Data Volumes & ePIC DAQ streaming

Jeff Landgraf

8/21/2023

Electron-Ion Collider

ePIC Data Volumes

- Data volumes / assumptions
- Examples of DAQ processing

Summary of Channel Counts

Detector Group	Channels					RDO	Fiber	DAM	Data Volume (RDO) (Gb/s)	Data Volume (To Tape) (Gb/s)
	MAPS	AC-LGAD	SiPM/PMT	MPGD	HRPPD					
Tracking	36B			202k		872	1744	24	27	26
Calorimeters	88M		123k			258	556	10	502	27
Far Forward	300M	2.3M	170k			178	492	5	15	8
Far Backward	146M		2k			50	100	6	150	1
PID		7.8M	320k		140k	241	523	39	2628	36
TOTAL	36.5B	10.1M	615k	202k	140k	1599	3415	84	3,322	98

Summary of Data Flow

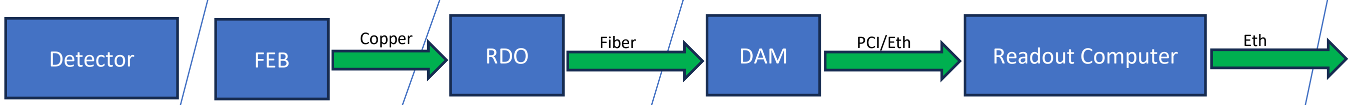

100% Occupancy	27.5-1760 Pb/sec

Aggregate	3.4 Tb/sec
Noise	3 Tb /sec
Signal from Physics + Background	400 Gb / sec

Aggregate	3.4 Tb/sec
Per fiber (Avg)	1.85 Gb/sec

Aggregate	99 Gb/sec
Collision Signal	38 Gb/sec
Synchrotron Rad	.01 Gb/sec
Electron Beam	4 Gb/sec
Hadron Beam	22 Gb/sec
Noise	35 Gb/sec

Summary of Channel Counts

Detector Group	Channels					RDO	Fiber	DAM	Data Volume (RDO) (Gb/s)	Data Volume (To Tape) (Gb/s)
	MAPS	AC-LGAD	SiPM/PMT	MPGD	HRPPD					
Tracking	36B			202k		872	1744	24	27	26
Calorimeters	88M		123k			258	556	10	502	27
Far Forward	300M	2.3M	170k			178	492	5	15	8
Far Backward	146M		2k			50	100	6	150	1
PID		7.8M	320k		140k	241	523	39	2628	36
TOTAL	36.5B	10.1M	615k	202k	140k	1599	3415	84	3,322	98

Summary of Data Flow

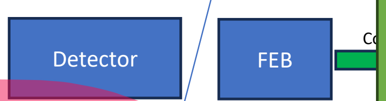

100% Occupancy	27.5-1760 Pb/sec

Aggregate	3.4 Tb/sec
Noise	3 Tb /sec
Signal from Physics + Background	400 Gb / sec

Aggregate	
Per fiber (Avg)	1.85 Gb/sec

Noise	35 Gb/sec

Sparse readout implies lots of header information

1. Geography
    1. Which Channel?
    2. Which ASIC?
    3. Which RDO?
    4. Which DAM?
    5. Which Time?
    6. What is the ADC / TOA / TOT?
2. Assume for Data Volume Calculations 64 bits / hit

→ 1760 Pb/sec assumes sparse readout of every channel
→ 27.5 Pb/sec assumes headers collapsed because not sparse!

Real work was all done by the Backgrounds Group

https://wiki.bnl.gov/EPIC/index.php?title=Background

I had to make many assumptions.

To find hit rates:

18x275 DIS was available
But scaled for 83kHz

I assumed “same” kinematics, but scaled to 500kHz rate…

cross-section	5x41 GeV	5x100 GeV	10x100 GeV	10x275 GeV	18x275 GeV
DIS ep	28.5ub	35ub	41ub	50ub	54ub
hadron beam (p) gas	77.3mb	76.8mb	76.8mb	78.5mb	78.5mb
electron beam gas	622.158 +/- 0.036 mb	622.158 +/- 0.036 mb	699.393 +/- 0.041 mb	699.393 +/- 0.041 mb	768.343 +/- 0.049 mb
DIS eA	ub	ub	ub	/	/
hadron beam (Au) gas	3418mb	3440mb	3440mb	/	/

