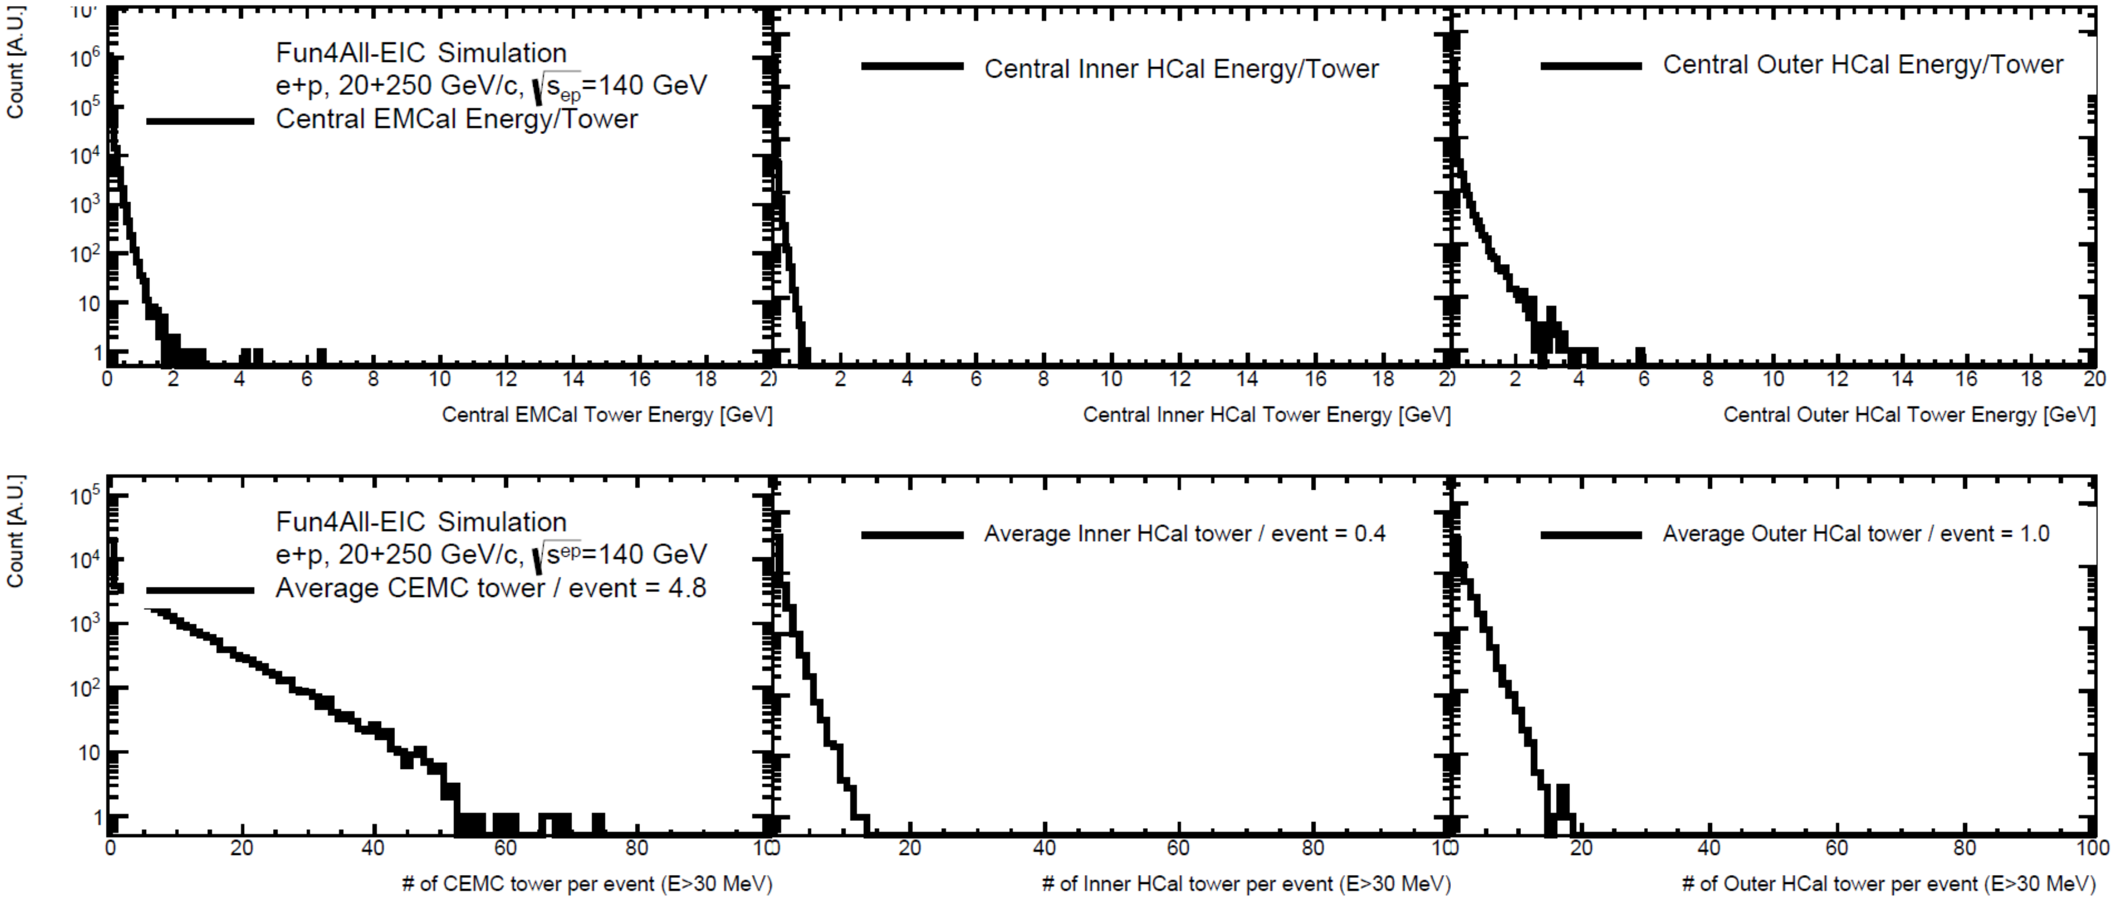

Multiplicities in trackers from YR

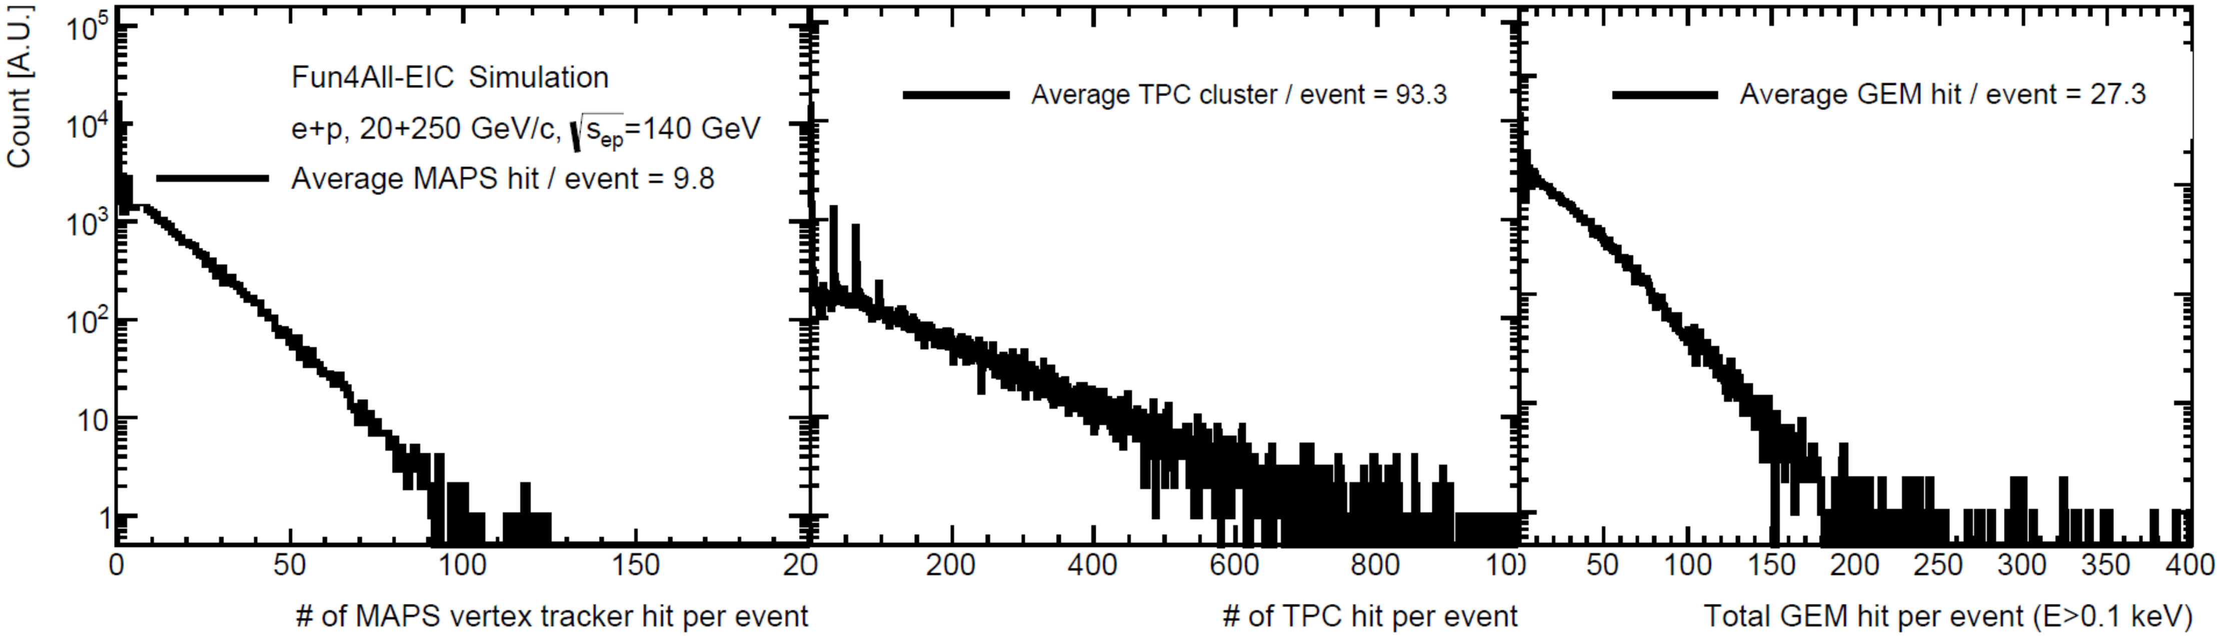

Streaming DAQ requirement

Collider:

- EIC will deliver beam in up to 1160 bunches (~100MHz)
- Expected physics rate ~500KHz

Detector:

- Streaming Design
  - Physics requires min-bias data
  - High collision rates
  - low detector occupancy
- O(100 Gbps) output rates

Streaming for EIC

- EIC data rates much smaller than LHC or RHIC
- Lower rates and lower multiplicities than p-p or ion-ion

- Detector can be optimized for streaming readout and AIML
  - Choice of readout
  - Optimization of detector

- Opportunity to really have a streaming readout ie save the streams and not only record reconstructed events of interest
- Allow to look at complicated event topologies
- Example for physics with backward photon tagger much easier to handle than triggered

Raw Data Requirements *(estimated)*

ECCE Runs	year-1	year-2	year-3
Luminosity	$10^{33}\text{cm}^{-2}\text{s}^{-1}$	$2 \times 10^{33}\text{cm}^{-2}\text{s}^{-1}$	$10^{34}\text{cm}^{-2}\text{s}^{-1}$
Weeks of Running	10	20	30
Operational efficiency	40%	50%	60%
Disk (temporary)	1.2PB	3.0PB	18.1PB
Disk (permanent)	0.4PB	2.4PB	20.6PB
Data Rate to Storage	6.7Gbps	16.7Gbps	100Gbps
Raw Data Storage (no duplicates)	4PB	20PB	181PB
Recon process time/core	5.4s/ev	5.4s/ev	5.4s/ev
Streaming-unpacked event size	33kB	33kB	33kB
Number of events produced	121 billion	605 billion	5,443 billion
Recon Storage	0.4PB	2PB	18PB
CPU-core hours (recon+calib)	191Mcore-hrs	953Mcore-hrs	8,573Mcore-hrs
2020-cores needed to process in 30 weeks	38k	189k	1,701k

