



WBS 6.10.09 DAQ and Scientific Computing


Jeff Landgraf and David Abbott

L3 CAMs

EIC Detector Comprehensive
Design Review

August 29-30, 2023

Electron-Ion Collider



Overview

- The technology and scale of the ePIC DAQ (Charge 1)
 - Streaming readout community
 - Streaming architecture
 - System components and operational concept
 - Component counts & expected data
- Requirements Flow Down (Charge 2)
- Status of Components (Charge 5)
- Schedule (Charge 6)

Charges

1. Given the detector progress over the last two years and the status of the ePIC detector, are the projected timelines of the Electron-Ion Collider detector feasible? Do there remain significant open detector technology questions?
2. Are the requirements for the detector and their flow down sufficiently comprehensive for this stage of the project to complete the design of the various detector technologies?
3. Are the interfaces between the elements of the design adequately defined for this stage of the project and to proceed with the detector long-lead procurement items?
4. Is the design of these long-lead procurement items sufficiently advanced and mature to start procurement in 2024? Are the technical specifications complete?
5. Is the projected design maturity of the further detector components likely to be accomplished by the end of 2024 for CD-2 and CD-3?
6. Is the overall schedule for completion of the design, production, and installation of detector components realistic?

EIC Streaming Readout Community

- **The EIC Streaming Readout (SRO) Consortium**

- Formed in 2018 with the goals of engaging the nuclear physics community to explore a streaming readout architecture for the EIC detectors.
- Series of ten workshops so far since 2018 (next in Oct 2023)
- Led to adoption of SRO for Yellow Report the physics program for EIC.

- **ePIC Electronics/DAQ Working Group**

- Conveners: F. Barbosa (JLab), J. Huang (BNL), J. Landgraf (BNL)
- Regular meetings to focus on design and requirements specifically for the new ePIC detector
- Interface to detectors, readout, DAQ, online processing, monitoring and control.

- **ePIC Software/Computing Working Group**

- Conveners: M. Diefenthaler (JLab), S. Joosten (ANL), T. Wenaus (BNL), W. Deconinck (UM)
- Coordinate software development for EIC physics

- **ePIC Software/Computing Streaming DAQ Working Group**


- Conveners: M. Battaglieri (JLab), J. Huang (BNL)
- Streaming Reconstruction

- **EIC Project L3 CAMs (DAQ/Online, Electronics)**

- D. Abbott (JLab), J. Landgraf (BNL), F. Barbosa (JLab)
- Project management (cost/schedule)
- Coordination between relevant EIC working groups (including Detectors)


ePIC Streaming Readout

Triggered DAQ



(Marco Battaglieri INFN)

Streaming Readout (SRO)



Features of ePIC SRO

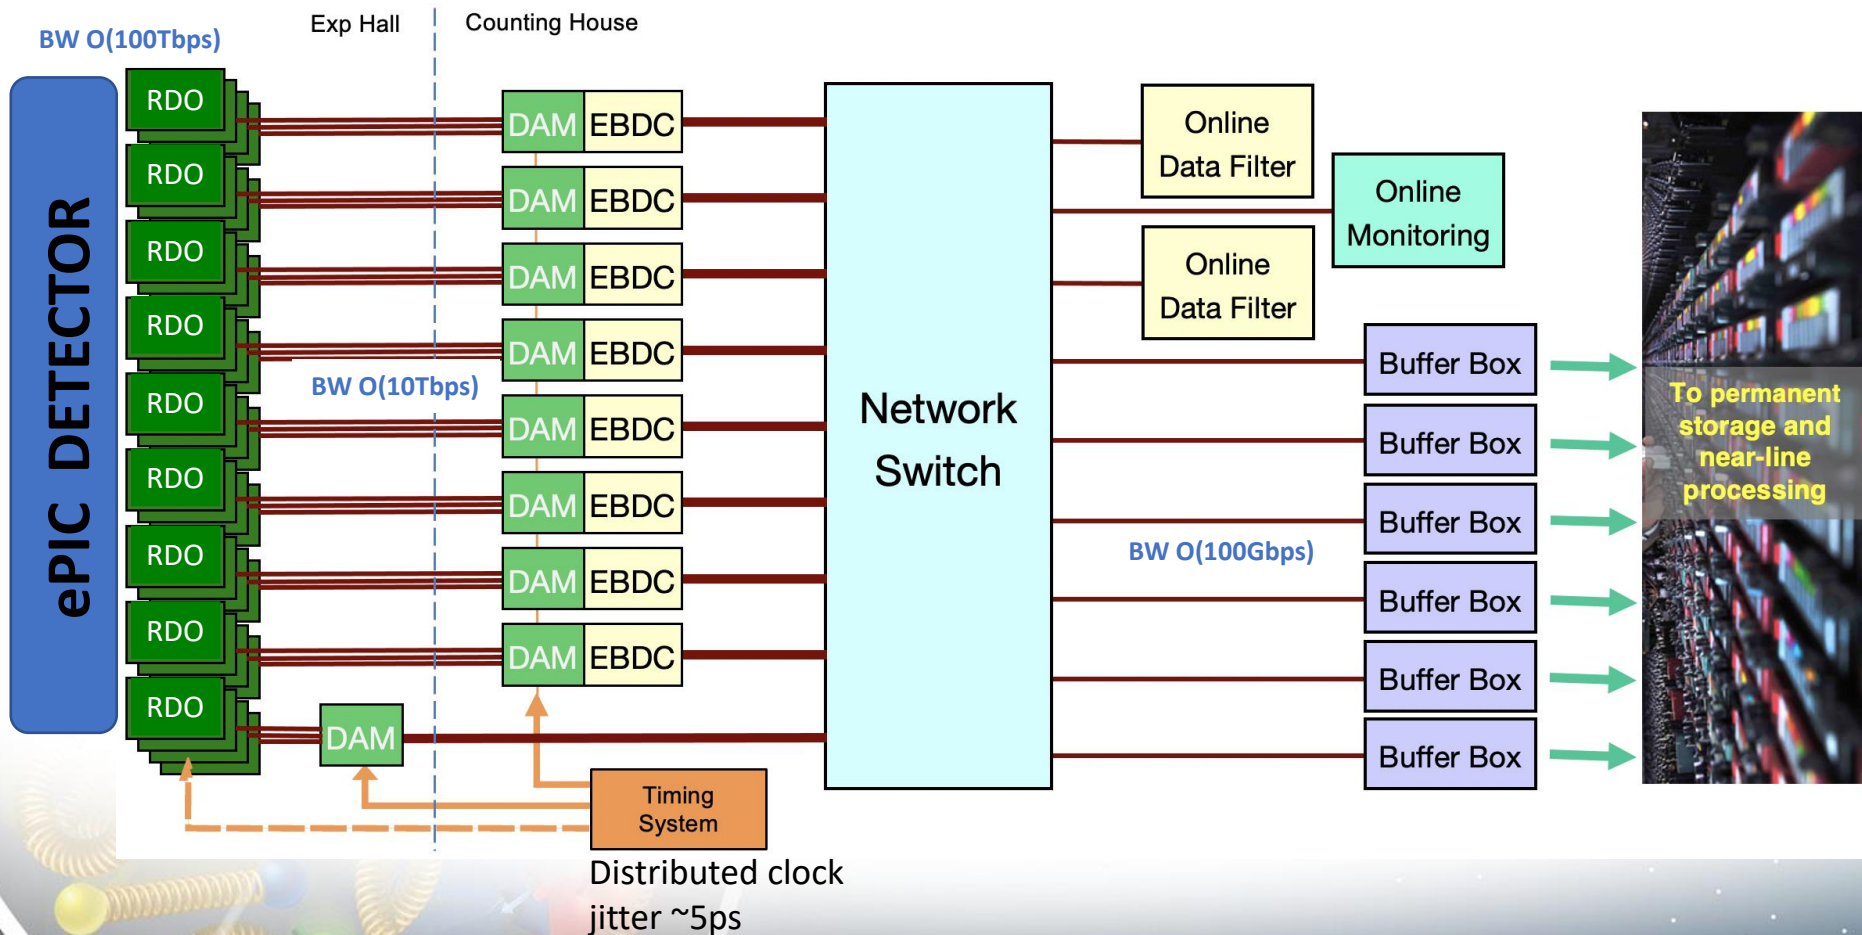
- No global trigger
 - No custom trigger electronics
 - Zero-suppress early (ASICs)
- Hits identified by time stamp rather than by event
- Flexibility in event selection
 - Can be performed with all channels available
 - Can be performed in CPU, FPGA, or GPU
- Cons
 - SRO has greater sensitivity to noise and background