This table shows the rates for electron+proton beam configurations:

Electron beam-gas rates consider larger region of -5 to +15 meters along the IP, hadron beam-gas rates consider region of -5.5 to +5 meters.

rates in kHz	5x41 GeV	5x100 GeV	10x100 GeV	10x275 GeV	18x275 GeV	Vacuum
DIS ep	12.5 kHz	129 kHz	184 kHz	500 kHz	83 kHz	
hadron beam gas	12.2kHz	22.0kHz	31.9kHz	32.6kHz	22.5kHz	10000Ahr
	131.1kHz	236.4kHz	342.8kHz	350.3kHz	241.8kHz	100Ahr
electron beam gas	2181.97 kHz	2826.38 kHz	3177.25 kHz	3177.25 kHz	316.94 kHz	10000Ahr
DIS eA	kHz	kHz				
hadron beam (Au) gas	7.36kHz	10.3kHz				
	79.1kHz	110.7kHz				

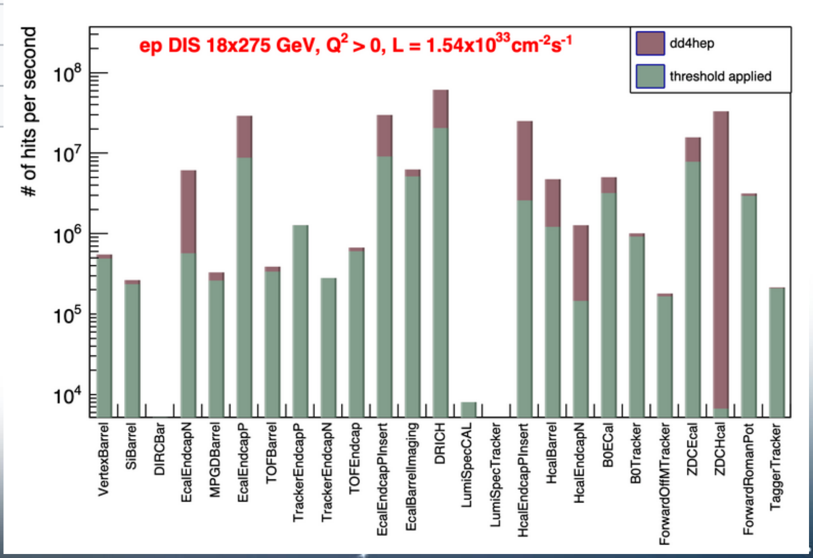

Summary of Channel Counts

Sub-detector	Threshold	Integration time	Sub-detector	Threshold	Integration time
VertexBarrel	0.65 keV	2 μs	EcalEndcapP	3.0 MeV	5 ns
SiBarrel	0.65 keV	2 μs	TOFBarrel	0.5 keV	50 ps
EcalEndcapN	5.0 MeV	5 ns	TrackerEndcap	0.65 keV	50 ps
MPGDBarrel	0.25 keV	20 ns	DIRCBar	0.2 p.e.	50 ps
EcalEndcapPInsert	3.0 MeV	5 ns	TOFEndcap	0.5 keV	50 ps
LFHCAL	500 keV	25 ns	PFRICH	0.5 p.e.	50 ps
HcalEndcapPInsert	500 keV	25 ns	DRICH	0.5 p.e.	50 ps
HcalBarrel	75 keV	25 ns	EcalBarrelScFi	2.5 MeV	5 ns
B0ECal	1 MeV	5 ns	HcalEndcapN	170 keV	25 ns
B0Tracker	1.0 keV	40 ps	ZDCEcal	1 MeV	5 ns
ForwardOffMTracker	1.0 keV	40 ps	ZDCHcal	100 MeV	25 ns
TaggerTracker	1.0 keV	5 ns	ForwardRomanPot	1.0 keV	40 ps

The simulations had nominal thresholds, but these are known to be wrong in some cases. For example, we do not believe the Synchrotron Radiation Numbers. This is being resolved but not yet available.