Yellow report streaming DAQ design

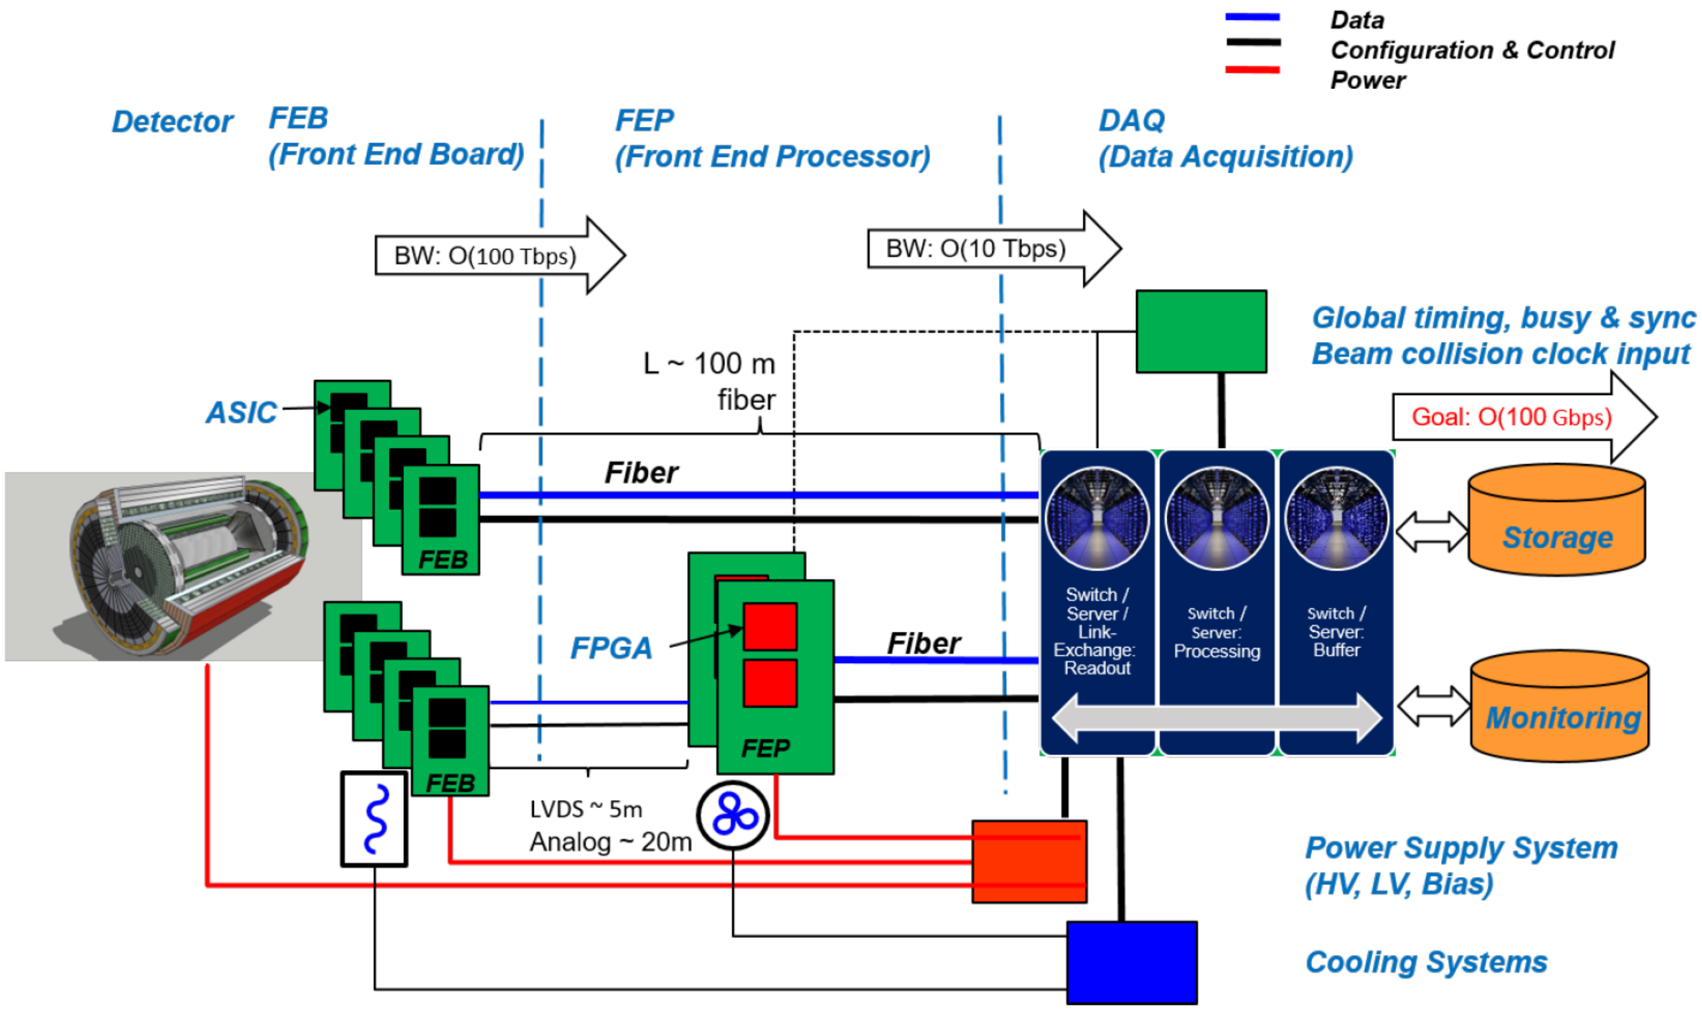

DAQ: Overview

- **Community solidly behind Streaming Data Acquisition System (SRO)**
  - Widely recommended by experts: EIC Computing Consortium, EIC Yellow Report
  - No need to wait for all signals from single crossing to read out data
  - Removes nearly all deadtime
  - Less restrictions for filter criteria and potentially less bias

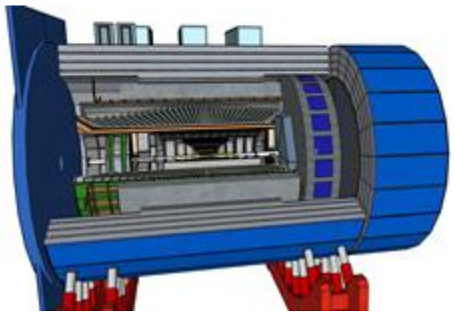

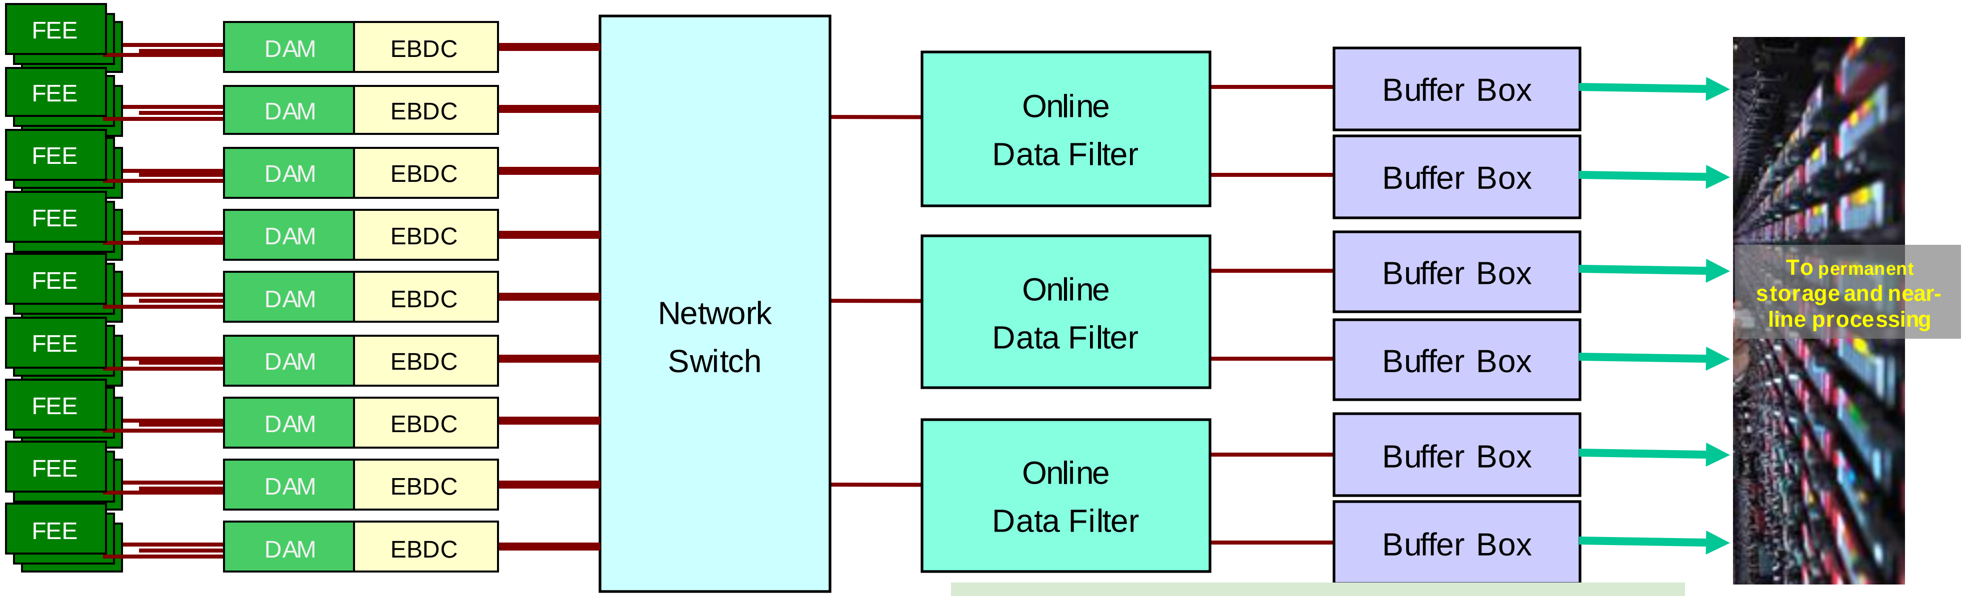

FEE = Front End Electronics
DAM = Data Aggregation Module
EBDC = Event Buffer / Data Compressor