ePIC Streaming DAQ/Computing Architecture

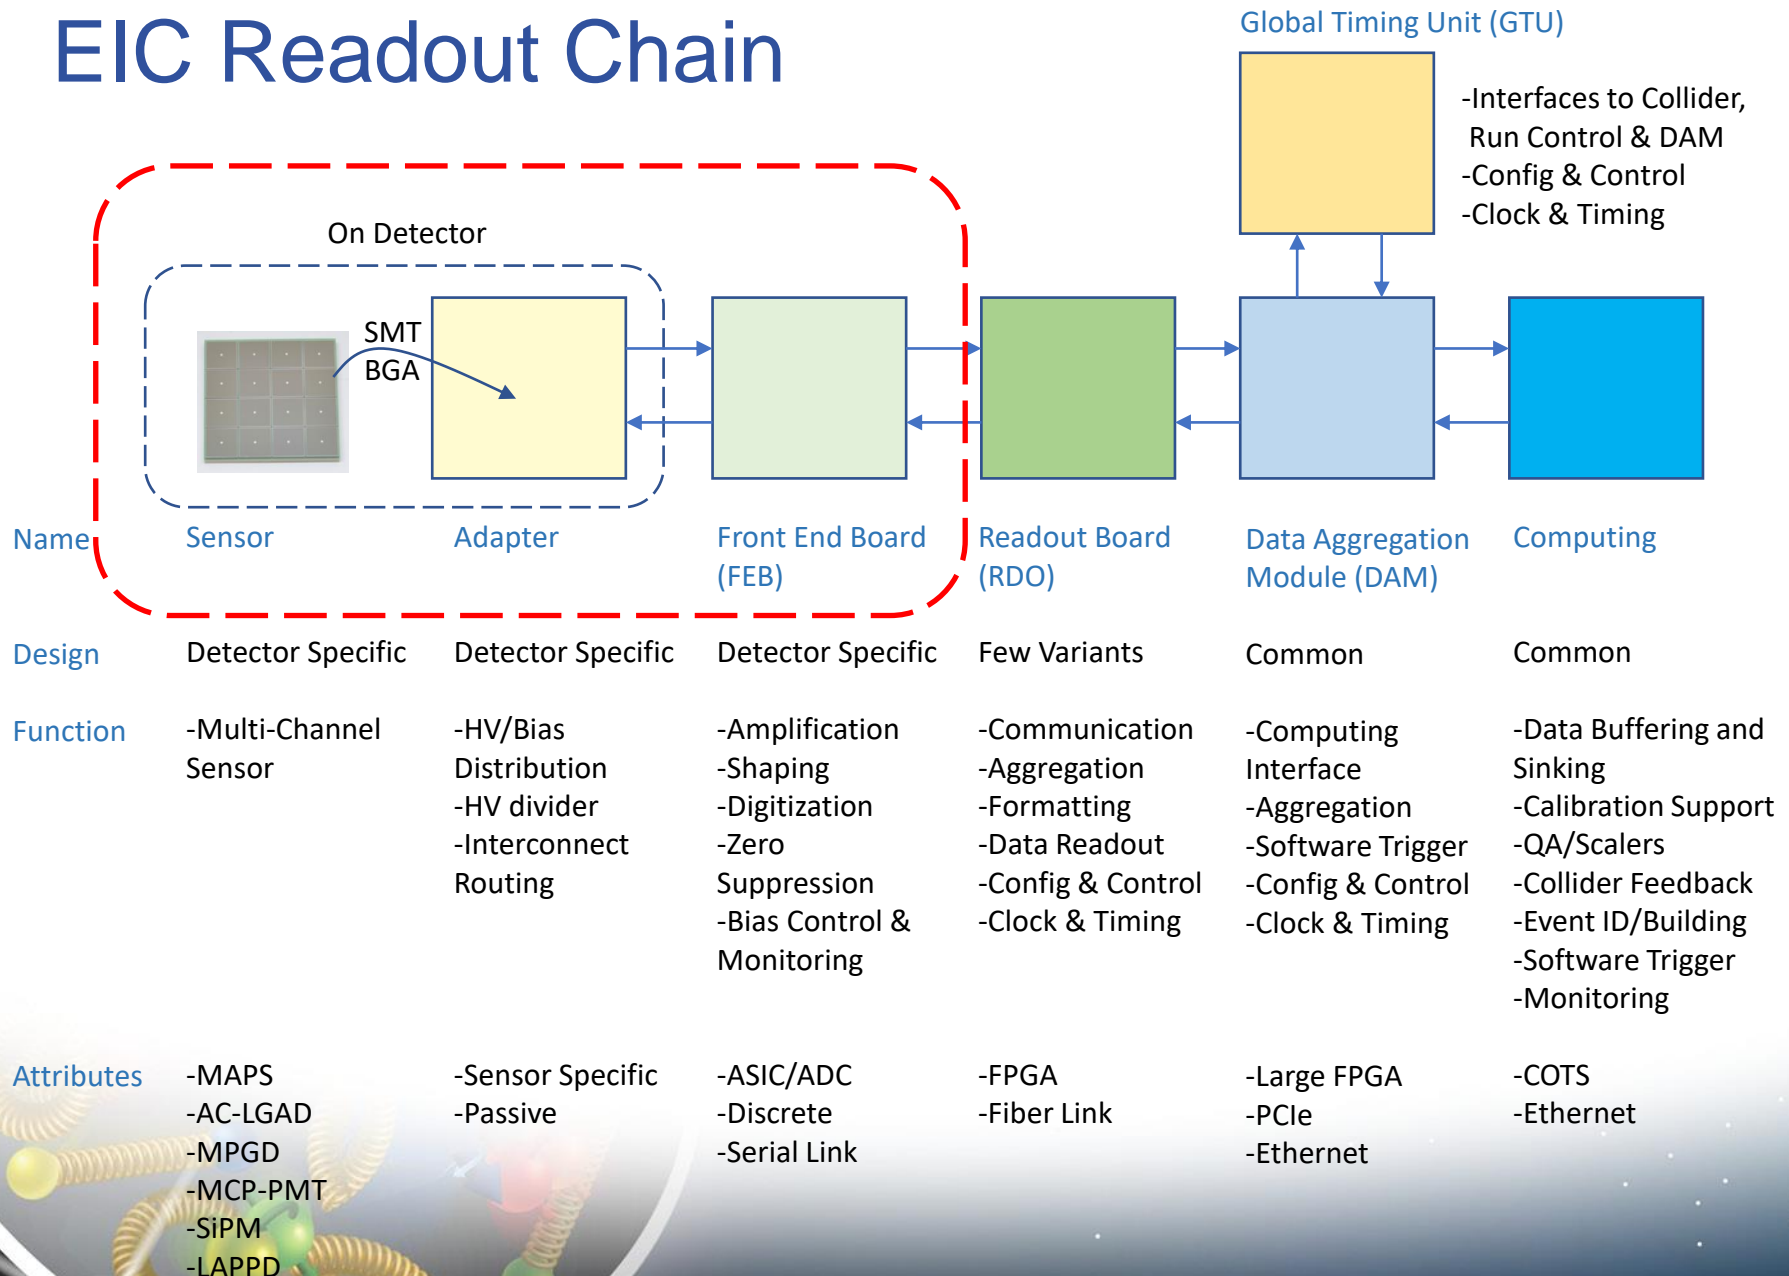
Bunch Crossing ~ 10.2 ns/98.5 MHz

Interaction Rate ~ 2 us/500 kHz

Low occupancy.