M	Data Volume (RDO) (Gb/s)	Data Volume (To Tape) (Gb/s)
4	27	26
0	502	27
	15	8
	150	1
9	2628	36
4	3,322	98

mputer

Eth

Aggregate	3.4 Tb/sec
Noise	3 Tb /sec
Signal from Physics + Background	400 Gb / sec

Aggregate	3.4 Tb/sec
Per fiber (Avg)	1.85 Gb/sec

Aggregate	99 Gb/sec
Collision Signal	38 Gb/sec
Synchrotron Rad	.01 Gb/sec
Electron Beam	4 Gb/sec
Hadron Beam	22 Gb/sec
Noise	35 Gb/sec

How many channels are hit by an incident particle?

- Charge Sharing (# of neighboring channels above threshold from single hit)
- Hit duration (Pre-amp shaping ensures that multiple time-bins are hit to allow for better time resolution)

Examples from Digitization worksheet:
https://docs.google.com/spreadsheets/d/1s8oXj36SqIh7TJeHFH89gQ_ayU1_SVEpWQNkx6sETKs/edit?usp=sharing

I used many sources, including digitization worksheet
DAQ WG presentations from detector experts
And my own "educated" guesses. Spreadsheet with these guesses is available at
https://indico.bnl.gov/event/20133/

Analog Features

Shaping time (FWHM)	2 us (target value)	Total hit time resolution	Charge sharing
			1.2 at 90deg
			1.2 at 90deg
~500 ns	~20 ns	~20 ns	3-5
5 ns		30 ps	2
5 ns		25 ps	2*2 in x*y
100 ns	few ns	few ns	yes
1 us			1 (maybe consider 1-2)
100 ns	few ns (4 ns for GlueX)	Few ns	By default
100 ns	few ns	few ns	yeah (1 particle ==> N tower cluster)
25ns	25ns		

100% Occupancy	27.5-1760 Pb/sec

Aggregate	3.4 Tb/sec
Noise	3 Tb /sec
Signal from Physics + Background	400 Gb / sec

Aggregate	3.4 Tb/sec
Per fiber (Avg)	1.85 Gb/sec

Aggregate	99 Gb
Collision Signal	38 Gb
Synchrotron Rad	.01 Gb/sec
Electron Beam	4 Gb/sec
Hadron Beam	22 Gb/sec
Noise	35 Gb/sec

“Noise”

I took from Digitization worksheet if available:
https://docs.google.com/spreadsheets/d/1s8oXj36SqIh7TJeHFH89gQ_ayU1_SVEpWQNkx6sETKs/edit?usp=sharing

Otherwise from DAQ WG presentations from detector experts.

Generally, I assume “noise” was uncorrelated electronics noise, so is susceptible to removal via cluster finding.

I also included the worse case dark current numbers from dRICH

M	Data Volume (RDO) (Gb/s)	Data Volume (To Tape) (Gb/s)
	27	26
	502	27
	15	8
	150	1
	2628	36
	3,322	98

Summary of Data Flow

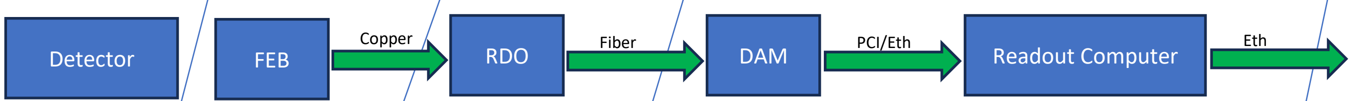

100% Occupancy	27.5-1760 Pb/sec

Aggregate	3.4 Tb/sec
Noise	3 Tb /sec
Signal from Physics + Background	400 Gb / sec

Aggregate	3.4 Tb/sec
Per fiber (Avg)	1.85 Gb/sec

Aggregate	99 Gb/sec
Collision Signal	38 Gb/sec
Synchrotron Rad	.01 Gb/sec
Electron Beam	4 Gb/sec
Hadron Beam	22 Gb/sec
Noise	35 Gb/sec

Summary of Channel Counts

Detector Group	Channels					RDO	Fiber	DAM	Data Volume (RDO) (Gb/s)	Data Volume (To Tape) (Gb/s)
	MAPS	AC-LGAD	SiPM/PMT	MPGD	HRPPD					
Tracking	36B			202k		872	1744	24	27	26
Calorimeters	88M		123k			258	556	10	502	27
Far Forward	300M	2.3M	170k			178	492	5	15	8
Far Backward	146M		2k			50	100	6	150	1
PID		7.8M	320k		140k	241	523	39	2628	36
TOTAL	36.5B	10.1M	615k	202k	140k	1599	3415	84	3,322	98