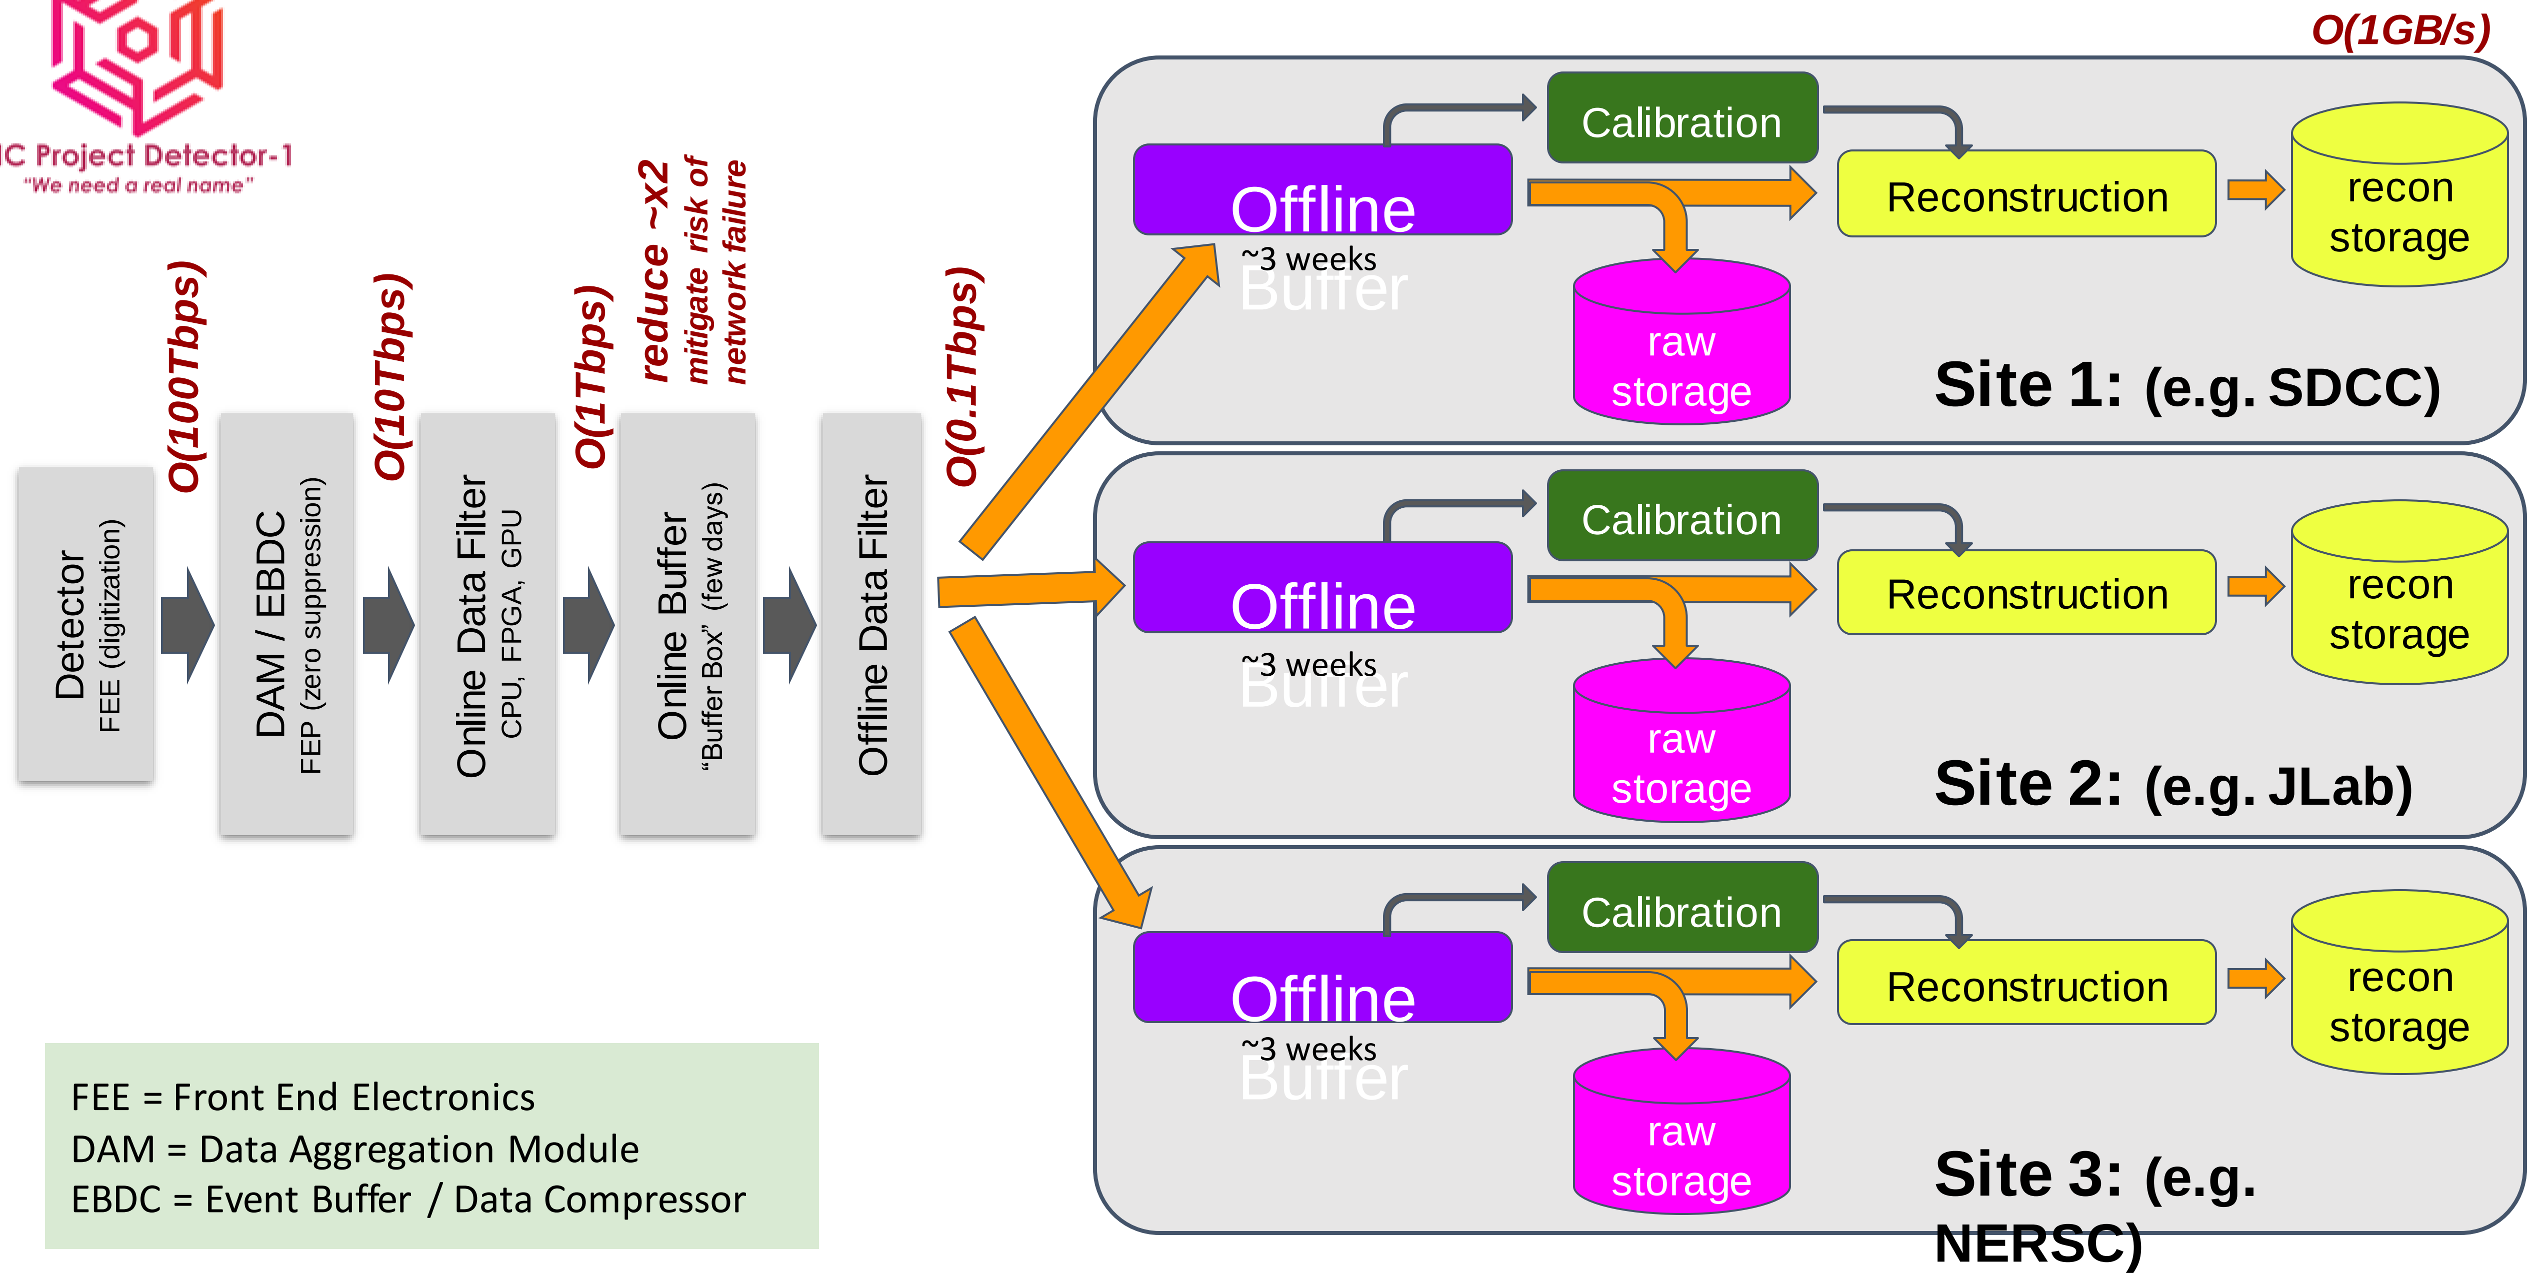

FEE = Front End Electronics
DAM = Data Aggregation Module
EBDC = Event Buffer / Data Compressor

Some Obvious Points

Detector	Channels
RICH	625k
Calorimeters	72k
Imaging Calorimeter (MAPS)	619M
Tracking MAPS	60B
Tracking MPGDs	144k
Far Forward (MAPS)	400M
Far Forward	1100k
Far Backward	4k
TOF	332k
Total MAPS	61B
Total Channels	2.3M

- Assume O(100Gbps) Bandwidth to tape
  - Hit Size ~ 64 bits (24 bits time, 24 bits position, 12 bits, 4 status)
  - 1 Hit / bunch crossing = 6.4Gbps (~5% of bandwidth assumed)
  - At 500Khz, average event size ~25KB (~3.1K hits / event)
- 2.3M channels + 60k MAPS sensors
  - In flux, but currently ~4.5k Fiber reading into ~120 FELIX boards
- Assume O(3.5Tbps) Bandwidth to DAQ computers
  - 2.3M channels + 60k MAPS sensors
    - ~4.5k Fiber reading into ~120 FELIX boards
    - ~ < 1Gbps / fiber (fiber capacity ~6 Gbps)
    - ~ < 25Gbps / FELIX (FELIX bandwidth ~100 Gbps)
- We will need to pay a lot of attention to noisy channels, flaky fibers, and any other potential noise sources

Significant SiPM Dark Currents expected in dRICH and pfRICH detectors

- Sensors are inside the 3T magnetic field,
  - SiPM sensors are envisioned.
  - Thresholds must be sensitive to single photons
- Dark Currents at this threshold ~3Khz / channel increasing to ~300Khz / channel after several years after which annealing will be performed
  - Mitigate by ~3 by applying timing window with respect to bunch crossing
  - Read up to 3.3Tbps into the DAQ computers but filter using
    - Software trigger
    - Potentially ML/AI if turns out practical
  - Software trigger reduces dark current volume to ~<15.5Gbps
  - As a potential mitigation the timing system & FELIX could be adopted to supply hardware trigger.

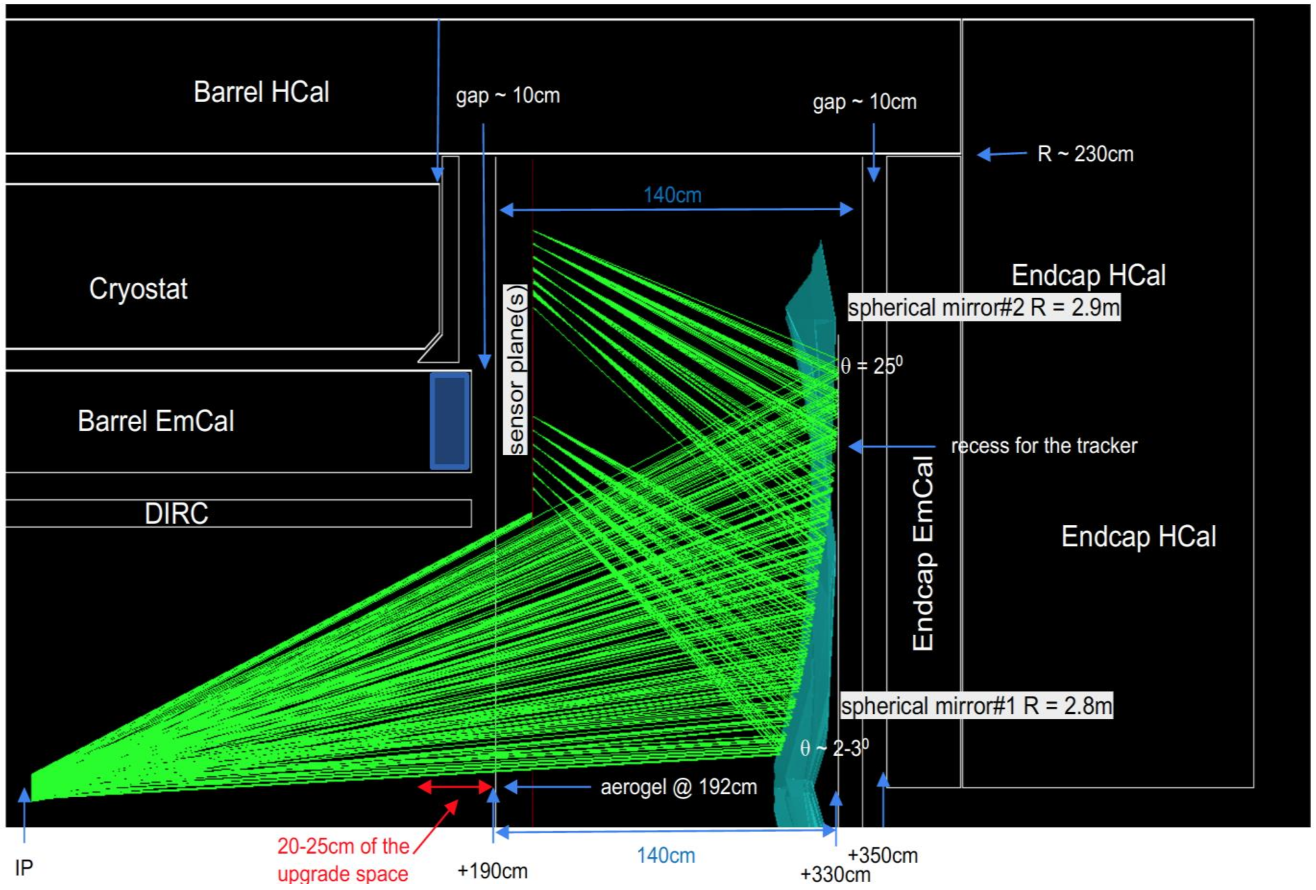

Detector Expert Estimates of data volumes

Detector	Channels	DAQ Input(Gbps)	DAQ Output(Gbps)
B0 Si	400M	<1	<1
B0 ac-lgad	500k	<1	<1
RP+Offm+ZDC	700k	<1	<1
FB Cal	4k	80	1
eCal	34k	5	5
hCal	39k	5.5	5.5
imCal	619M	4	4
Si Tracking	60B	5	5
Micromegas Tracking	66k	2.6	.6
GEM Tracking	28k	2.4	.5
uRWELL Tracking	50k	2.4	.5
dRICH	300k	1830	14
pfRICH	225k	1380	12
DIRC	100k	11	11
TOF	332k	3	.8
Totals		3400	61

Far forward detectors have low acceptance

Far backward do low Q tagging and Luminosity measurements but have high signal rate due to bremsstrahlung. The data is used primarily for histogramming, but also have subset that will be readout in concert with central detector collisions

Calorimeters with SiPM readout have higher thresholds and time-clustering in FEE

MAPS have enormous numbers of very quiet pixels. Also, they read out over 100-200ns time (<< 2us)

