Summary of Streaming Computing Workfest at the collaboration meeting

STREAMING COMPUTING MODEL WG:

MARCO BATTAGLIERI, LANDGRAF JEFF, TAKU GUNJI (QNSI-TOKYO)

SOFTWARE AND COMPUTING COORDINATION:

MARKUS DIEFENTHALER, TORRE WENAUS





東京大学 大学院
理学系研究科・理学部
SCHOOL OF SCIENCE, THE UNIVERSITY OF TOKYO

Streaming DAQ and Computing

- **Compute–Detector Integration to Maximize & Accelerate Science**
  - Capture every collision signal, including background
    - Event selection using all available detector data for holistic reconstruction
  - Rapid turnaround of 2-3 weeks for data for physics analyses
    - Timeline driven by alignment and calibration
  - We need to deploy real-time reconstruction and real-time calibration/alignment
    - **Joint efforts between detectors, DAQ, and computing**

Streaming readout for continuous data flow of the full detector information.	**AI** for rapid processing (autonomous alignment, calibration, and validation).	**Heterogeneous computing** for acceleration (CPU, GPU).

Markus Diefenthaler

Focus at the workfest

- **This workfest was focused on the near-term goals of the SRO group.**
  - Finalizing the requirements document for streaming computing.
  - Building testbeds
    - **Streaming Orchestration**
      - Developing Echelon0 (E0)-Echelon1 (E1) streaming workflows utilizing Rucio and PanDA.
    - **Alignment and Calibration workflows**
      - Executing complex calibration workflows with their dependencies.
    - **Streaming Reconstruction**
      - Raw data stream to collision event identification to reconstruction and analysis
    - **Streaming analysis, and potentially others**
      - Demonstrate simulation data production streaming to Echelon2 (E2) site.

Rucio: Distributed *data management* (replication, transfer, catalog)

PanDA: Distributed *workload management* (job scheduling, execution, recovery)

Agenda of the workfest

08:00	**Streaming Computing R&D in Japan**	*Prof. Shinsuke Ota*	*L102, Thomas Jefferson National Accelerator Facility*	08:00 - 08:20
	Discussion		*L102, Thomas Jefferson National Accelerator Facility*	08:20 - 08:30
	ePIC Calibration and Alignment: Workflow Overview and Requirements	*Marco Battaglieri et al.*	*L102, Thomas Jefferson National Accelerator Facility*	08:30 - 08:50
	ePIC Alignment and Calibration: Calorimeter Use Case	*Carlos Munoz Camacho*	*L102, Thomas Jefferson National Accelerator Facility*	08:50 - 09:00
09:00	**ePIC Alignment and Calibration: Tracking Use Case**	*Dr Joe Osborn*	*L102, Thomas Jefferson National Accelerator Facility*	09:00 - 09:10
	Autonomous Alignment and Calibration Workflows	*Torri Jeske*	*L102, Thomas Jefferson National Accelerator Facility*	09:10 - 09:30
	Discussion		*L102, Thomas Jefferson National Accelerator Facility*	09:30 - 09:50
10:00	**Coffee Break**		*L102, Thomas Jefferson National Accelerator Facility*	09:50 - 10:20
	Streaming Orchestration Using PanDA and Rucio	*Maxim Potekhin et al.*	*L102, Thomas Jefferson National Accelerator Facility*	10:20 - 10:40
	Discussion		*L102, Thomas Jefferson National Accelerator Facility*	10:40 - 11:00
11:00	**JANA2 Updates for Streaming Computing**	*Nathan Brei*	*L102, Thomas Jefferson National Accelerator Facility*	11:00 - 11:20
	Discussion		*L102, Thomas Jefferson National Accelerator Facility*	11:20 - 11:30
	Streaming Reconstruction in EICrecon	*Dr Takuya Kumaoka*	*L102, Thomas Jefferson National Accelerator Facility*	11:30 - 11:50
	Discussion		*L102, Thomas Jefferson National Accelerator Facility*	11:50 - 12:00

Streaming computing R&D in Japan

Alignment and Calibration workflows

Streaming Orchestration

Streaming Reconstruction

Streaming computing R&D in Japan

SPADI Alliance

Shinsuke Ota

Signal processing and data acquisition infrastructure alliance
toward the standardization for sustainable developments

>140 researchers (23 institute)

Strategy committee
Chair : Shinsuke OTA
Vice-Chair: Ryotaro Honda and Hidetada Baba