Summary of Data Flow

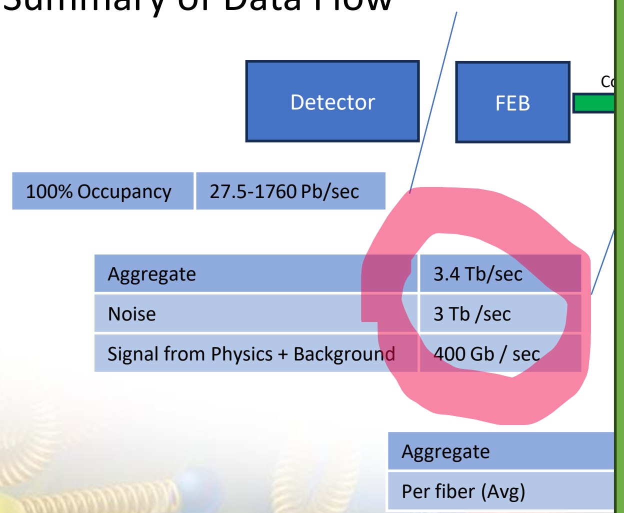

Significant Filtering Challenges:

- At RDO level, "noise" dominates signal by factors up to x7.5
- dRICH – SiPMs requiring single photon sensitivity
  - Dark currents increase with radiation damage
  - Expect several years before annealing necessary to reduce dark currents
- Electron Bremsstrahlung in Far Backward
  - Bremsstrahlung will produce up to 18 particles / BX.
  - Signal needs to be summarized, but full data is only needed in conjunction with central detector collisions
- Main Strategy for dRICH/Far Backward
  - Supply enough bandwidth to account for maximum data volume to the DAM boards
  - Apply cross-detector correlation filter in nearby detectors in DAM / Readout computers to reduce recorded data volume

Summary of Chan…

Detector Group	MAPS								Data Volume (RDO) (Gb/s)	Data Volume (To Tape) (Gb/s)
Tracking	36B								27	26
Calorimeters	88M								502	27
Far Forward	300M								15	8
Far Backward	146M		2k			50	100	6	150	1
PID		7.8M	320k		140k	241	523	39	2628	36
TOTAL	36.5B	10.1M	615k	202k	140k	1599	3415	84	3,322	98

Data Filtering

- Cluster finding accounts for the balance of the calculated noise reduction, and for the reduced signal data volume.
- Additional methods are being considered including AI/ML techniques for pattern recognition and/or data compression
- These functions will be performed in the DAM boards and the Online computing farm

Summary of Data Flow

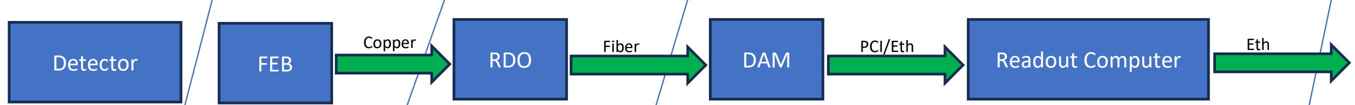

100% Occupancy	27.5-1760 Pb/sec

Aggregate	3.4 Tb/sec
Noise	3 Tb /sec
Signal from Physics + Background	400 Gb / sec

Aggregate	3.4 Tb/sec
Per fiber (Avg)	1.85 Gb/sec

Aggregate	99 Gb/sec
Collision Signal	38 Gb/sec
Synchrotron Rad	.01 Gb/sec
Electron Beam	4 Gb/sec
Hadron Beam	22 Gb/sec
Noise	35 Gb/sec

Summary of Channel Counts

Detector Group	Channels					RDO	Fiber	DAM	Data Volume (RDO) (Gb/s)	Data Volume (To Tape) (Gb/s)
	MAPS	AC-LGAD	SiPM/PMT	MPGD	HRPPD					
Tracking	36B			202k		872	1744	24	27	26
Calorimeters	88M		123k			258	556	10	502	27
Far Forward	300M	2.3M	170k			178	492	5	15	8
Far Backward	146M		2k			50	100	6	150	1
PID		7.8M	320k		140k	241	523	39	2628	36
TOTAL	36.5B	10.1M	615k	202k	140k	1599	3415	84	3,322	98

Summary of Data Flow

After online processing we expect roughly equal contributions from collision hits, beam background hits, and noise.