dRICH/pfRICH subject to SiPM dark currents

Synchrotron radiation studies from YR

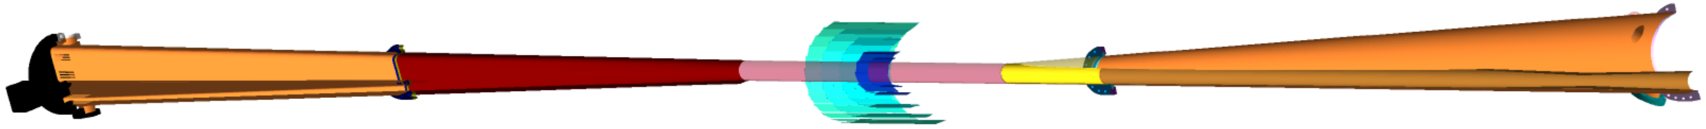

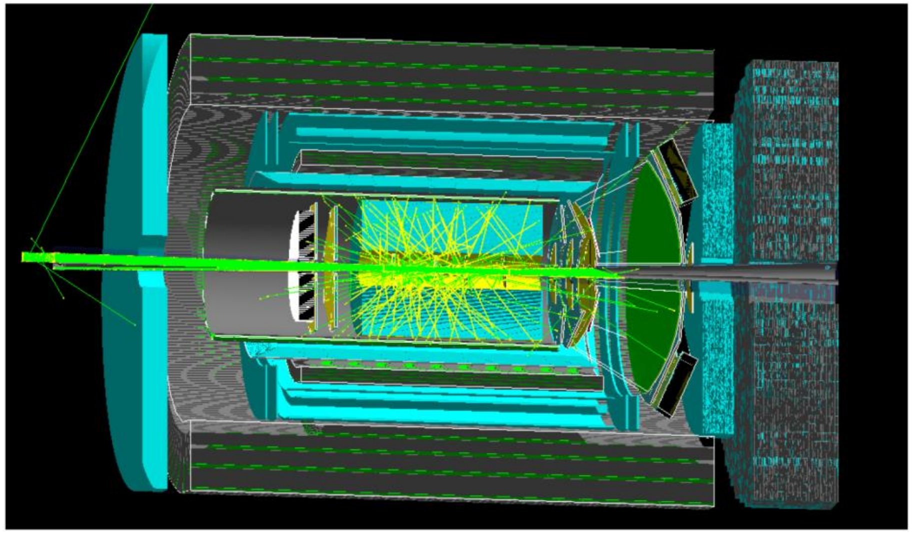

YR Geant4 simulation using synchrotron radiation input files from Synrad

Synchrotron radiation in trackers

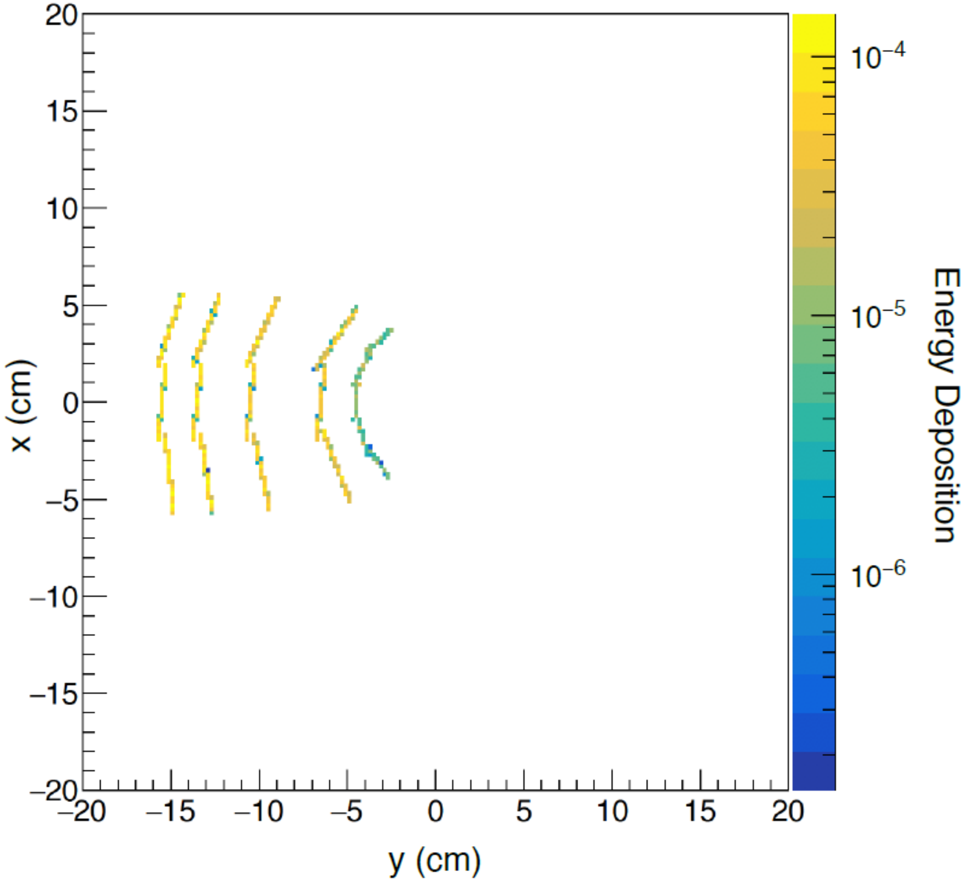

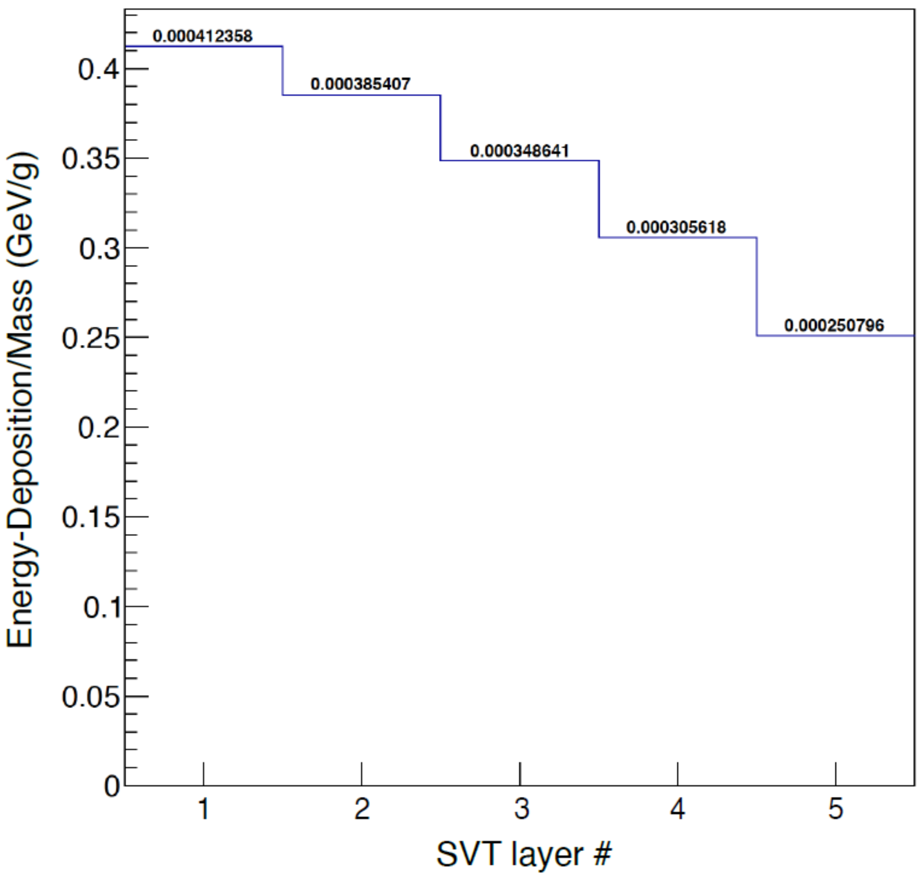

Energy deposit in 5 layers silicon vertex tracker

Beam Gas & Synchrotron Radiation Simulations

(Elke Aschenauer, Adam Jaroslav, Zhengqiao Zhang, Deepak Samual, Reynier Cruz Torres)

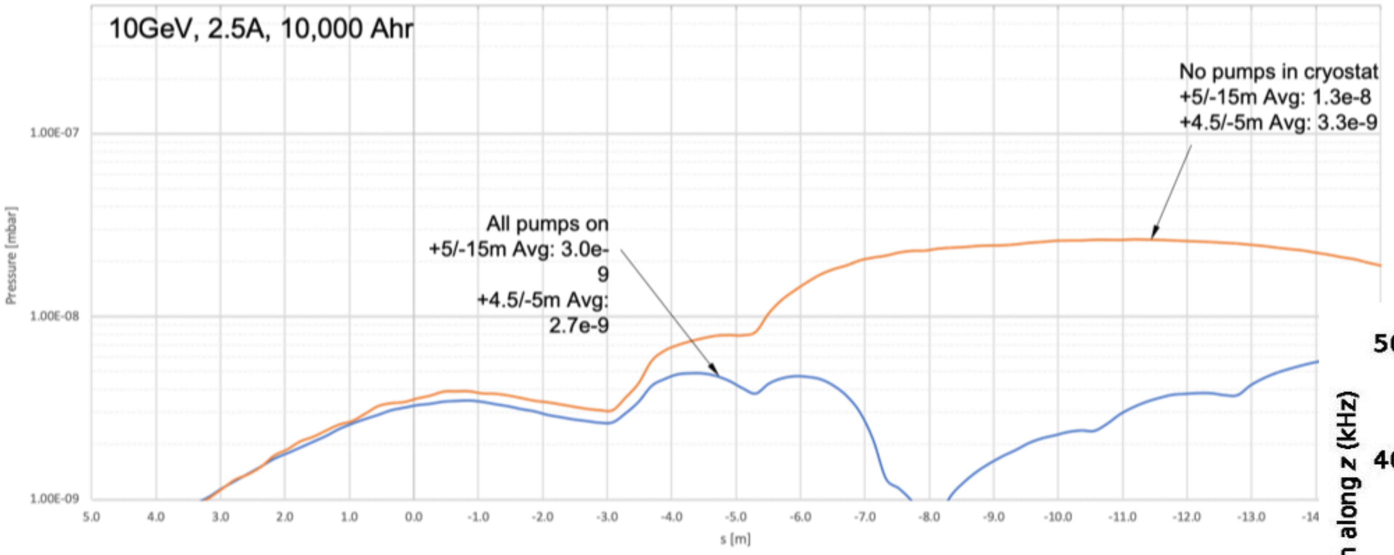

Vacuum

“For all simulations, physics, detector performance and backgrounds, it is very important to simulate all the different beam effects, like crossing angle, divergence…”