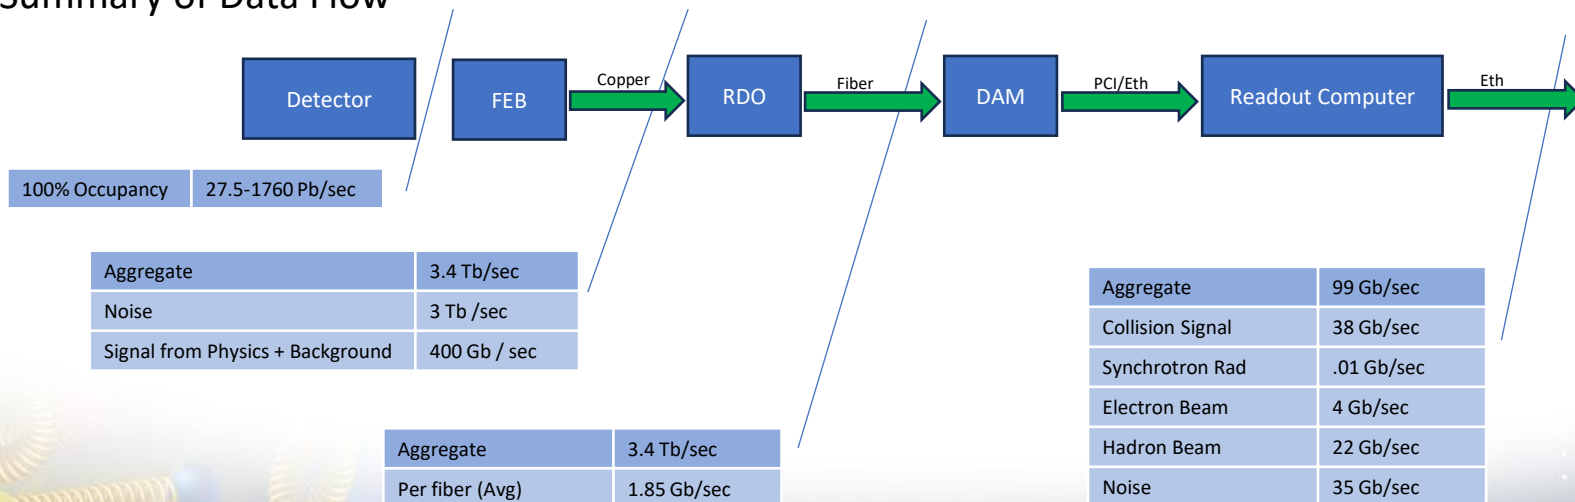
EIC Readout Chain



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	37	27	26
Calorimeters	500M		104k			448	1126	19	502	27
Far Forward	300M	2.6M	170k			178	492	8	15	8
Far Backward	82M		2k			49	98	4	150	1
PID		7.8M	344k		140k	1281	2603	47	2628	36
TOTAL	36.9B	10.4M	620k	202k	140k	2828	6063	115	3,322	98


Summary of Data Flow



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	37	27	26
Calorimeters	500M		104k			448	1126	19	502	27
Far Forward	300M	2.6M	170k			178	492	8	15	8
Far Backward	82M		2k			49	98	4	150	1
PID		7.8M	344k		140k	1281	2603	47	2628	36
TOTAL	36.9B	10.4M	620k	202k	140k	2828	6062	115	2,222	98

Summary of Data Flow



Significant Filtering Challenges:

- After several years in a radiation environment, “noise” dominates signal at the RDO 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 / Bunch Crossing.
 - 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

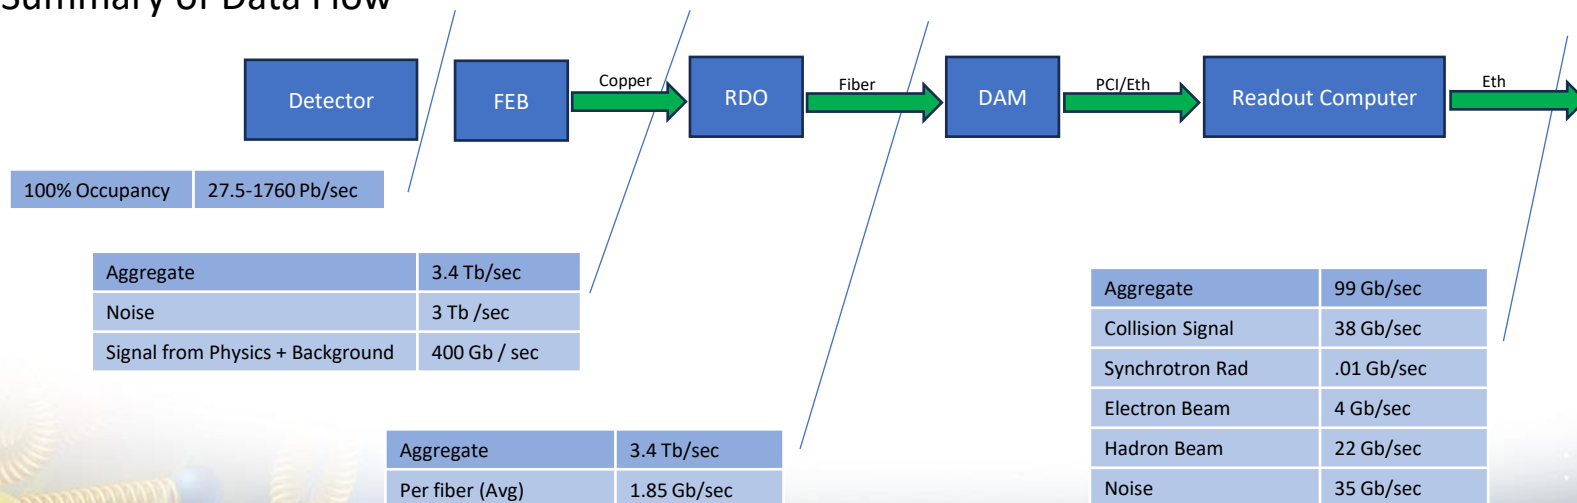
Summary of Channels

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

Detector Group									Data Volume (RDO) (Gb/s)	Data Volume (To Tape) (Gb/s)
	MAPS									
Tracking	36B									27
Calorimeters	500M									26
Far Forward	300M									502
Far Backward	82M		2k			49	98	4	15	27
PID		7.8M	344k		140k	1281	2603	47	150	1
TOTAL	36.9B	10.4M	620k	202k	140k	2828	6063	115	2628	36
									3,322	98