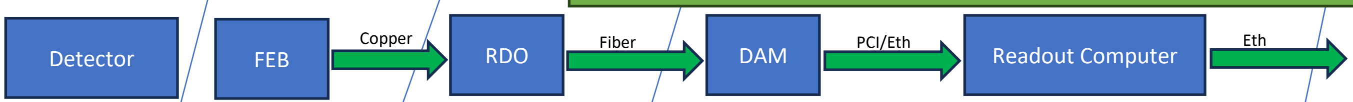

100% Occupancy	27.5-1760 Pb/sec

Aggregate	3.4 Tb/sec
Noise	3 Tb /sec
Signal from Physics + Background	400 Gb / sec

Aggregate	3.4 Tb/sec
Per fiber (Avg)	1.85 Gb/sec

Aggregate	99 Gb/sec
Collision Signal	38 Gb/sec
Synchrotron Rad	.01 Gb/sec
Electron Beam	4 Gb/sec
Hadron Beam	22 Gb/sec
Noise	35 Gb/sec

Summary of Channel Counts

Next Steps:

1. Get thresholds under control
2. Get to the ASIC / FEB granularity to study bottlenecks
3. Automate the data volume calculations, both to make the assumptions explicit, and to track detector changes
4. Need to track the development of detectors to ensure that the noise estimates are correct

						RDO	Fiber	DAM	Data Volume (RDO) (Gb/s)	Data Volume (To Tape) (Gb/s)
						872	1744	24	27	26
						258	556	10	502	27
						178	492	5	15	8
						50	100	6	150	1
PID		7.8M	320k		140k	241	523	39	2628	36
TOTAL	36.5B	10.1M	615k	202k	140k	1599	3415	84	3,322	98

Summary of Data Flow

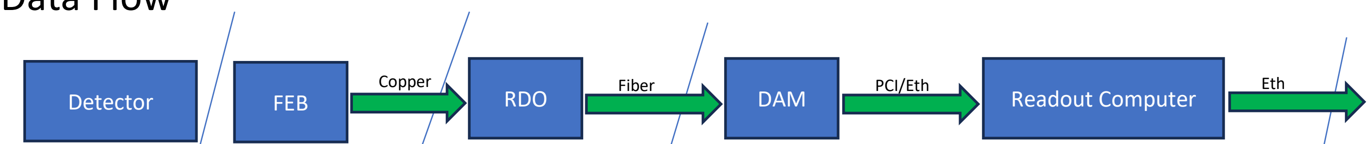

100% Occupancy	27.5-1760 Pb/sec

Aggregate	3.4 Tb/sec
Noise	3 Tb /sec
Signal from Physics + Background	400 Gb / sec

Aggregate	3.4 Tb/sec
Per fiber (Avg)	1.85 Gb/sec

Aggregate	99 Gb/sec
Collision Signal	38 Gb/sec
Synchrotron Rad	.01 Gb/sec
Electron Beam	4 Gb/sec
Hadron Beam	22 Gb/sec
Noise	35 Gb/sec

DAQ Computing

- Time Frames (~1ms)
  - Up to ~500 events
  - ~10MB output data
  - ~3.4MB from RDO average / DAM
  - ~100kB to Tape average / DAM

- Routing data
- Formatting data
- Processing data
  - DAM FPGA & CPUs
  - Cluster finding
  - Software triggering
  - Sanity Checkers
  - QA Monitoring
  - Metadata
  - Slow controls integration
- One goal we should have is to ensure transparency and appropriate control of the algorithms used for all algorithms that might impact physics