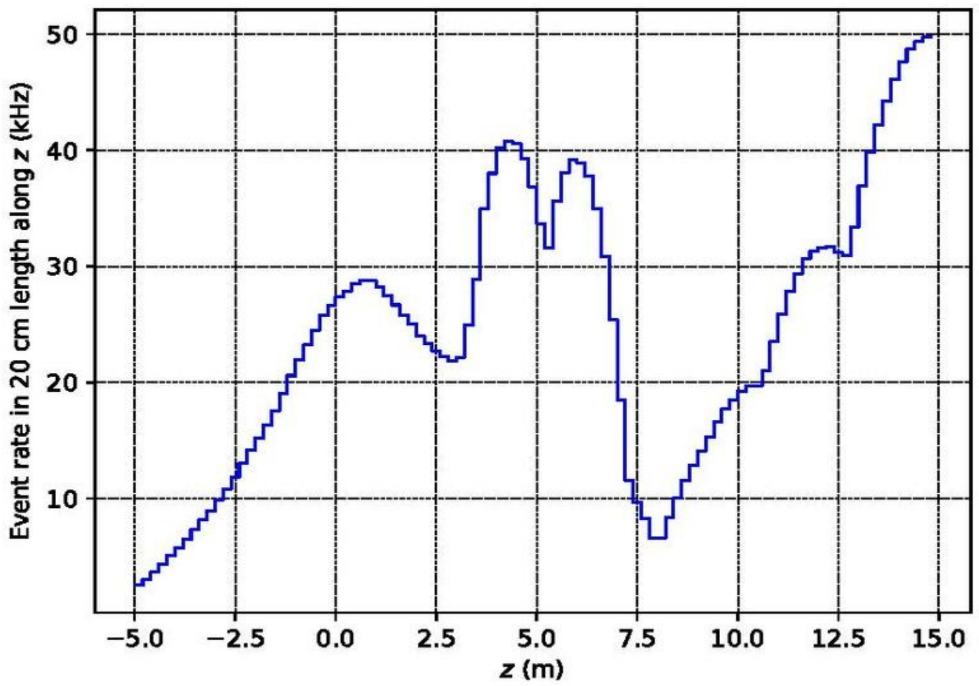

Electron Beam Gas Production

Beam Gas from YR

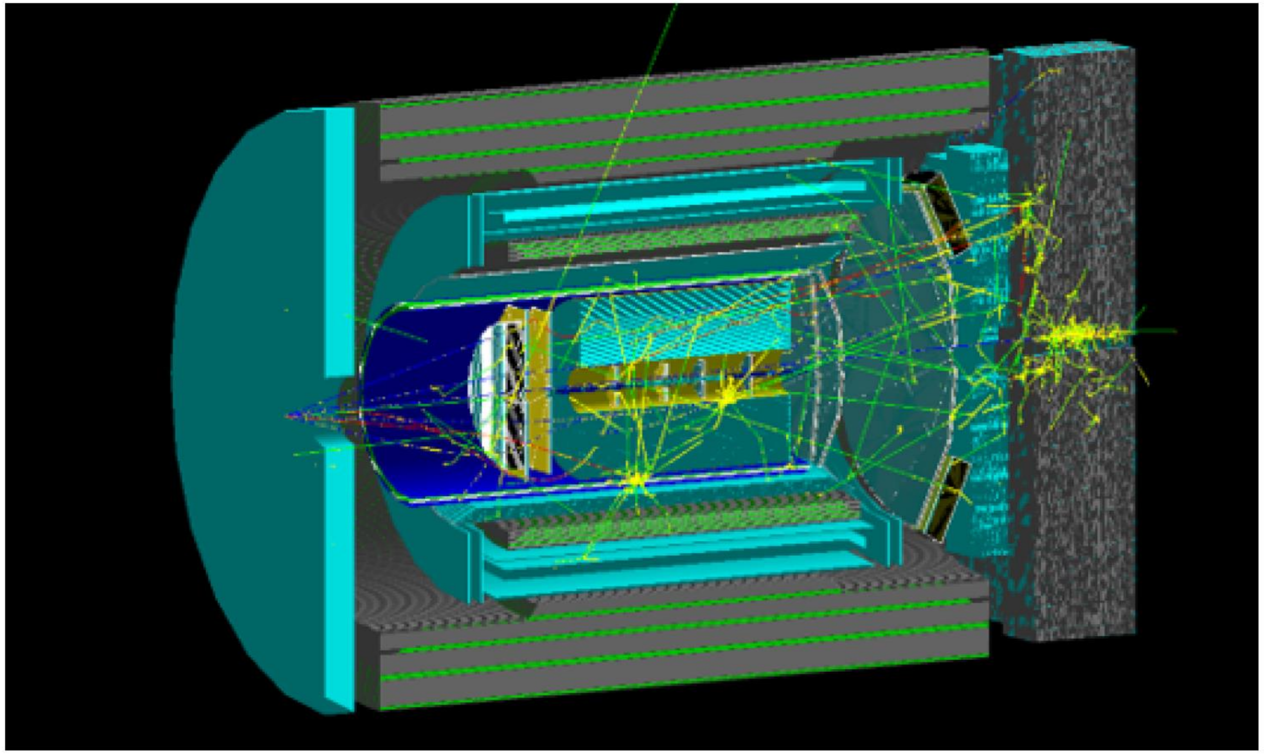

Beam gas data rates from YR

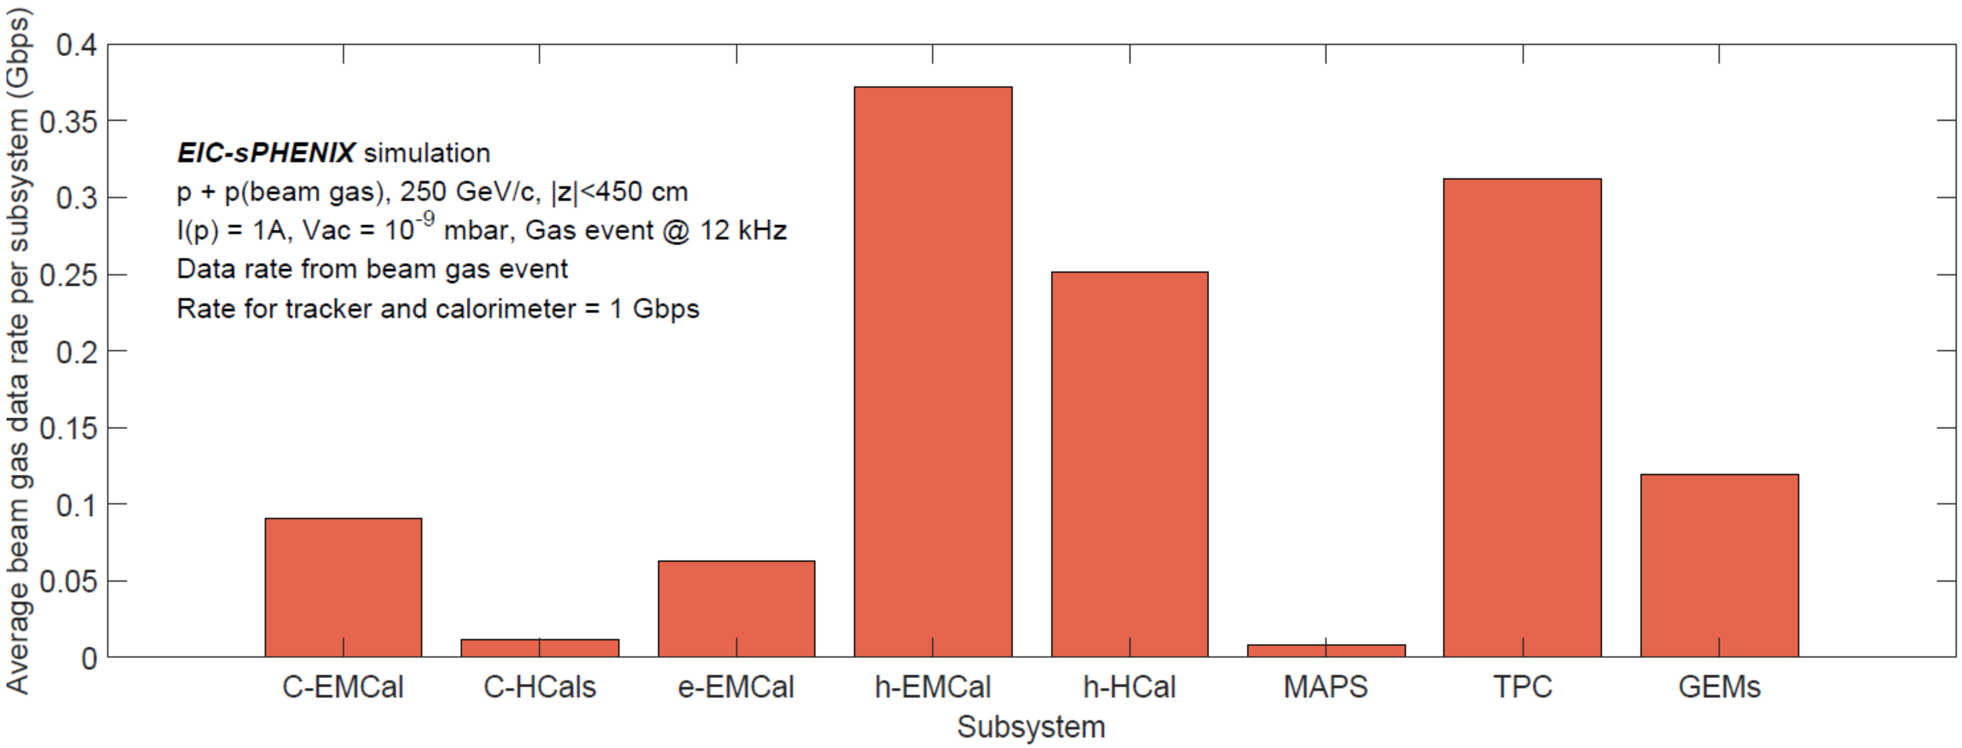

Beam Gas & Synchrotron Radiation Simulations

(Elke Aschenauer, Adam Jaroslav, Zhengqiao Zhang, Deepak Samual, Reynier Cruz Torres)

Synchrotron Radiation

- Studied effect of 5 um gold coating on beam pipe.
- Did not translate these to data volumes because misleadingly high
  - Weighted Montecarlo statistics too low.
  - Thresholds set at assumed minimum detector sensibility, not by desired zero-suppression threshold (used .2KeV for MPGD, .4KeV for MAPS)