Summary of Data Flow

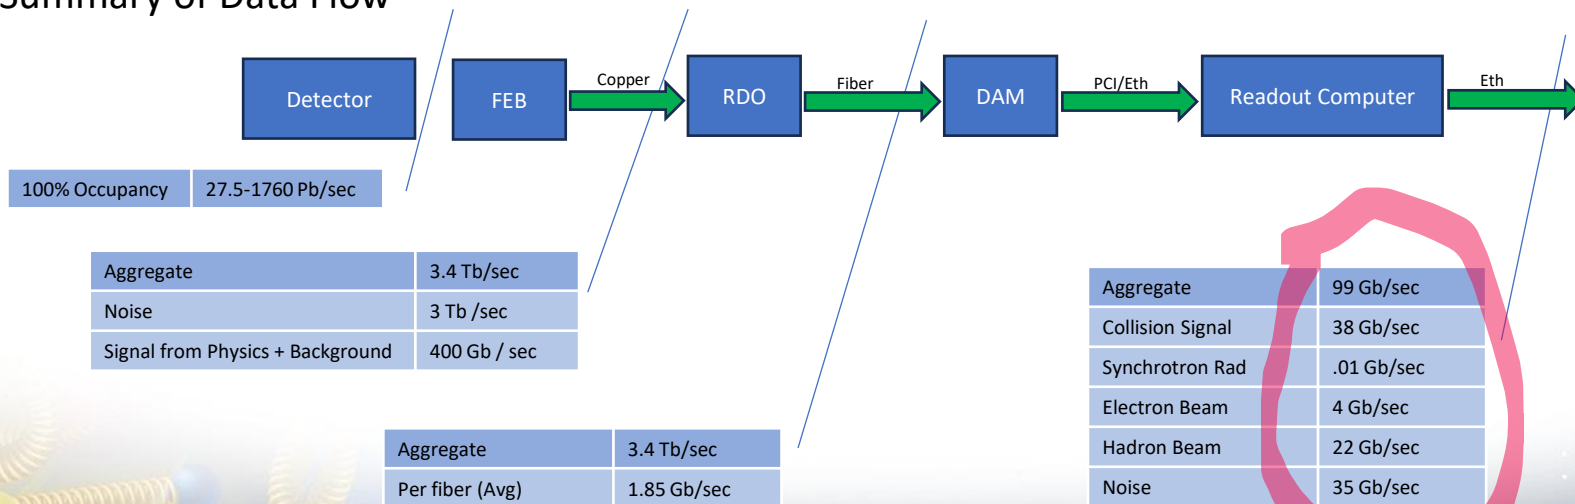


Summary of Channel Counts

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	37	27	26
Calorimeters	500M		104k			448	1126	19	502	27
Far Forward	300M	3.6M	470k			178	492	8	15	8
Far Back							98	4	150	1
PID							2603	47	2628	36
TOTAL							6063	115	3,322	98

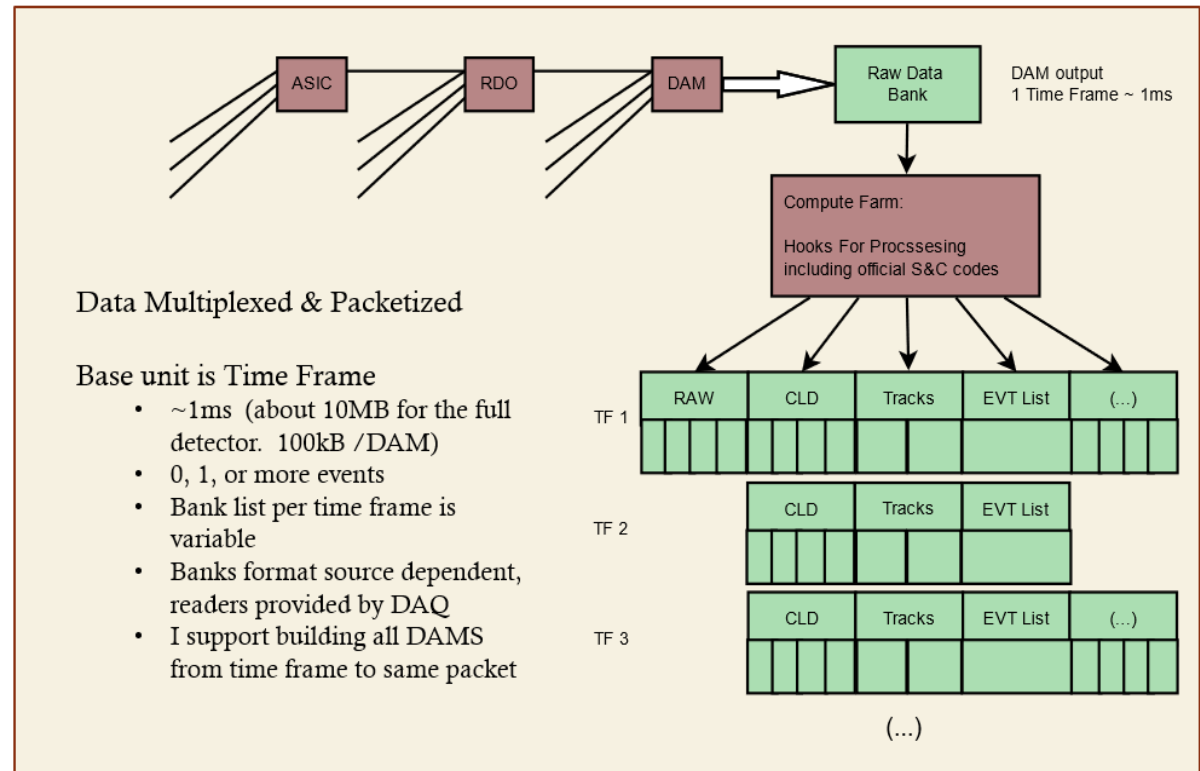
After online processing we expect roughly equal contributions from collision hits, beam background hits, and noise.

Summary of Data Flow



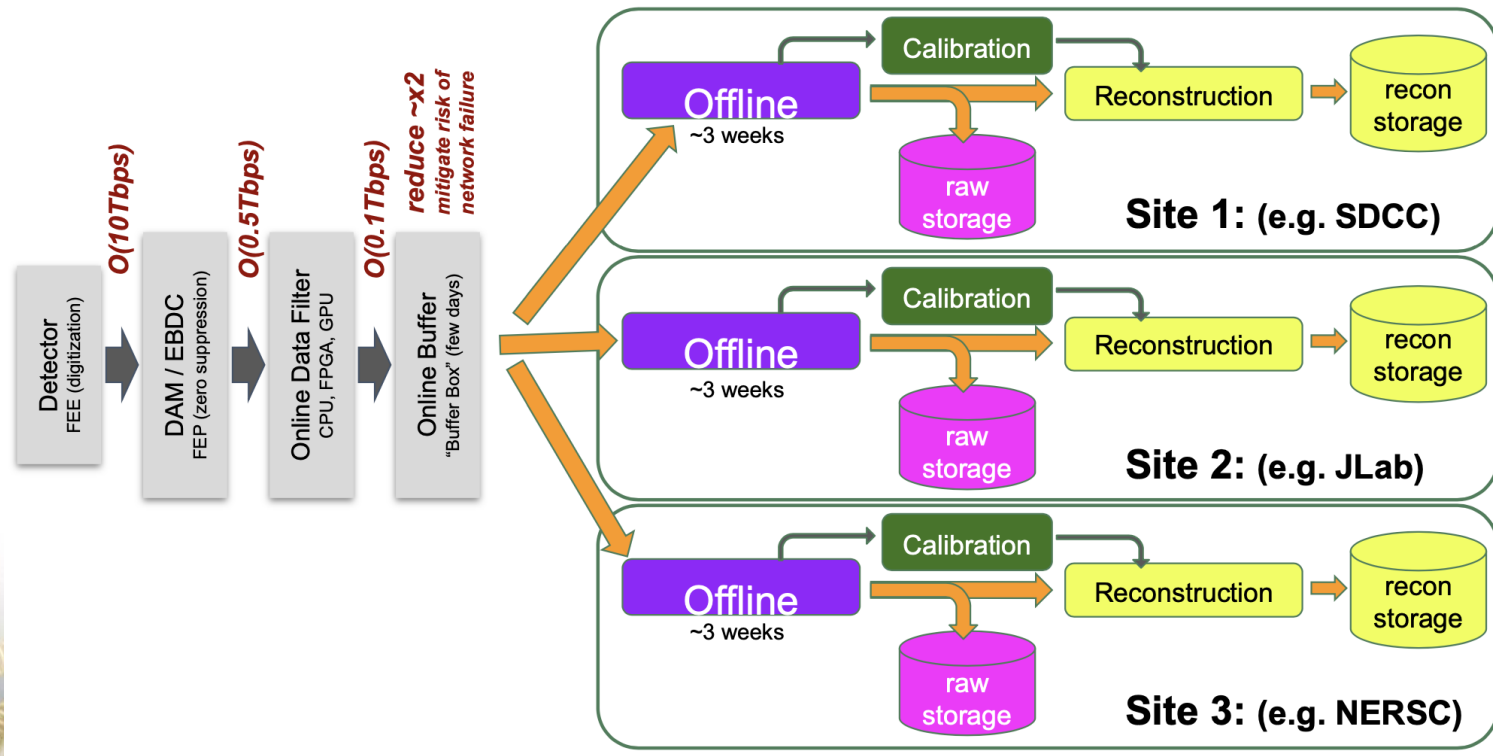
DAQ Computing

- 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
- Scalers / continuously running DAQ components



Boundary between Online and Offline

- The streaming architecture allows for some blurring of the offline and online processing with respect to calibrations, QA monitoring, and potential seed analysis to improve reconstruction turn around. These aspects are not yet fully specified
- Require an interface to allow sharing of code between offline / online processing
- Local buffering provides elasticity in the transfer of data to computing facilities
- Offline buffering will allow several weeks for calibration and reconstruction.