- Scalers / continuously running DAQ components

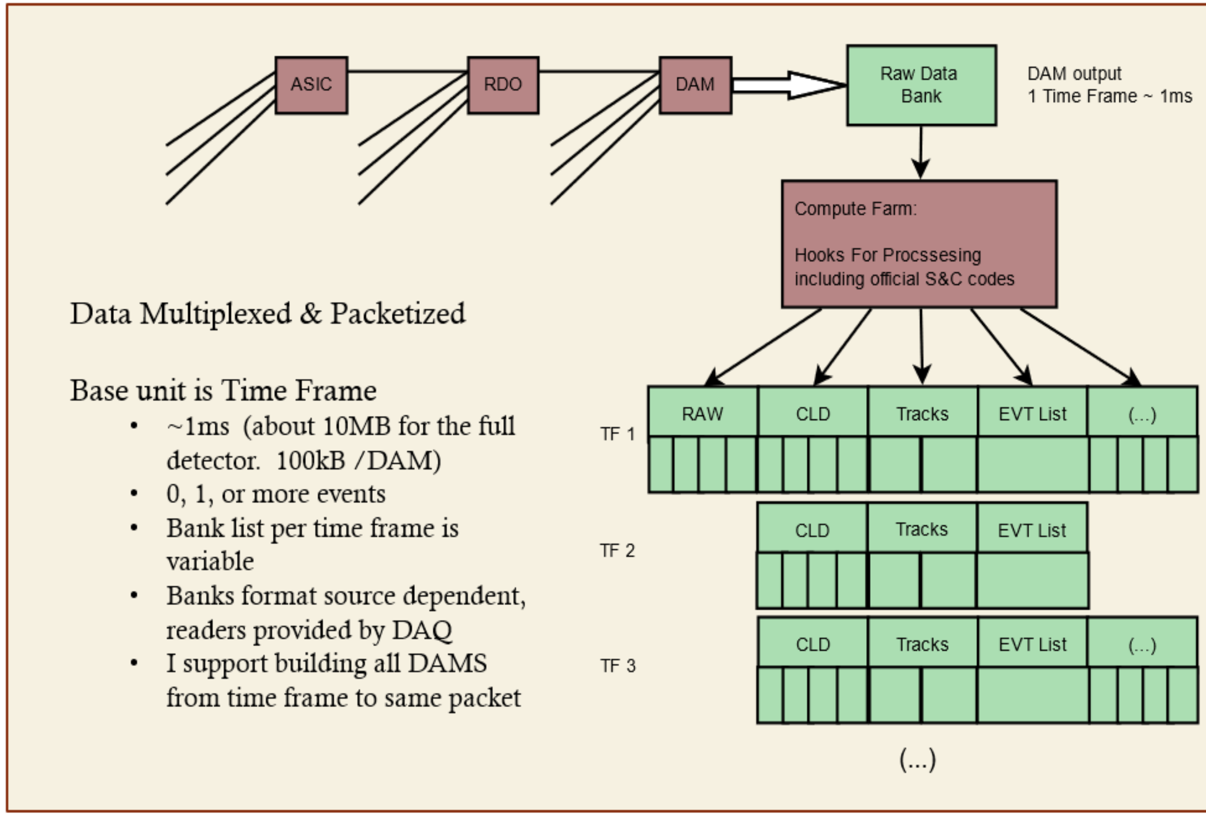

DAQ Processing: CLD data bank for Barrel TOF

Barrel: Data Compression & Processing 1

- each barrel stave is 128 strips in z and 64 strips in phi
  - the local stave coordinate system is thus a plane of 64 x 128 "pixels"
- per-channel processing
  - gain correction is applied to the ADC data
  - t0 correction is applied to the TDC data
  - slewing corrections is applied to TDC data
  - ⇒ obviously unphysical data is removed (cuts down the noise significantly)
- cluster finder runs on this (locally x-y) plane and looks for strip patterns
  - more than 1 adjacent strips with the same timing information form a valid particle (as opposed to random noise)
  - timing data should correspond to possible collisions, out-of-time hits are assumed to be noise
  - morphological cuts: e.g. middle pixels should be higher than neighbors, etc
  - we think this gives us at least x100 noise rejection
    - a better number needs a slow simulator & reconstruction
- hits are formed and saved with the following information
  - **coarse counter C** 17 bits (relative tick from the start of the timeframe ; 17 bits is up to 1.3 ms ⇒should be enough)
  - **local x-coordinate** as a fixed point number of 7.5 bits (relative coordinate system of the stave)
  - **local y-coordinate** as a fixed point number of 6.5 bits (relative coordinate system of the stave)
  - **fine hit time T** as a fixed point number in 10.5 bits (timing from TDC)
  - **summed up ADC** (charge) is 12 bits
  - **flags** 4 bits
  - total #bits per hit is 71 ⇒ but let's call it 10 bytes