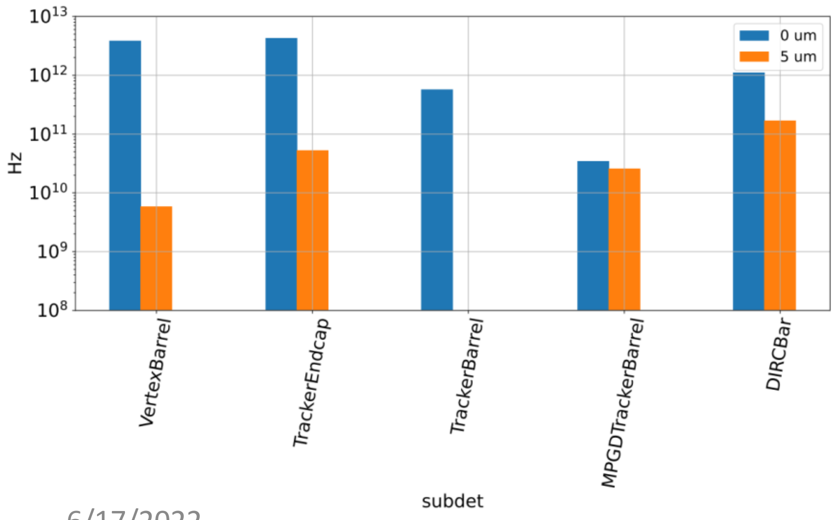

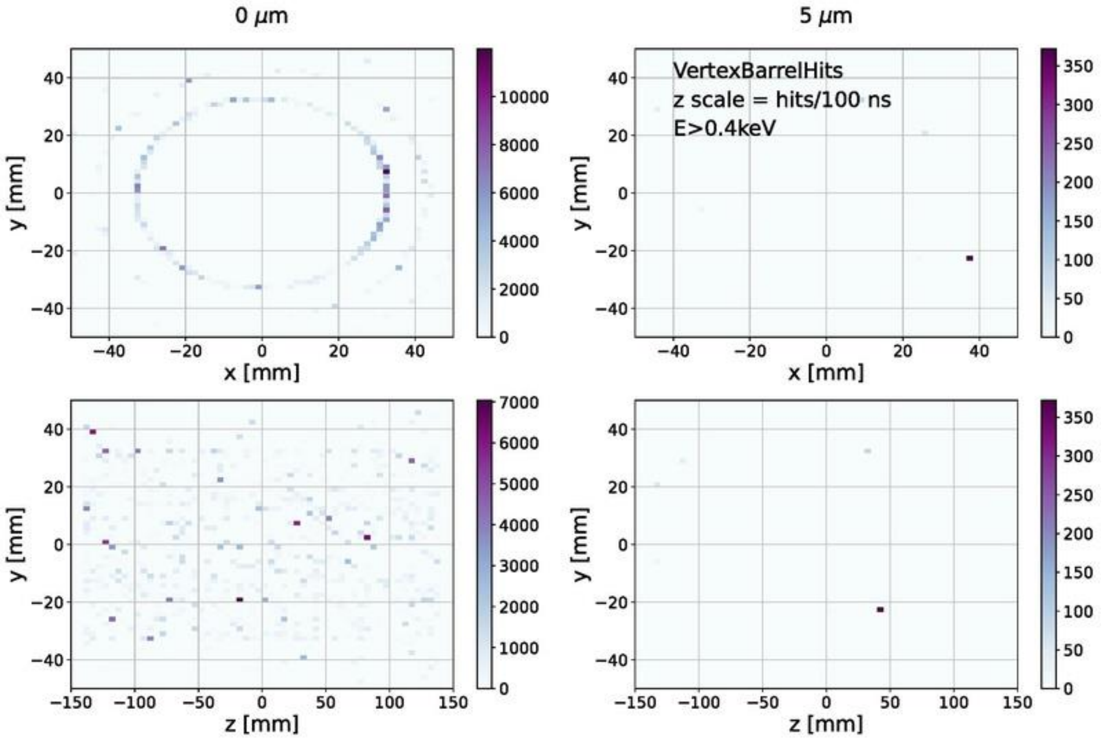

Data rates from YR

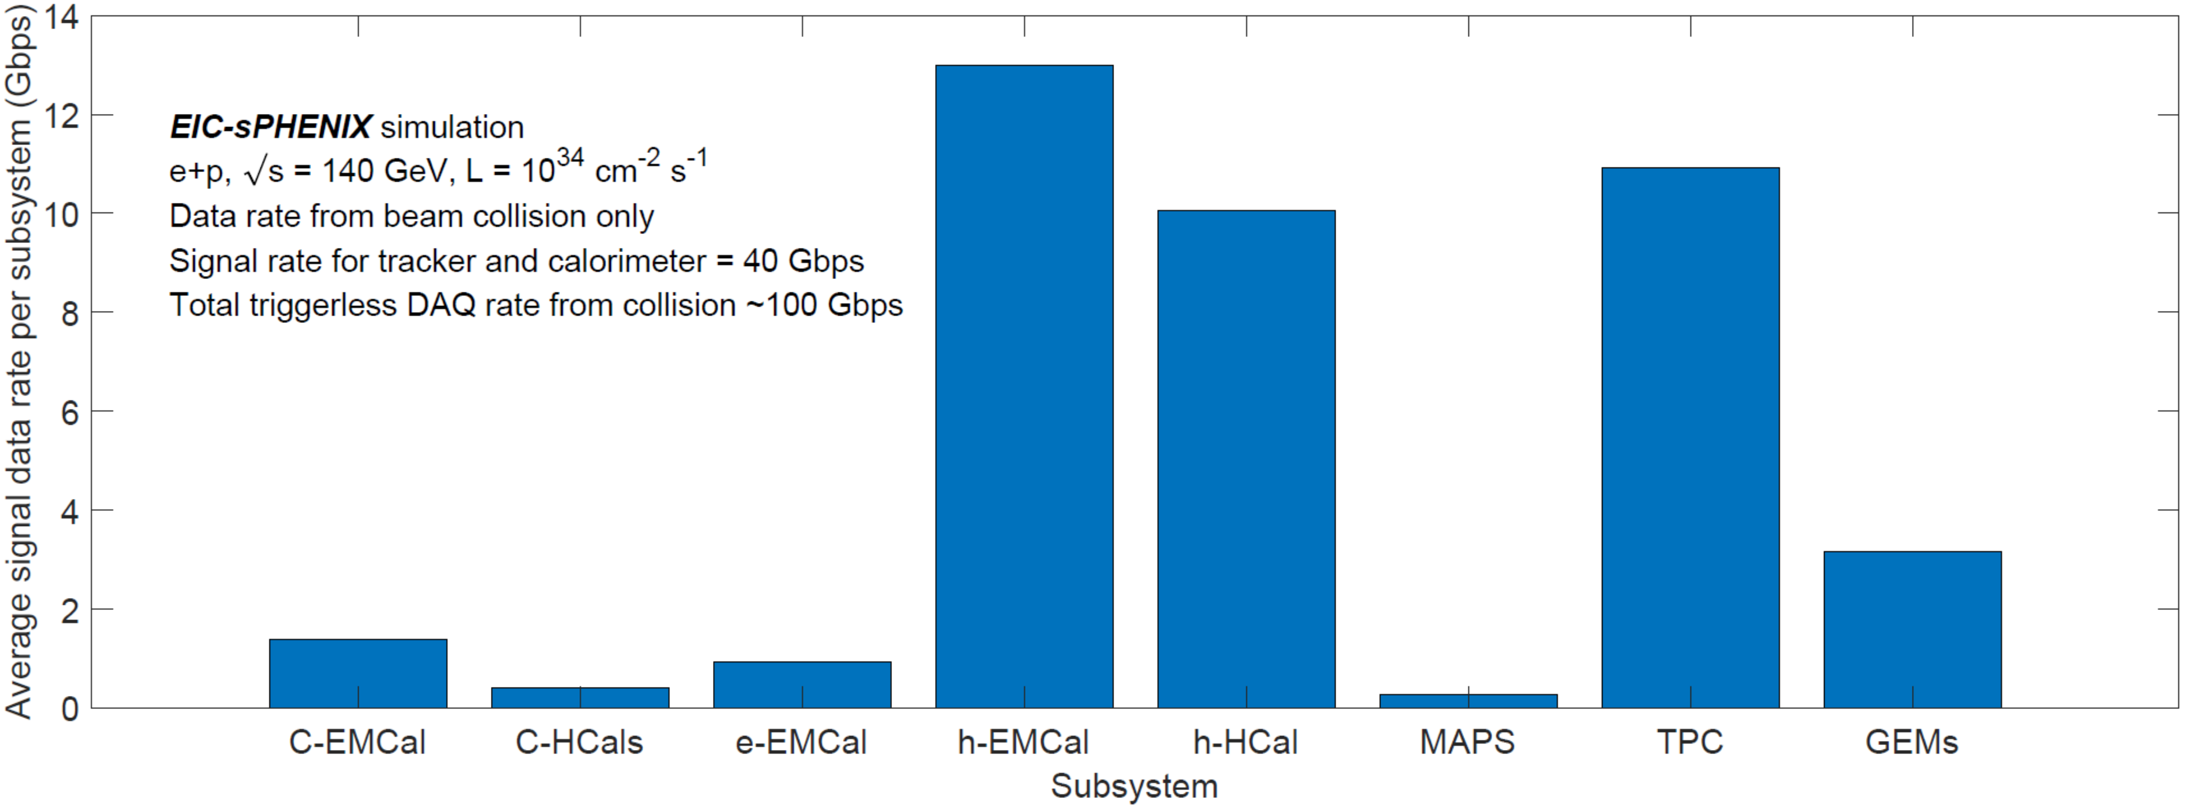

Simulated Data Rates for ATHENA

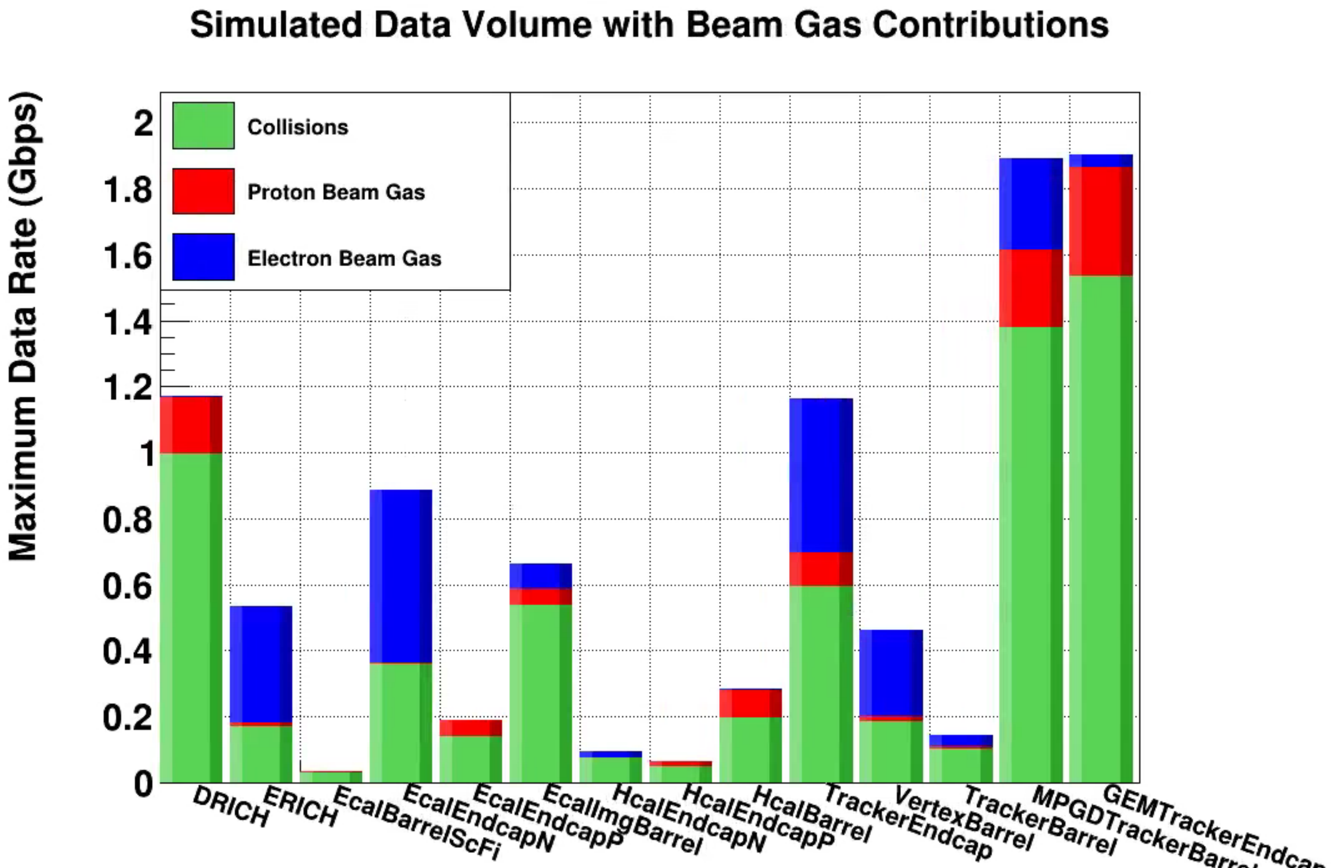

Group as per detector expert table and compare

Detector	Channels	DAQ Input(Gbps)	DAQ Output(Gbps)	DAQ Output (Sim Collision + Beam Gas)
B0 Si	400M	<1	<1	-
B0 ac-lgad	500k	<1	<1	-
RP+Offm+ZDC	700k	<1	<1	-
FB Cal	4k	80	1	[1]
eCal	34k	5	5	1.2
hCal	39k	5.5	5.5	.45
imCal	619M	4	4	.7
Si Tracking	60B	5	5	1.8
Micromegas Tracking	66k	2.6	.6	1.9
GEM Tracking	28k	2.4	.5	1.9
uRWELL Tracking	50k	2.4	.5	-
dRICH	300k	1830	14	1.2
pfRICH	225k	1380	12	.55
DIRC	100k	11	11	[1.2]
TOF	332k	3	.8	[3.5]
Totals		3400	61	13.4

Front End Boards (FEB)

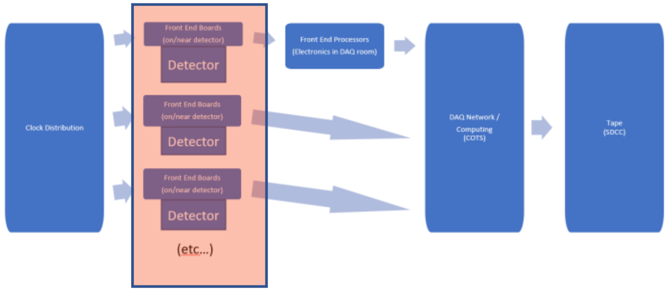

- The collider performance:
  - ~500KHz of collisions
  - ~60-100Gbps zero suppressed data
  - ~15 KB/event
  - ~100 bytes/bunch crossing
- We have an enormous number of channels but the Silicon MAPS readouts test the relevance of the concept of channel.
- Challenging data compression scheme
  - Noise reduction
  - Zero suppression
  - Background elimination

Detector	Readout Technology	Channel Count
Silicon Tracking	Si MAPS	37B
GEM/MMG Layer	GEM	217K
Cylindrical MPGD *	GEM	60M
HP-DIRC	MAP/MT	100-330k
ECAL	SiPM	1.7K
HCAL	SiPM	24K
ECAL imaging	Si MAPS	480M
dRICH	PMT/SiPM	350K
mRICH	PMT/SiPM	330K
B0	Si MAPS	32M + 320K
Off-Momentum	AC-LGAD (eRD24)	750K
Roman Pots	AC-LGAD (eRD24)	500K
ZDC	LGAD + ASIC eRD27	225+366
TOF	AC-LGAD	15M

SBS detector and experiments

- GEn
- GEn RP
- GEp 5
- TDIS

- Trigger rates up to 5 KHz
- Data rates from 1 to 3 GB/s

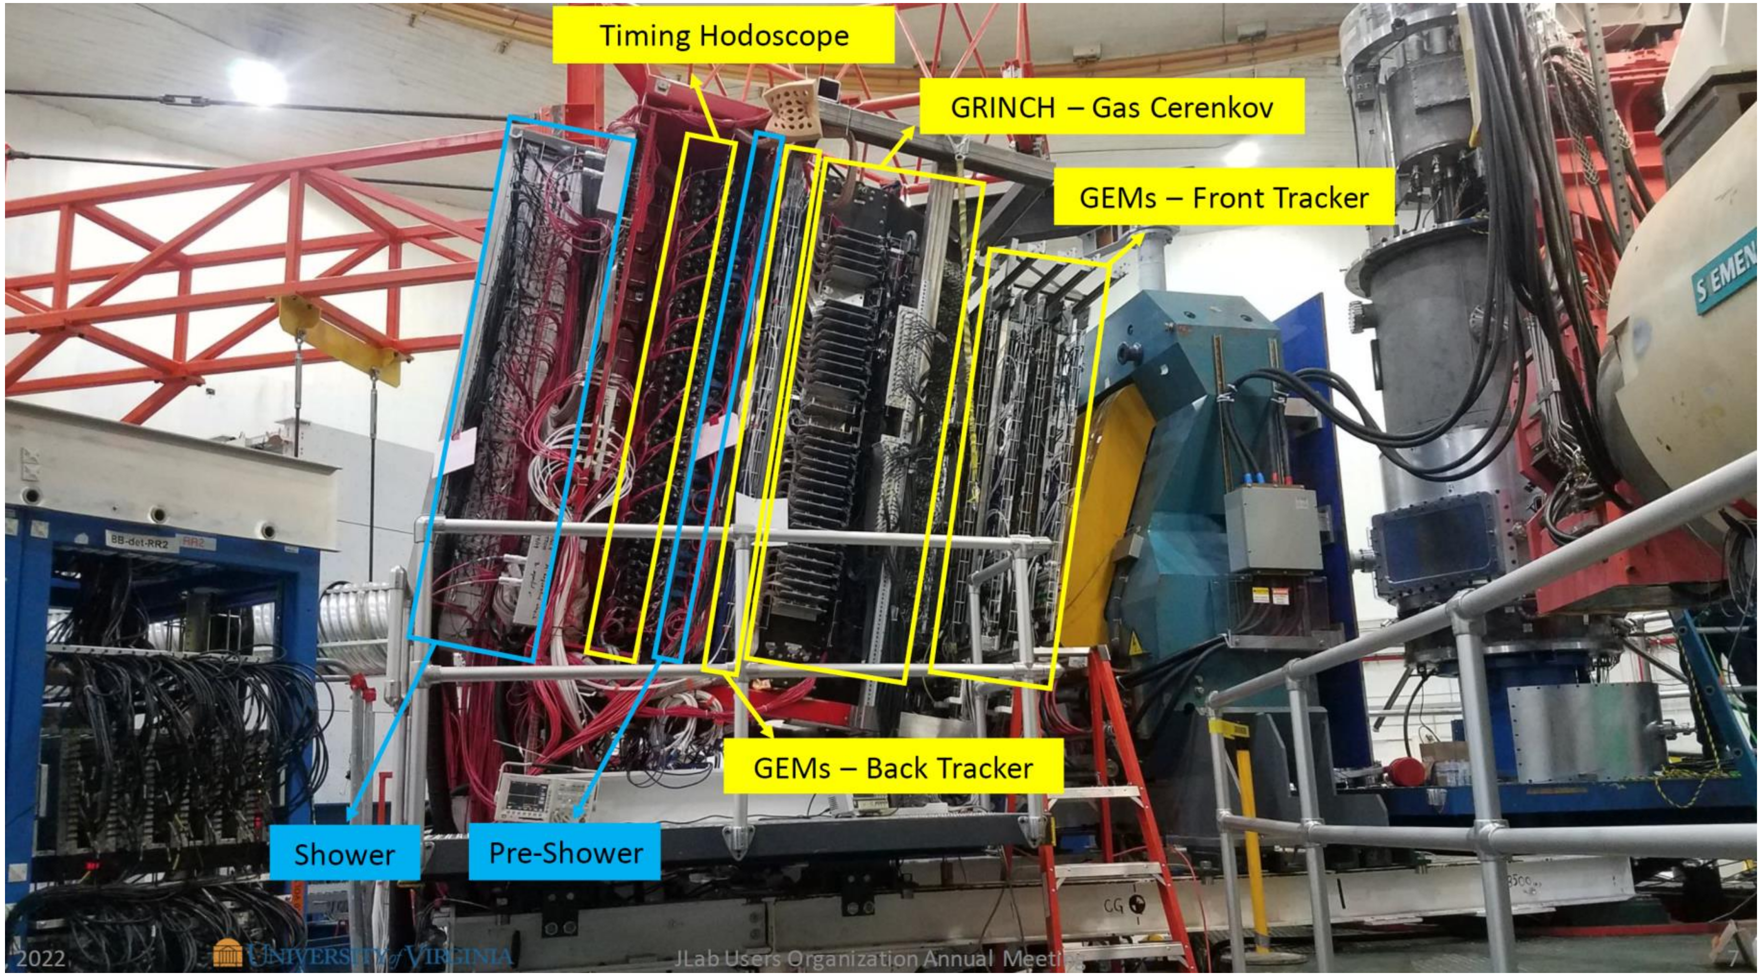

Streaming readout applications

- SBS tracking
  - To go streaming would need to replace APV by other chip (SAMPA or VMM)
  - Increase bandwidth to farm