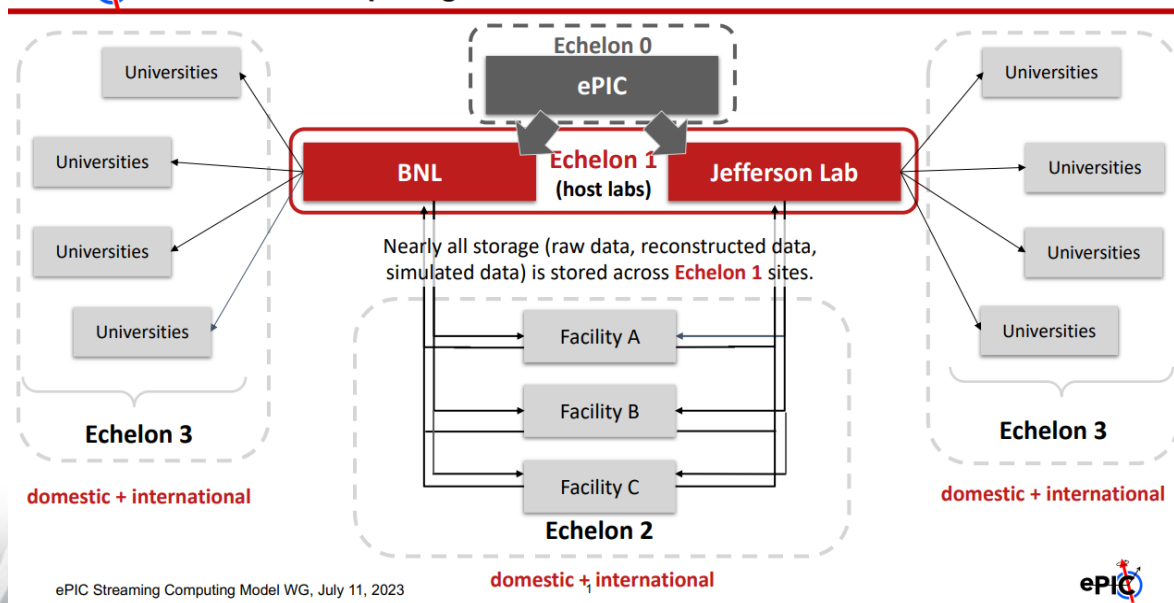
ePIC Scientific Computing

- The organization of the scientific computing is handled by the ePIC collaboration
- The computing group is actively implementing the scientific software as well as defining the data model to best interact with the streaming DAQ system.
- A primary goal for this collaboration is to provide an environment that can provide a real-time interplay with the running experiments.

For more information reference this review or planned reviews of the ePIC software

- August 22, 2022 – EIC Software Infrastructure Review
<https://indico.bnl.gov/event/16676/>
- October 19-20, 2023 - ePIC software and Computing model review
- December, 2023 - ePIC presents the updated computing model to the EIC Resource and Review Board

ePIC Distributed Computing Model



ePIC DAQ Requirements Summary

General Requirements

Read out all detector data and transfer to external facilities

Prefer COTS solutions to avoid cost, risk, obsolescence

Functional Requirements

Sufficient space, power, environment, and proximity to sensors

Streaming, continuous operation, while supporting independent detector operation

Define components and specify the DAQ must provide sufficient links and processing power

Provide common hardware, software, and network for detector slow controls efforts

Performance Requirements

Specific Environmental/Power requirements


Online Computing Network must support $\geq 100\text{Gb/s}$ per node and $\geq 400\text{Gb/s}$ external

Readout fibers from Front ends must support $\geq 10\text{Gb/s}$ per fiber and 10Tb/s aggregate

Timing system must deliver clock with $<100\text{ps}$ jitter, and support $<5\text{ps}$ jitter for selected detectors

72 hours of elasticity buffer in online computer nodes

Sufficient computing power to perform needed online processing



System Component Status

Charge #5 – Design Maturity


Detailed Electronics Design starts with CD3. The current status is that we are evaluating system design and interactions between components

Timing Protocol

- Defines the data flow between the GTU, RDO, and DAM boards
- Optical Protocol Requirements
 - Timing (≤ 5 ps Hi-Res Detectors)
 - Timing (≤ 100 ps Lo-Res Detectors)
 - Timing phase stability on power cycle
 - Configuration and slow control monitoring passthrough for ASICs/FEBs
 - Real-Time Command / Control
 - Data Transfer (≥ 10 or 25 Gb/s per RDO)
- Optical protocol choices under evaluation
 - ePIC – simple custom protocol
 - GPT in FPGA
 - Dedicated Clock/Reconstructed Clock (2 or 3 fibers per RDO)

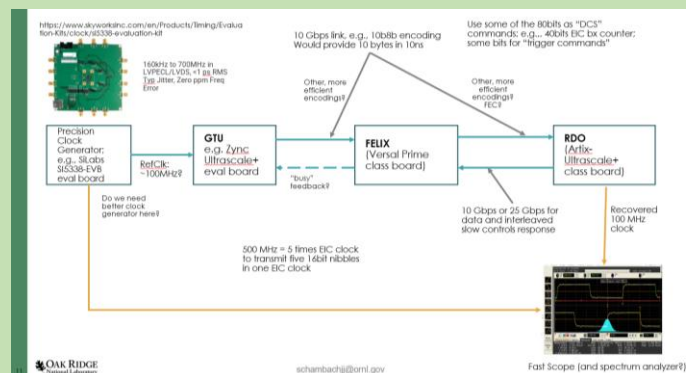
RDO Status:

- Mocking up with devkits for initial studies
- Nominal 2.5" x 2.5" for common elements + connector space
- FPGA defined by price (Artix Ultrascale+ Class)
- Power 3-5 Watts. 2 LV +5V for optical interfaces, lower for FPGA/digital components
- Accept 2 or 3 fibers for potential Hi-Res Timing




GTU status:

- Mocking up with devkits for initial studies
- Interface to collider (Bunch crossing clock, Spin Patterns per bunch crossing)
- Distribution of timing to ~100 DAMs + ~220 RDOs of Hi-Res Timing Detectors
- Interface to Run Control