Tonko Ljubicic 9 Feb 2023 (at the DAQ WG)

Kinds of data calibration data needed by DAQ

- Pedestals
- Gain corrections
- T0 Corrections
- Slewing Corrections

The details will be detector specific

They can be tied to actual voltage settings as well

These may need special runs and automated processing

DAQ Processing: Pedestal Bank

- We will not be able to readout black detectors, so pedestal runs in ePIC are time-interleaved with ZS data (or blocked data)
- Synchronous commands would be sent via the GTU -> RDO -> ASIC to send bursts of non-ZS data
- These would be split out from the ZS data to make NonZS data banks

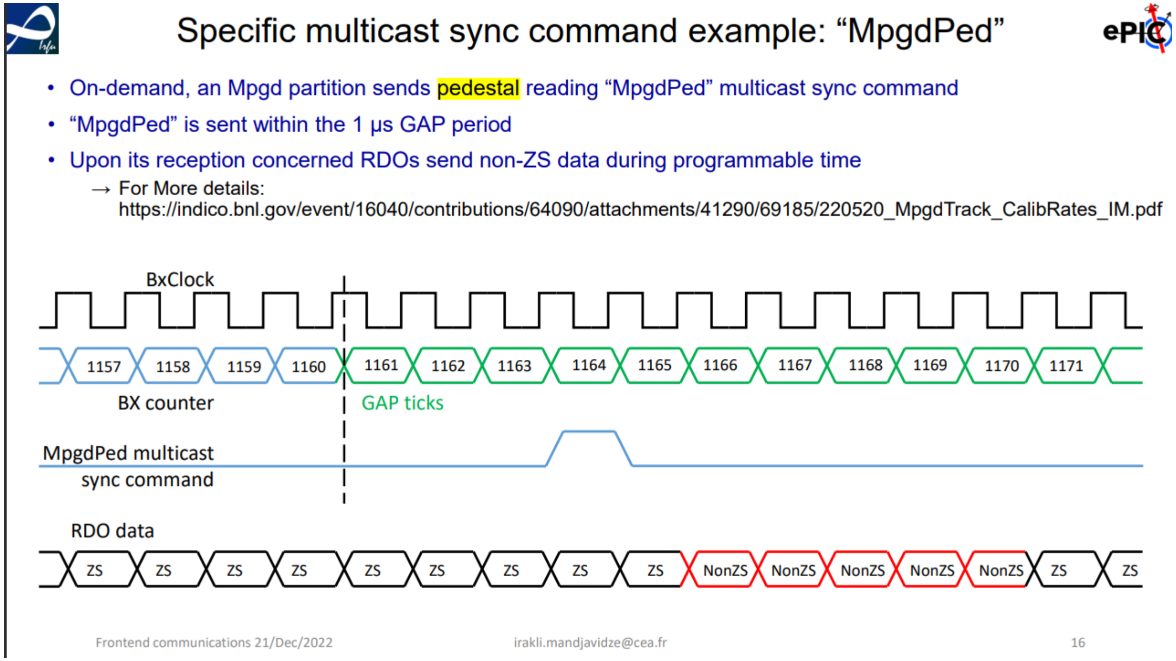

Other DAQ Processing

- Slow Controls Bank:
  - Each timeframe (or at some other periodicity) we might write data banks with important collider information (spin patterns), Magnet info, Basic monitoring info from detectors to ensure that this information is always available during analysis
- Lumi information Bank:
  - The Far Backward detectors will produce histograms of the full data in order to determine bunch by luminosities
- Event candidate lists
  - DAQ will need to evaluate potential events for the cross-detector filters for FB and dRICH. These can be stored for later evaluation (or use)
  - DAQ will also need monitoring to ensure, for example, that the synchronization of the data is robust. This may also require event candidate lists
- Reconstruction “seed tasks” to simplify offline processing
  - The algorithms for these (and possibly for some of the previously discussed tasks) should be provided, and monitored by the collaboration (or the reconstruction team) but run in the context of DAQ
  - We will need a way to define the offline resources available at various stages and allow for transparent use of analysis codes.
  - We also would wish to have a scheme to make the codes in place at a given time transparent to the collaboration
  - This is not easy, for example:
    - Release schedules have far different constraints
    - Have strict constraints in DAQ for avoiding crashes, and for limiting resource use
    - Have separate firewall protected enclaves from normal analysis processing

Summary

- Explained the expected data volume
- Explained the limitations / uncertainties / tasks needed to complete to improve the data volume calculations
- Discussed some of the expected processing that would be handled early in the DAQ / Software chain

- Questions?