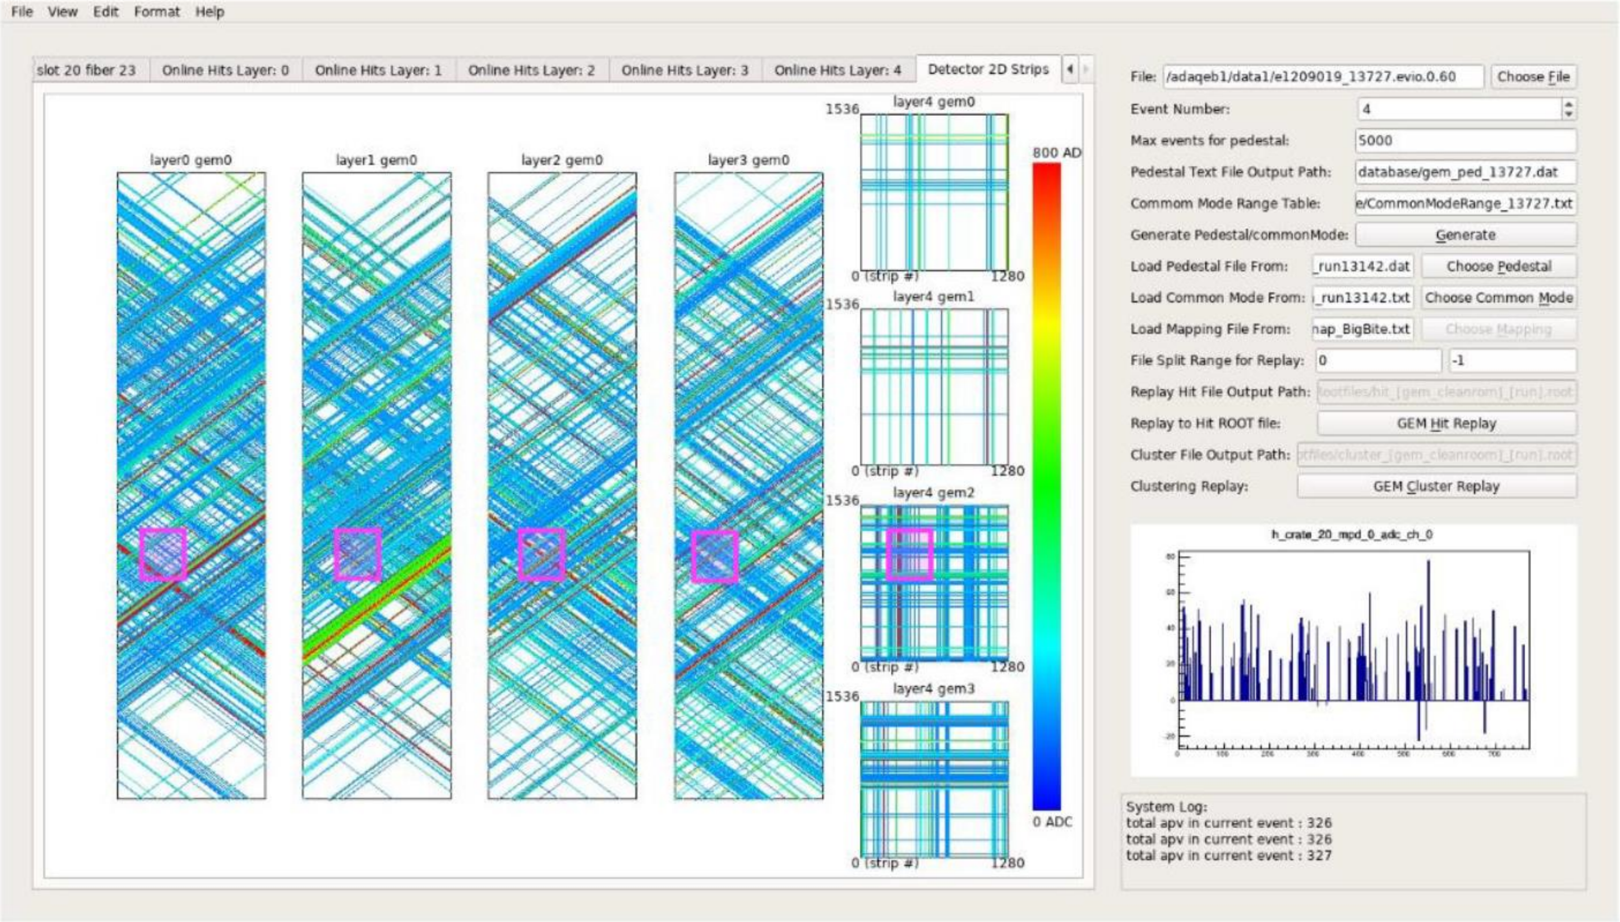

TDIS experiment

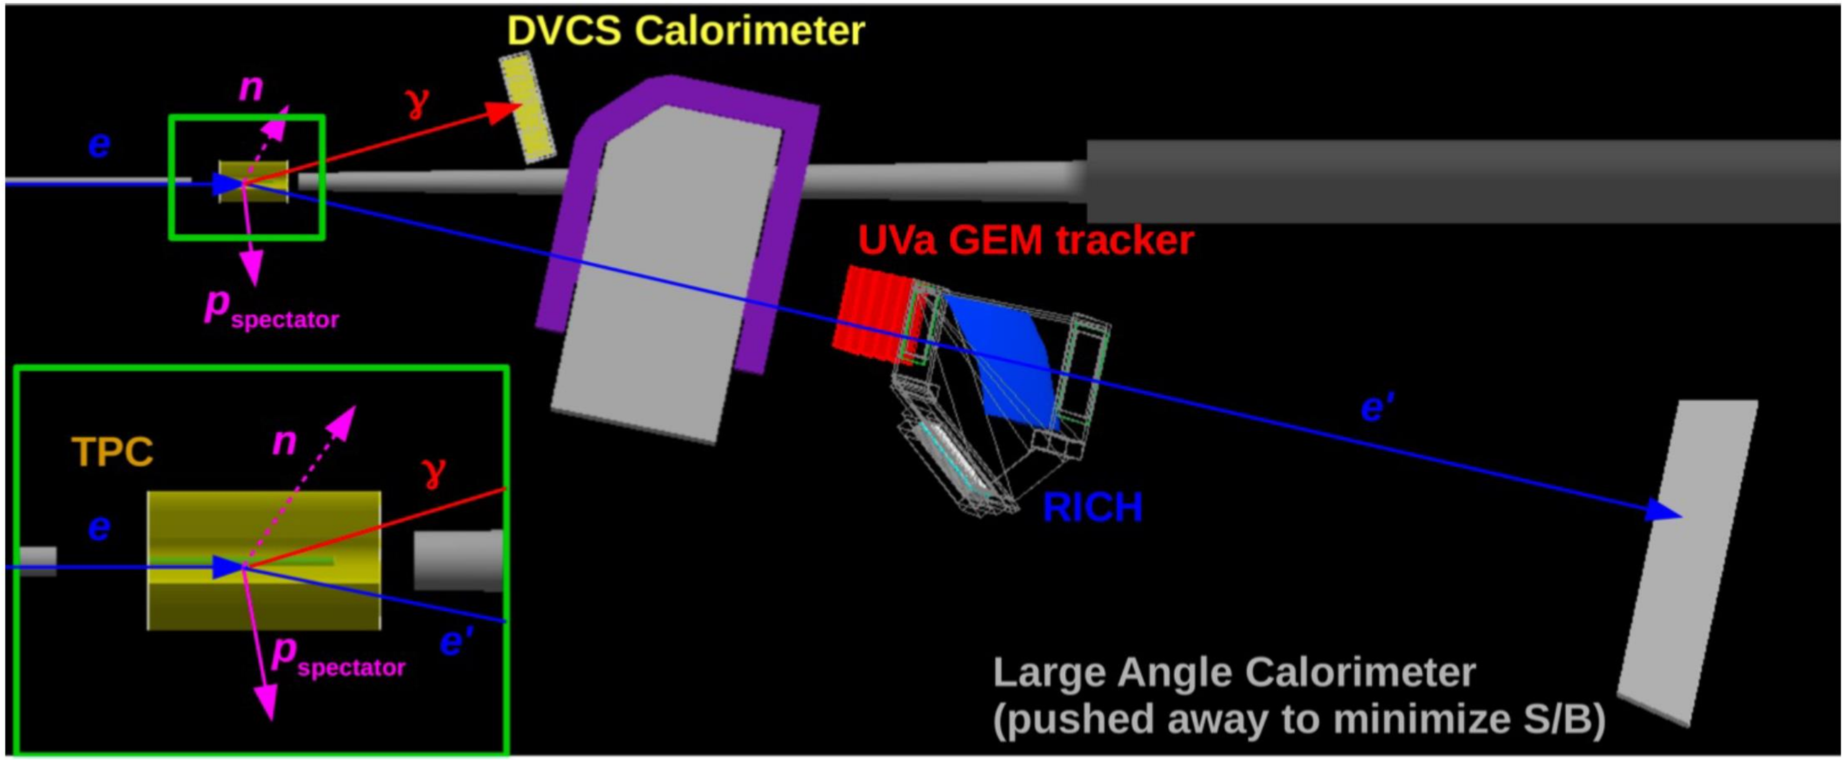

TDIS experiment

[TDIS in Hall A: Experimental layout]

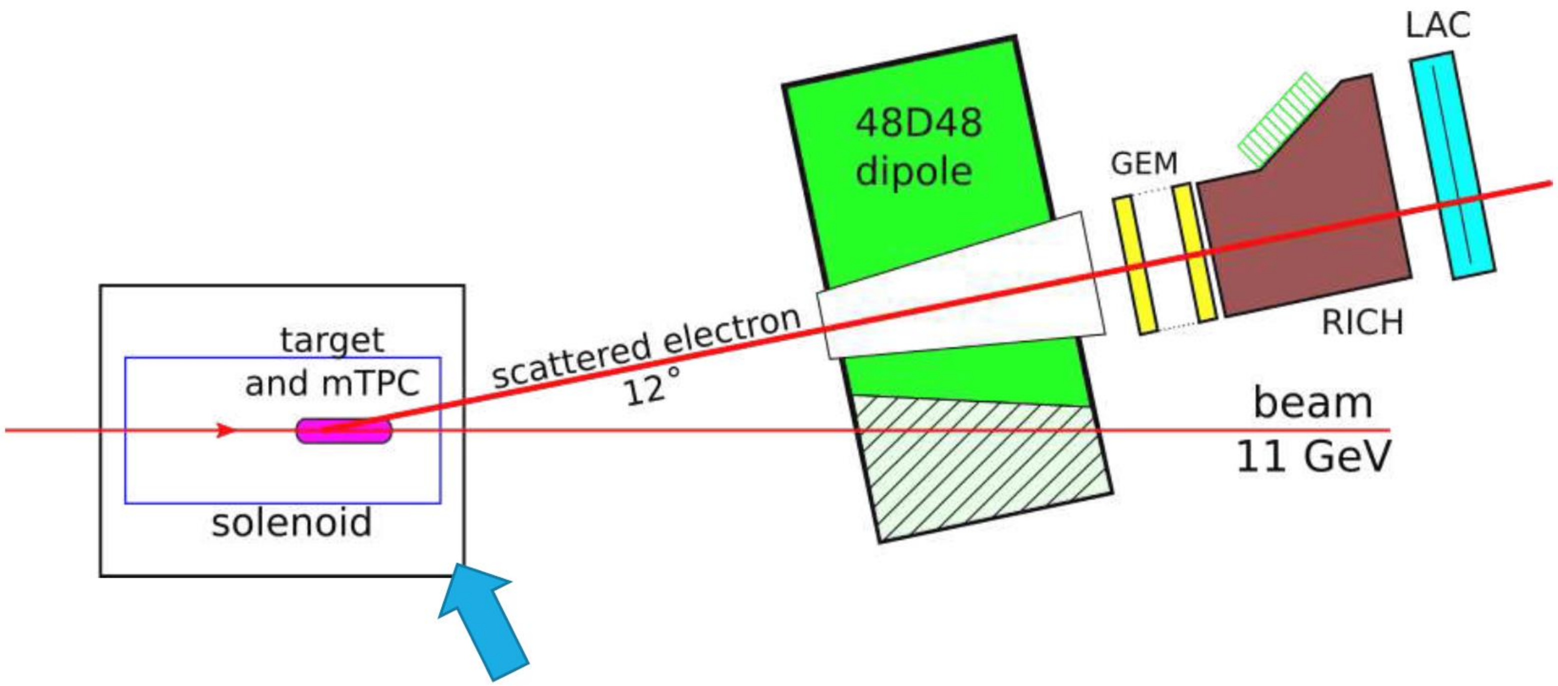

A new detector → mTPC: multi-Time Projection Chamber to measure low-momentum recoil protons