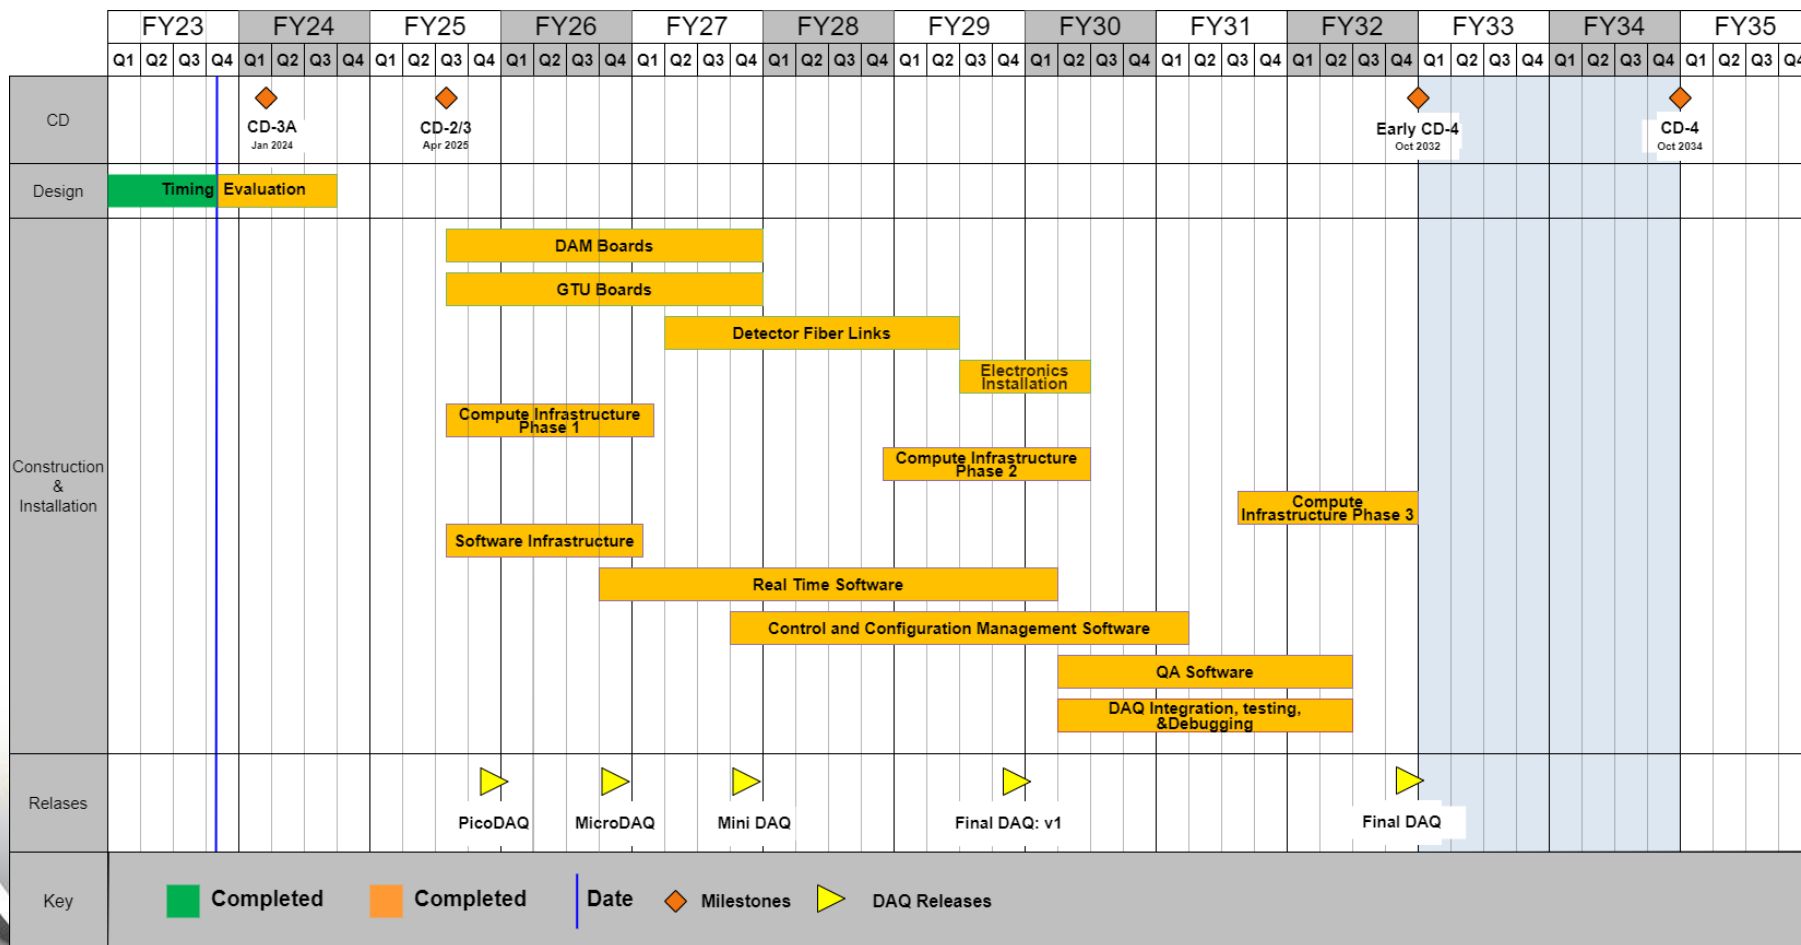
DAM Status:

- Obtained Board From Omega
- Versa FPGA
- 24 x 10 PCIe Gen4 x16 or 100Gb Ethernet
- 24 Firefly links 10Gb/s or 25Gb/s options
- 48 link upgrade planned
- Additional front panel I/O



EIC DAQ/Computing schedule


- The DAQ and computing schedule is planned to be somewhat accelerated in relation to the full detector construction.
- A functional DAQ will be necessary for small scale detector testing as well as commissioning and pre-ops.
- Hardware and software development early in the construction phase will produce a series of releases with increasing functionality
- Computing resources are staged in order to make purchases when the hardware is needed, taking advantage of price savings and performance improvements in time



Conclusion

1. Given the detector progress over the last two years and the status of the ePIC detector, are the projected timelines of the Electron-Ion Collider detector feasible? Do there remain significant open detector technology questions?
 - We have presented the technologies and status of their development for the ePIC DAQ
2. Are the requirements for the detector and their flow down sufficiently comprehensive for this stage of the project to complete the design of the various detector technologies
 - We have summarized the DAQ requirements
3. Are the interfaces between the elements of the design adequately defined for this stage of the project and to proceed with the detector long-lead procurement items?
 - The interfaces are defined, but there are no LLP directly defined by 6.10.09
4. Is the design of these long-lead procurement items sufficiently advanced and mature to start procurement in 2024? Are the technical specifications complete?
 - N/A, no LLP for 6.10.09
5. Is the projected design maturity of the further detector components likely to be accomplished by the end of 2024 for CD-2 and CD-3?
 - The system should be fully defined by the end of 2024, although electronics board design starts with CD3
6. Is the overall schedule for completion of the design, production, and installation of detector components realistic?
 - We have presented the scope and schedule of the DAQ system

Backup



EPIC Detector Scale and Technology Summary:

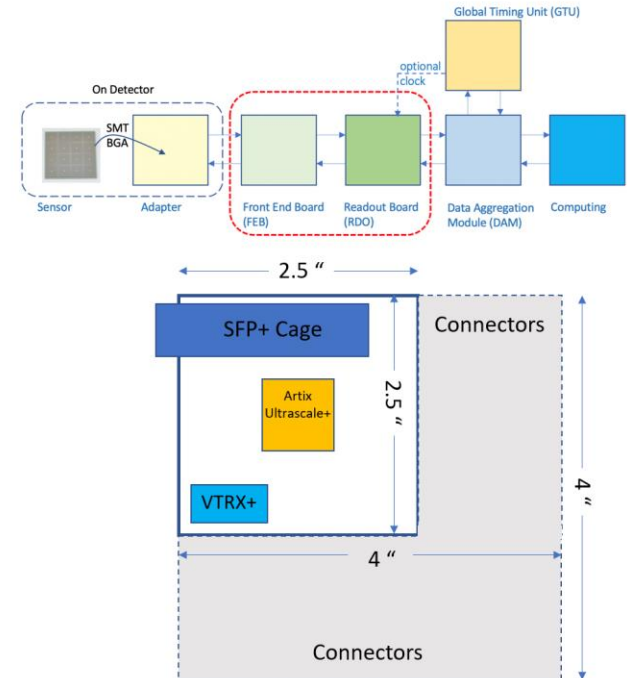
Detector System	Channels	RDO	Gb/s (RDO)	Gb/s (Tape)	DAM Boards	Readout Technology	Notes
Si Tracking: 3 vertex layers, 2 sagitta layers, 5 backward disks, 5 forward disks	7 m ² 36B pixels 5,200 MAPS sensors	400	26	26	17	MAPS: Several flavors: curved its-3 sensors for vertex Its-2 staves / w improvements	Fiber count limited by Artix Transceivers
MPGD tracking: Electron Endcap Hadron Endcap Inner Barrel Outer Barrel	16k 16k 30k 140k	32 32 120 288	1	.2	20	uRWELL / SALSA uRWELL / SALSA MicroMegas / SALSA uRWELL / SALSA	64 Channels/Salsa, up to 8 Salsa / FEB&RDO 256 ch/FEB for MM 512 ch/FEB for uRWELL
Forward Calorimeters: LFHCAL HCAL insert* ECAL W/SciFi Barrel Calorimeters: HCAL ECAL SciFi/PB ECAL ASTROPIX Backward Calorimeters: NHCAL ECAL (PWO)	63,280 8k 19,000 1,536 5,760 500M pixels 3,256 3,240	74 9 64 9 32 230 18 12	502	28	19	SiPM / HG2CROC SiPM / HG2CROC SiPM / Discrete SiPM / HG2CROC SiPM / HG2CROC Astropix SiPM / HG2CROC SiPM / Discrete	Assume HGCROC 56 ch * 16 ASIC/RDO = 896 ch/RDO Assume FLASH FEB 16 ch * 16 FEB/RDO = 256 ch/RDO *HCAL insert not in baseline Assume similar structure to its-2 but with sensors with 250k pixels for RDO calculation.
Far Forward: B0: 3 MAPS layers 1 or 2 AC-LGAD layer 2 Roman Pots 2 Off Momentum ZDC: Crystal Calorimeter 32 Silicon pad layer 4 silicon pixel layers 2 boxes scintillator	300M pixel 1M 1M (4 x 135k layers x 2 dets) 640k (4 x 80k layers x 2 dets) 400 11,520 160k 72	10 30 64 42 10 10 10 2	15	8	8	MAPS AC-LGAG / EICROC AC-LGAD / EICROC AC-LGAD / EICROC APD HGCROC as per ALICE FoCal-E	3x20cmx20cm 600^cm layers (1 or 2 layers) 13 x 26cm layers 9.6 x 22.4cm layers There are alternatives for AC-LGAD using MAPS and low channel count DC-LGAD timing layers
Far Backward: Low Q Tagger 1 Low Q Tagger 2 Low Q Tagger 1+2 Cal 2 x Lumi PS Calorimeter Lumi PS tracker	1.3M pixels 480k pixels 700 1425/75 80M pixels	12 12 1 1 24	150	1	4	Timepix4 Timepix4 (SiPM/HG2CROC) / (PMT/FLASH) Timepix4	
PID-TOF: Barrel Endcap	2.2M 5.6 M	268 134	728	1	17	AC-LGAD / EICROC (strip) AC-LGAD / EICROC (pixel)	bTOF 128 ch/ASIC, 64 ASIC/RDO eTOF 1024 pixel/ASIC, 24-48 ASIC/RDO (41 ave)
PID-Cherenkov: dRICH pRICH DIRC	344,064 69,632 69,632	1240 17 24	1865 24 11	17 12 6	30	SiPM / ALCOR HRPPD / EICROC (strip or pixel) HRPPD / EICROC (strip or pixel)	Worse case after radiation. Includes 30% timing window. Requires further data volume reduction software trigger

RDO Discussions

- Physical Characteristics
 - Locations
 - Size
 - Power / Cooling
 - Radiation Requirements
 - Link distances/rates
- Optical Protocol Requirements
 - Timing (5ps / 50-100ps)
 - I2C for ASICs/FEBs
 - Real-Time Command / Control
 - Data Transfer (10Gb/s / 25 Gb/s)
- Optical Protocol Choices
 - ePIC – simple custom protocol
 - GPT in FPGA
 - Dedicated Clock/Reconstructed Clock (2 or 3 fibers/RDO?)

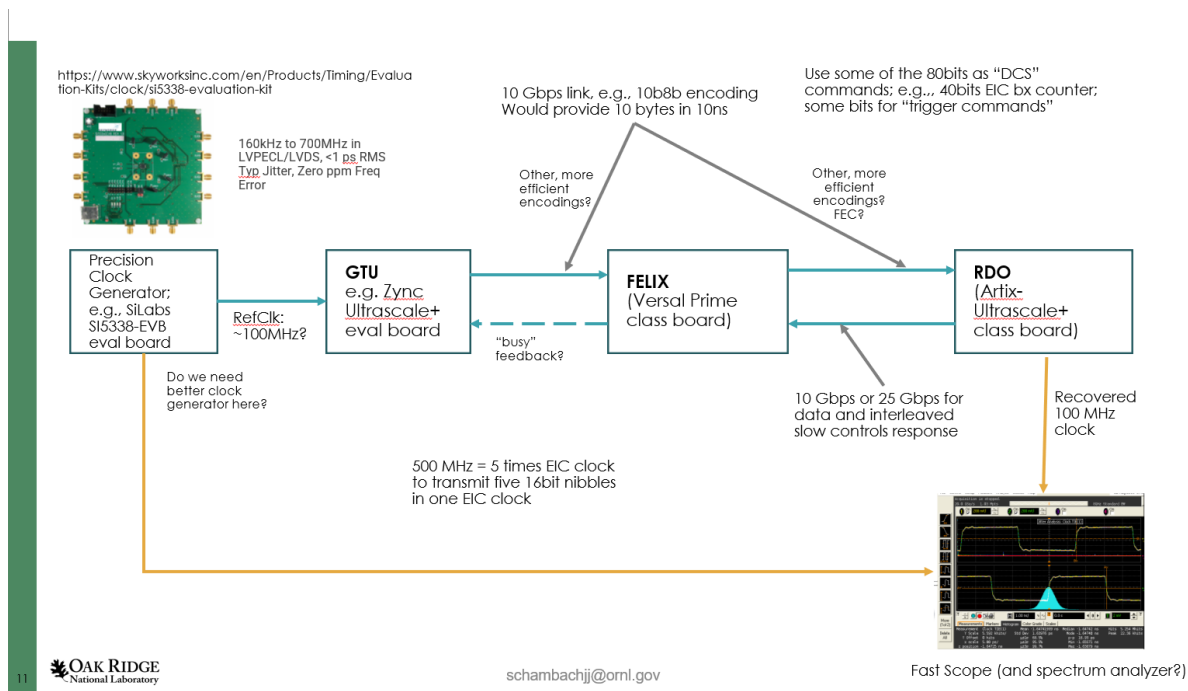
RDO Specifications/Guidance

- Nominally 2.5 in² for common RDO components including FPGA and optical link options (for example if it is integrated on the same PCB as the FEB). For standalone RDO, allow for up to 4 in² to provide space for copper-based connectors to FEBs.
- Power requirements: 3-5 Watts. Allow for at least two LV levels nominally 5V (for optics) and a lower voltage for FPGA power and ASIC signal management.
 - Consider using radiation tolerant switching voltage regulators (e.g. from CERN).
- Multiple optical link interfaces allow for flexible implementation of the RDO as either a standalone readout solution or use with the DAM boards. They can also be used for accepting an alternative low-jitter clock input.
 - Samtec Firefly connectors are also a potential option as they have a footprint similar to VTRX and also provide MTP options.