Backup

EPIC Detector Scale and Technology Summary:

Detector System	Channels	RDO	Gb/s (RDO)	Gb/s (Tape)	DAM Boards	Readout Technology	Notes
Si Tracking: 3 vertex layers,	7 m^2	400	26	26	10	MAPS:	Fiber count limited by Artix Transceivers
2 sagitta layers,	36B pixels					Several flavors:	
5 backward disks,	5,200 MAPS sensors					curved its3 sensors for vertex	
5 forward disks						Its-2 staves / w improvements	
MPGD tracking: Electron Endcap	16k	32	1	.2	14	uRWELL/ SALSA	64 Channels/Salsa, up to 8 Salsa / FEB&RDO
Hadron Endcap	16k	32				uRWELL/ SALSA	
Inner Barrel	30k	120				MicroMegas/ SALSA	256 ch/FEB for MM
Outer Barrel	140k	288				uRWELL	512 ch/FEB for uRWELL
Forward Calorimeters: LFHCAL	64k	74	502	28	10	SiPM/ HG2CROC	Assume HGCROC 56ch * 16 ASIC/RDO = 896ch/RDO
ECAL W/SciFi	16k	64				SiPM/ Discrete	Assume FLASH FEB 16ch * 16 FEB/RDO = 256ch/RDO
Barrel Calorimeters: HCAL	8k	9				SiPM/ HG2CROC	
HCAL insert	8k	9				SiPM/ HG2CROC	
ECAL SciFi/PB	8k	32				SiPM/ HG2CROC	
ECAL ASTROPIX	88M pixels	40				Astropix	Assume similar structure to its-2 but with sensors with
Backward Calorimeters: NHCAL	16.2k	18				SiPM/ HG2CROC	250k pixels for RDO calculation.
ECAL (PWO)	3k	12				SiPM/ Discrete	
Far Forward: B0: 3 MAPS layers	300M pixel	10	15	8	5	MAPS	3x20cmx20cm
1 or 2 AC-LGAD layer	300k or 600k	30				AC-LGAG / EICROC	600^cm layers (1 or 2 layers)
2 Roman Pots	1M (4 x 135k layers x 2 dets)	64				AC-LGAD / EICROC	13 x 26cm layers
2 Off Momentum	650k (4 x 80k layers x 2 dets)	42				AC-LGAD / EICROC	9.6 x 22.4cm layers
ZDC: Crystal Calorimeter	400	10				APD	There are alternatives for AC-LGAD using MAPS and low
32 Silicon pad layer	11.52k	10				HGCROC as per ALICE FoCal-E	channel count DC-LGAD timing layers
4 silicon pixel layers	160k	10					
2 boxes scintillator	72	2					
Far Backward: Low Q Tagger 1	33M pixels	12	150	1	4	Timepix4	
Low Q Tagger 2	33M pixels	12				Timepix4	
Low Q Tagger 1+2 Cal	700	1			1		
2 x Lumi PS Calorimeter	1425/75	1			1	(SiPM/HG2CROC) / (PMT/FLASH)	
Lumi PS tracker	80M pixels	24				Timepix4	
Photon Detector							
PID-TOF: Barrel	2.2M	268	728	1	12	AC-LGAD / EICROC (strip)	bTOF 128 ch/ASIC, 64 ASIC/RDO
Endcap	5.6 M	134				AC-LGAD / EICROC (pixel)	eTOF 1024 pixel/ASIC, 24-48 ASIC/RDO (41 ave)
PID-Cherenkov: dRICH	320k	200	1865	17	20	SiPM/ ALCOR	Worse case after radiation. Includes 30% timing window. Requires further data volume reduction
pfRICH	70k	17	24	12	1	HRPPD / EICROC (strip or pixel)	software trigger
DIRC	70k	24	11	6	6	HRPPD / EICROC (strip or pixel)	

Electron-Ion Collider