TDIS mTPC

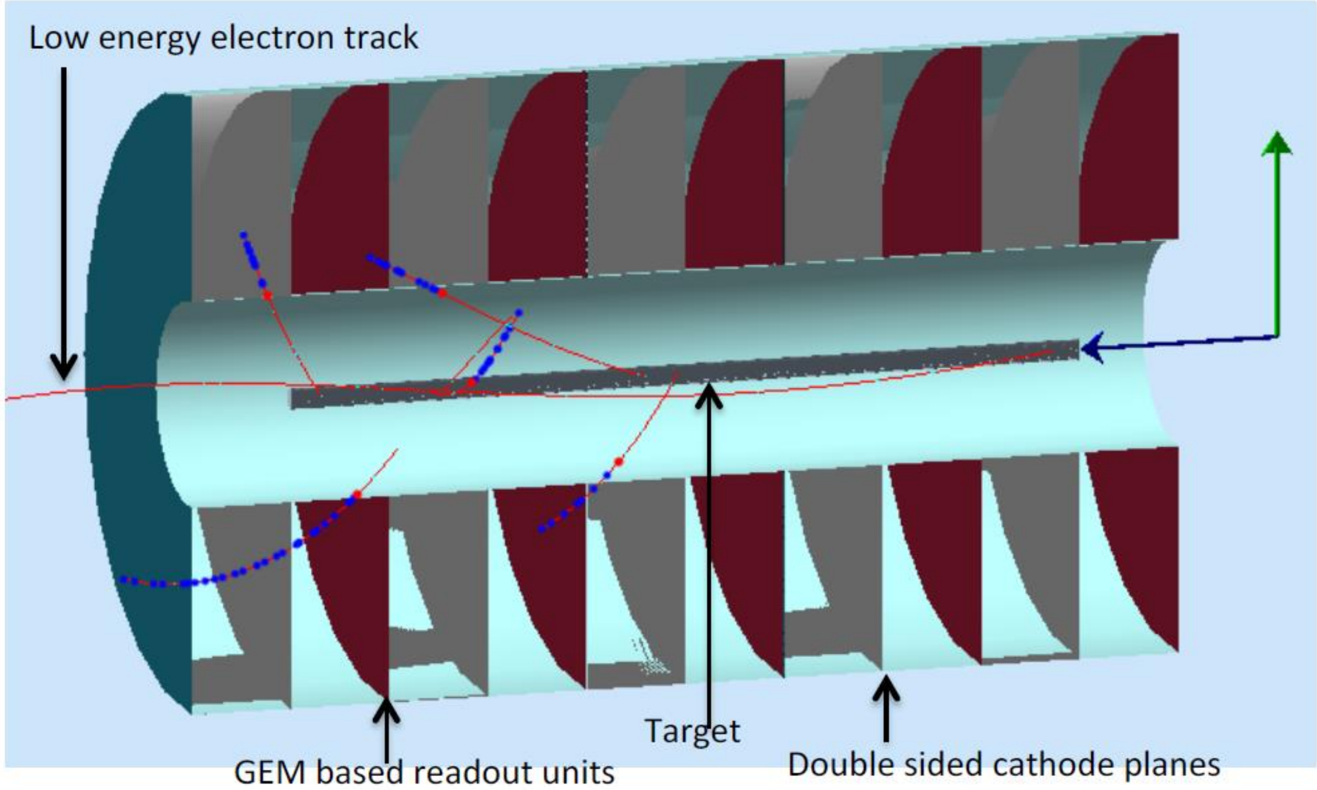

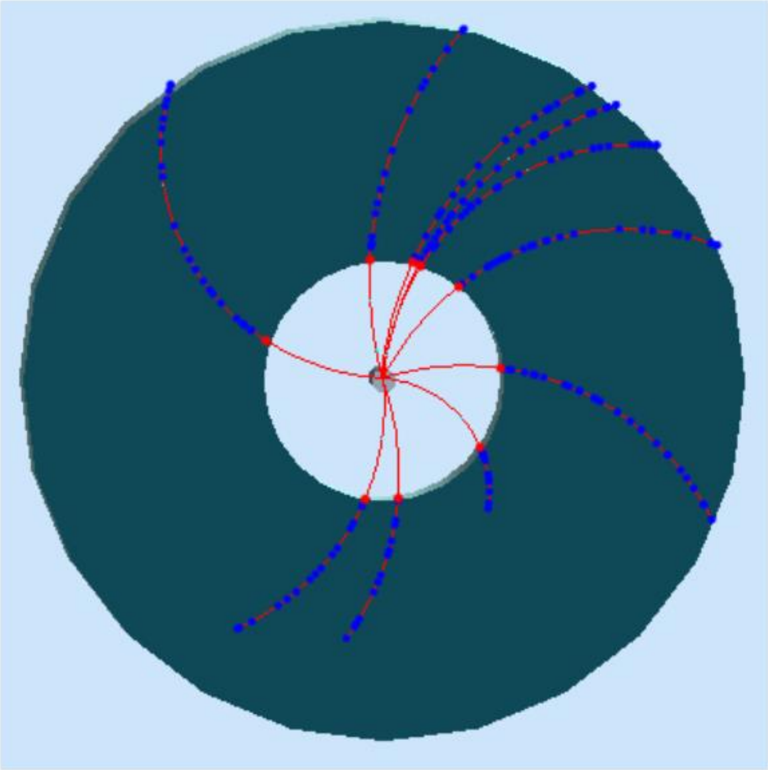

Ptcl tracks

TDIS SAMPA readout

Streaming Data Acquisition

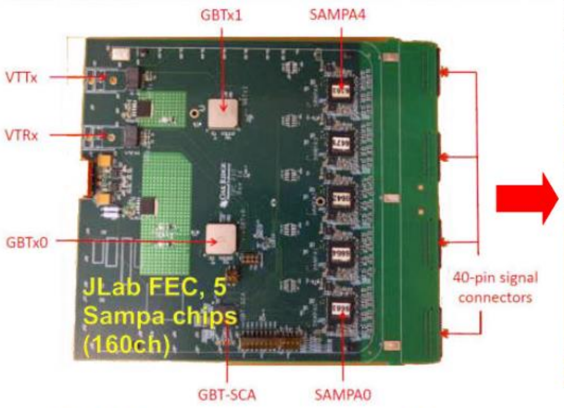

SAMPA - charge-sensitive pre-amp, ADC, DSP (zero-suppresion e.g.)

- Design/prototyping/testing
  - E. Jastrzembski, E. Pooser, G. Heyes (JLab)
- SAMPA chip
  - M. Bregant (U. Sao Paulo) and streaming readout developed for ALICE TPC upgrade

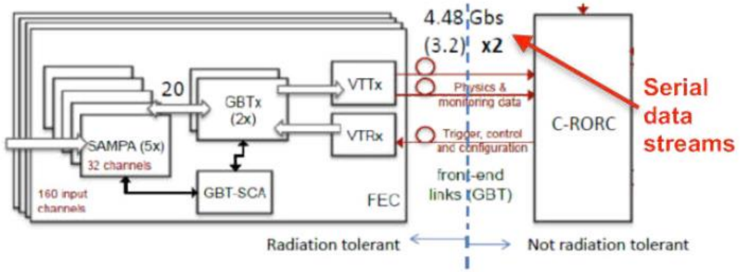

FEC – Front End Card (160 ch / FEC) (5 FEC = 800 ch)
C-RORC – Common Read Out Receiver Card (PCIe)
GBTx – Giga Bit Transceivers
GBT-SCA – GBTx Slow Controls Adapter
VTTx, VTRx – Fiber optic transceivers

Readout Electronics Updates:

- Obtain radiation hard components from CERN
- 2nd generation data transmission and power conversion components
  - lpGBT, VTRX+ (for High Luminosity LHC)
  - bPOL12V, bPOL2V5, linPOL12V (for HL LHC)

mTPC Simulation Status: digitization

Credit: R. Montgomery (University of Glasgow)

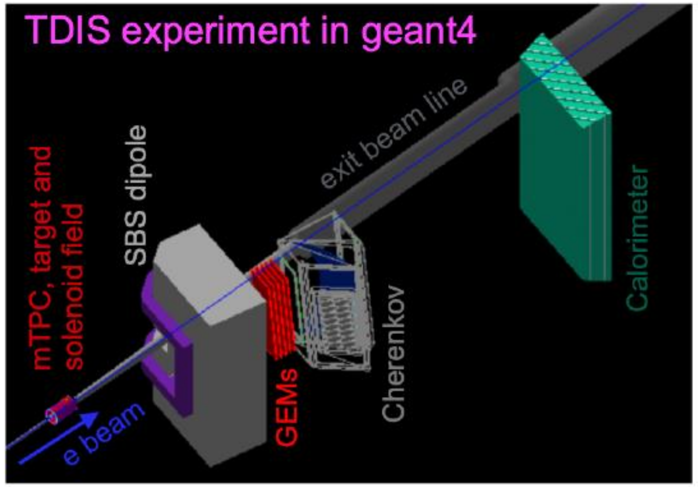

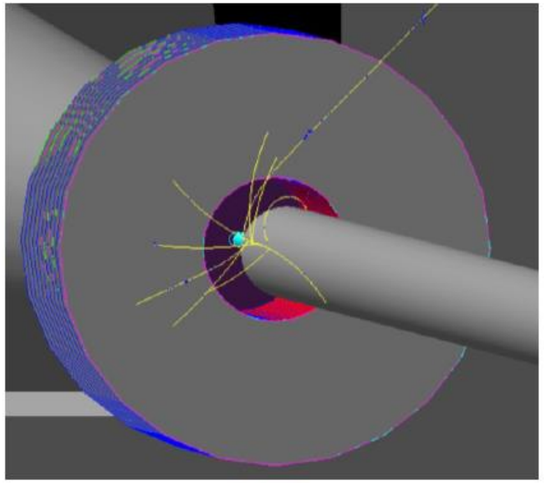

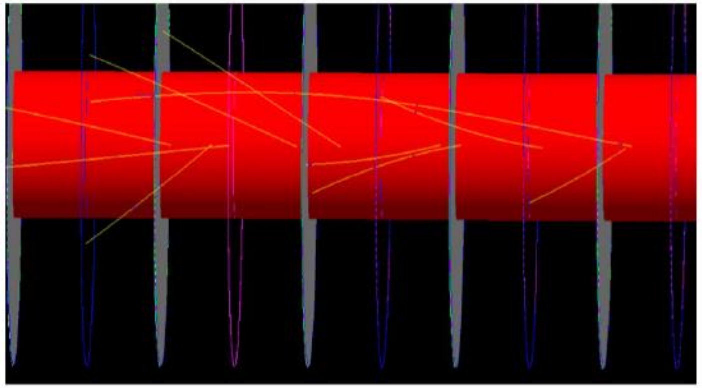

SBS geant4 framework g4sbs is used for simulation studies

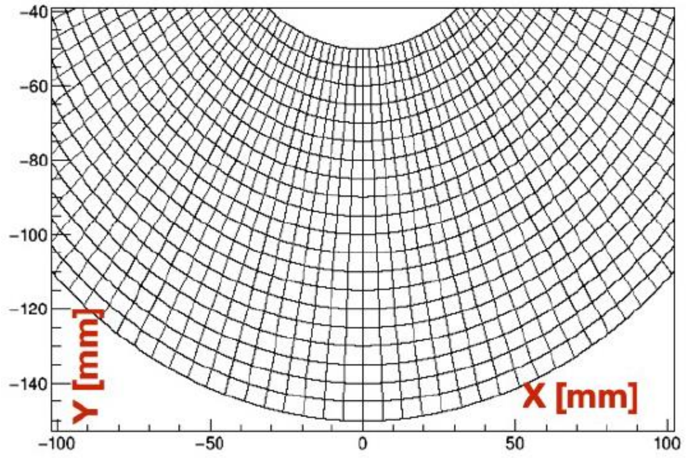

- Example readout pad layout
- 22 rings in radial direction
- 122 pads/ring (area increases with radius)

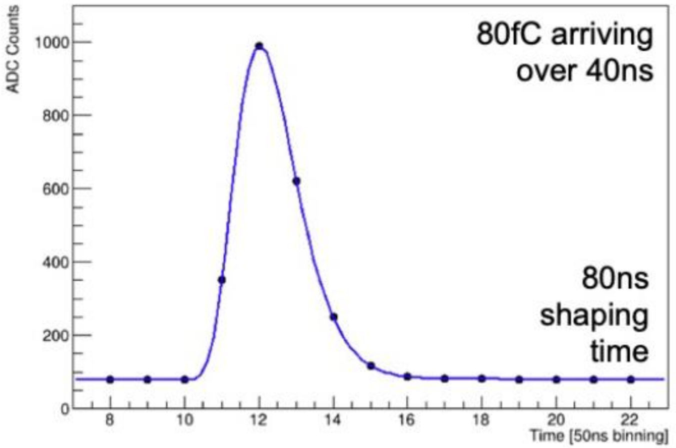

- Markers are sampled points
- Curve is convoluted output pulse shape from SAMPA impulse responses to charge over time window

- Digitisation of signals extracted from mTPC plus SAMPA readout
- Entire chain of signal considered from energy deposition, to charge diffusion, to SAMPA shaping
- Tracking studies using digitised output is underway
- Updated background rate studies are underway

TDIS

- Around 6 KHz trigger rate, 0.8 MHz rate in mTPC
- About 3GB/s without processing
- Pulse fitting reduce to about 1 GB/s

- More information in
- https://indico.jlab.org/event/503/contributions/9162/attachments/7498/10433/Hall_A_TDIS_Collab_02-22_fnlcor.pdf
- TDIS DAQ, 3 September 2020 - Hall A Wiki (jlab.org)

SoLID

- 300 to 500 KHz of trigger rates
- SIDIS if can take 200 kHz single electrons can record all physics (pion + kaons)
- High occupancy in trackers due to photon conversion (10 to 15%)
- Parasitic experiment such as Ay measurement would benefit of streaming DAQ

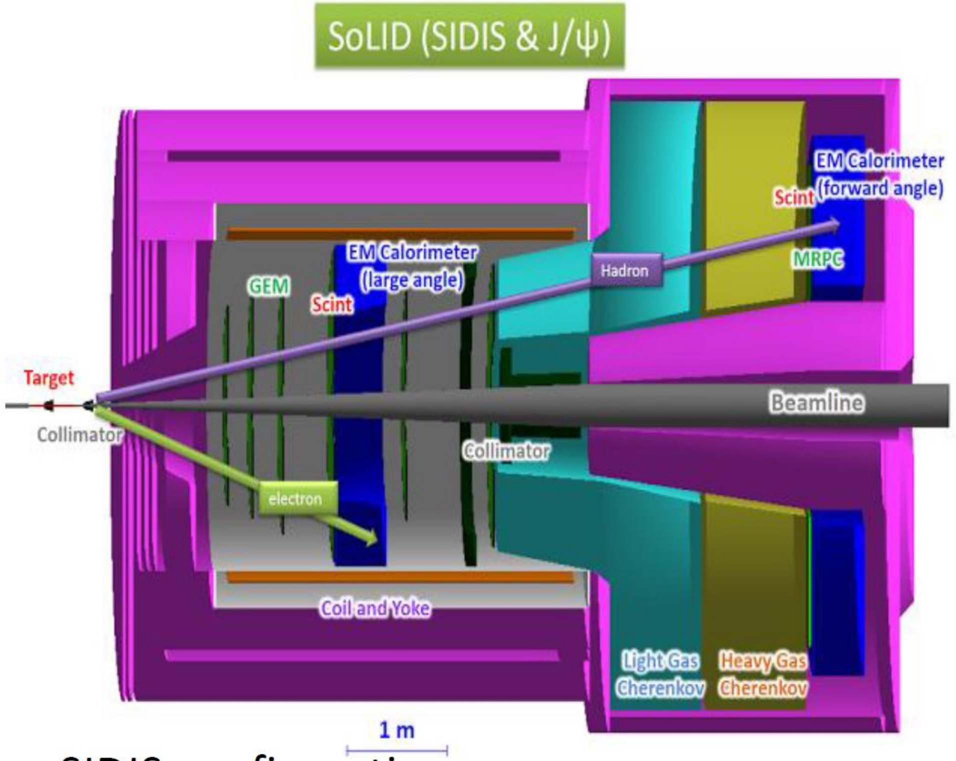

- SIDIS configuration
- Trigger :
  - Electron trigger : Calorimeter+Light Gas Cerenkov
  - Pion trigger
    - Scintillator + calorimeter
  - Main trigger : coincidence eπ
- 100 KHz of coincidence eπ
- J/ψ 30KHz triple coincidence ee^-e^+

The JLab issue / possible strategy

- Our detector/experiments (5 to 50 M$) are smaller scale than CERN (ALICE, ATLAS 1B$?)
- Cannot afford 20 M$ (rough cost for LHCb Allen system) investment for a single experiment
- Some progress on high performance computing but not clear what would be available for a trigger on streaming data
- Ideally make a budget request for a generic online trigger that could be used for triggered or streaming data by Hall A,B,C and D
- If ressources secured, can envision SoLID or TDIS in streaming mode

Conclusion

- Technologies advances in electronics, network and storage makes us closer and closer to full streaming readout DAQ
- Streaming systems implemented at CERN like LHCb
  - Very high rate capability 30 to 40 MHz
  - Online reconstruction allows continuous calibration and data monitoring
  - Full reconstruction significantly reduce the results turn around
  - Still triggered in the end (only signal of interest recorded)
  - Today streaming could be a bit more expensive but
- EIC will be a streaming readout DAQ system
  - First opportunity to be able to record streams since luminosity is low
  - Wealth of physics opportunities
- Streaming can be implemented at JLab and would be useful for Hall A/C, B ,D but mostly matter of funding vs physics yield
- One strategy possible for Hall A/C is presented to have a streaming system for several experiments
- More information at streaming workshop readout readout

Streaming Readout X (17-19 May 2022): Timetable · Jefferson Lab Indico (Indico) (jlab.org)