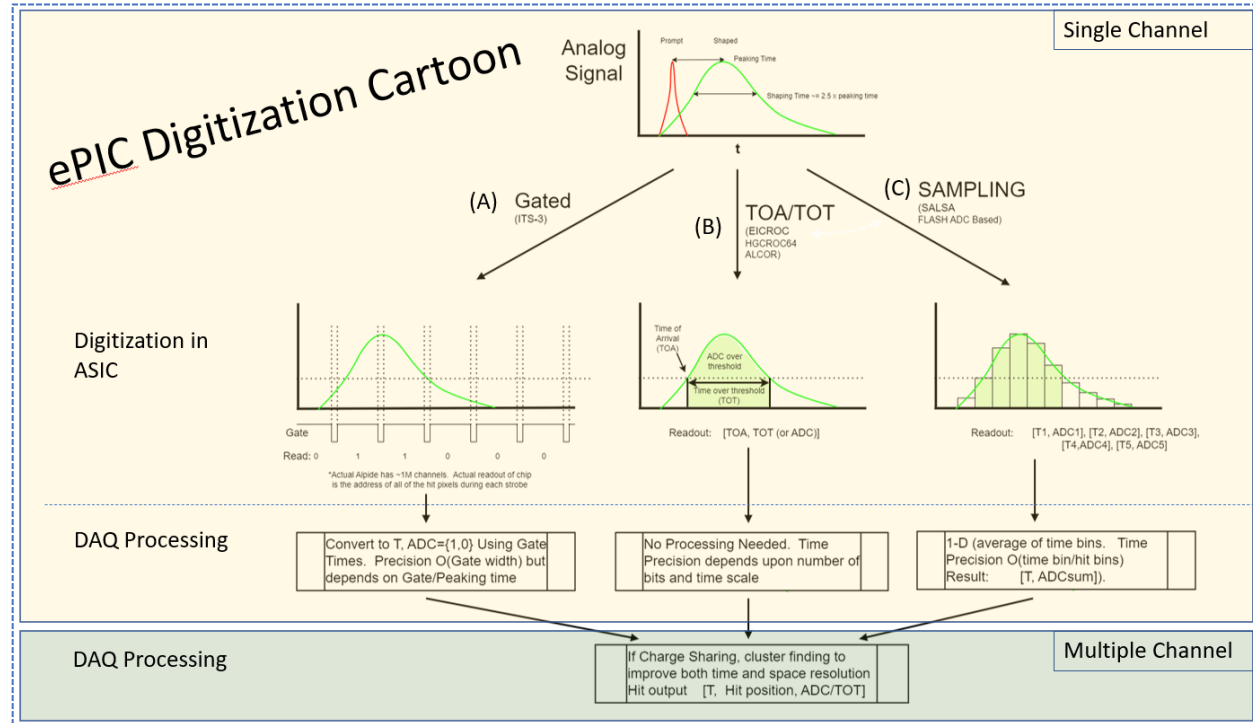
RDO Discussions

- RDO / Timing subgroup to mock up GTU/DAM/RDO to answer these questions
 - Slow startup due to time constraints of group members
 - Devkits being distributed
 - Arrangements for FELIX 182 slow but progressing
- Hope to start getting answers this fall
 - Protocol document
 - Timing measurements



Data Volumes

- Distilled volumes from
 - Background Group Simulations provide hits/sec above detector thresholds
 - Synchrotron Radiation
 - Hadron Beam
 - Electron Beam
 - DIS (18x275 scaled to 500khz)
 - Noise estimates from
 - “digitization spreadsheet”
 - E&DAQ WG presentations
 - Bits / hit = Charge Sharing * Hit
- Duration * Bits Per Hit
 - Cluster finding (where possible) reduces hit size to BitsPerHit
 - Cluster finding (when possible) reduces noise $\times 10^{-5}$
 - Software “triggering” for Far Backward/dRICH
- Charge sharing / Time Sharing / noise estimates should be considered VERY preliminary

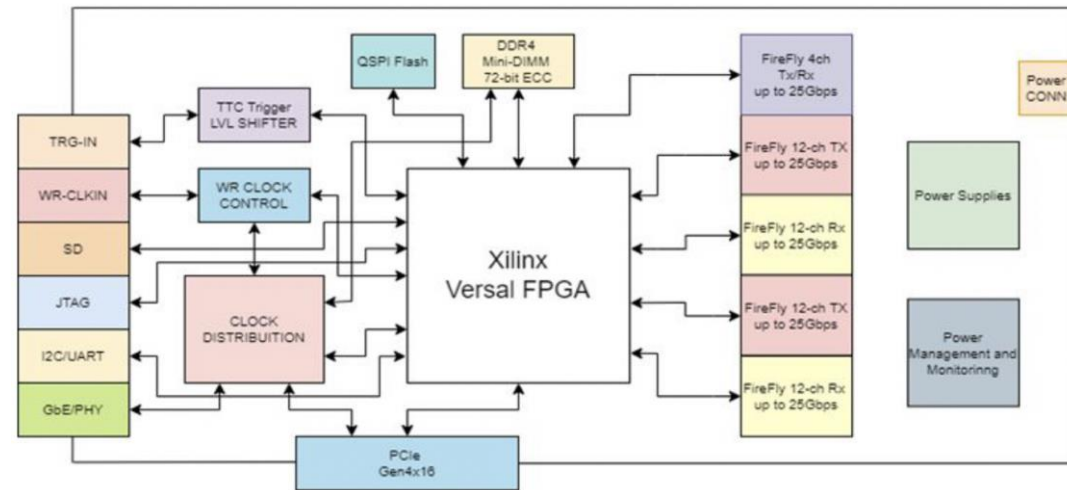


DAM Option - Front-End Link Exchange (FELIX)

- Originally developed for ATLAS at CERN
- Design team is based at Brookhaven
 - Current version (FLX-172) in use at BNL for sPHENIX
- New board being prototyped now for ATLAS HL-LHC running.

FPGA: Xilinx Versal Prime XCVM1802

- PCIe Gen4 x16
- 24 Firefly links with 3 config options
 - 24 links @ 25Gbs
 - 24 links @ 10Gbs
 - 12 @ 25Gbs + 12 @ 10Gbs
- 4 additional Firefly links
- Front Panel: Fiber I/O
 - Additional 3 Inputs, 3 Outputs
- 1 DDR4 Mini-UDIMM
- USB-JTAG/USB-UART




Block Diagram of FLX-182 Prototype

Notes: The Xilinx Versal Prime is a SoC.- Dual Core ARM Cortex
Power usage ~133W. No power through PCIe.
Can be implemented as a stand alone device (no Server)

FELIX timelines

10 boards were produced in early 2023.


DAQ Working Group has recently acquired one of these boards for testing



General technical support for the FELIX Hardware should be available through the end of HL-LHC operation.
(i.e. into the late 2030s for Run 5)

Plan for 48-ch FELIX



- FPGA: Versal Premium, e.g. VP1552
- Transceivers: Up to 100+ GTYP/GTM
- PCIe Gen 5 up to 16 lanes
- If FPGA is available as planned, design will start in Q1 of 2023, first board is expected to be available in Q3 2023.



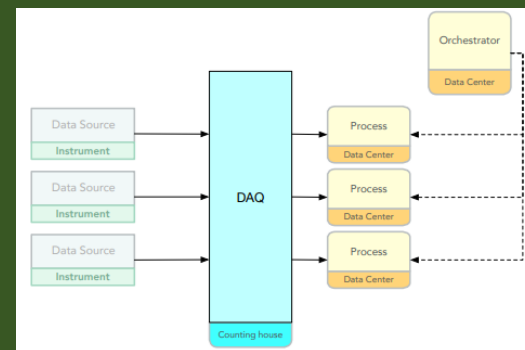
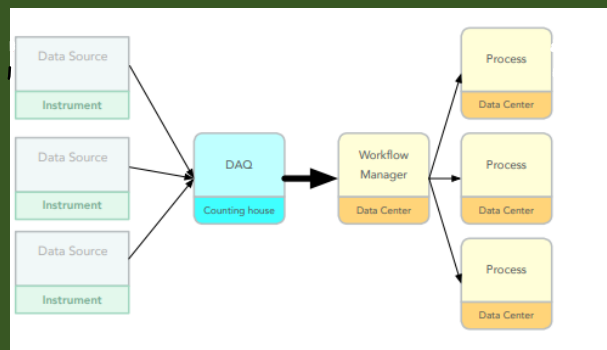
S&C Streaming DAQ WG

- Two streaming concepts
 - (Electronic) No Level-0 trigger
 - (Institutional) Rapid turn around of reconstructed data (streaming as opposed to batch processing)
- The focus of the group is the institutional requirements and organization
 - International Collaboration
 - Tasks to be done and associated institutional requirements
 - Support for streaming model
 - Automated data handling
 - Automated calibration tools
 - Automated monitoring tools
- We expect this may lead also to links between the DAQ and the streaming reconstruction for rapid turn around of critical information such as calibrations

Batch options: (reconstruction is human driven)



Streaming options: (reconstruction is data driven)



Graham Heyes