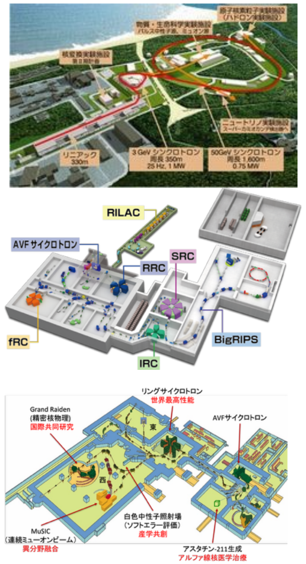

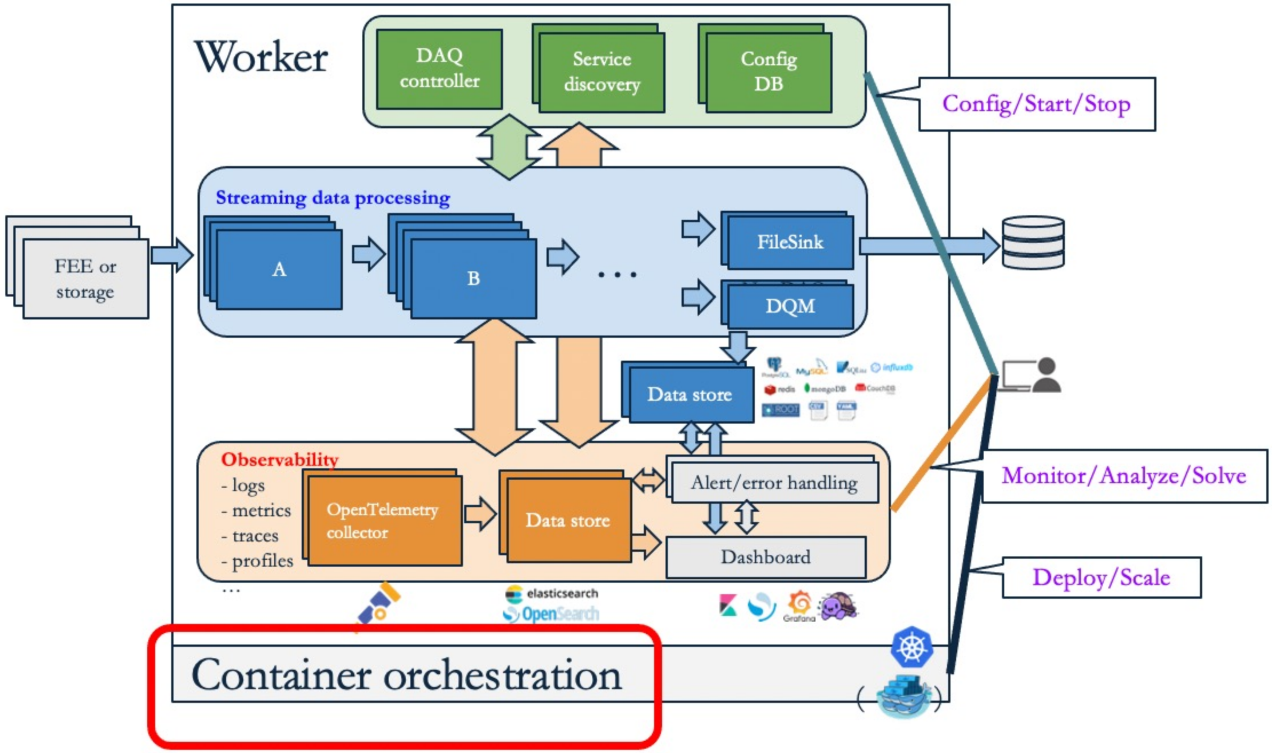

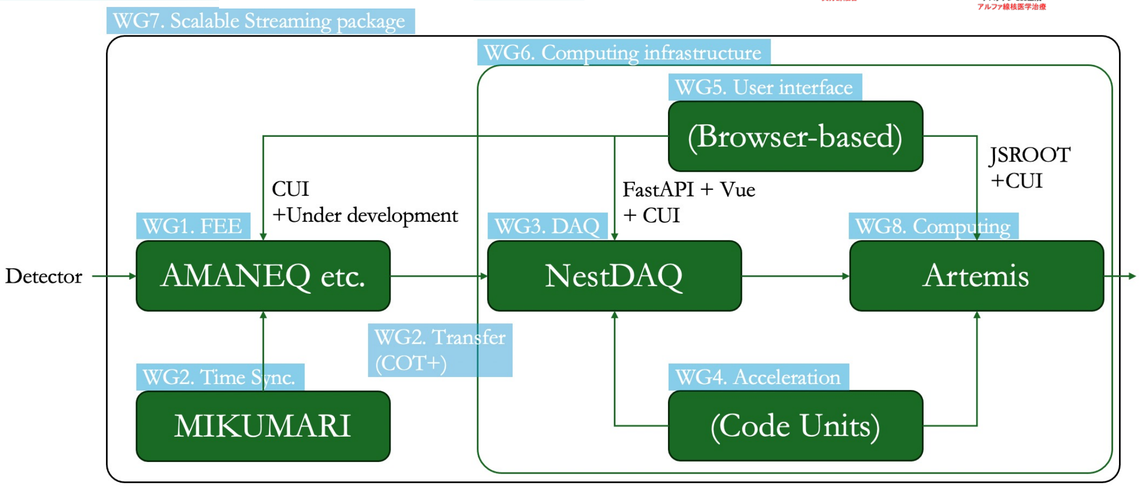

- **Toward the scalable system covering from table top to the EIC**
  - **Common framework for Observability**
  - **Orchestration of the distributed (and sub) Containers**
  - **Runnable over errors**
- **Will contribute to the testbed for orchestration, reconstruction and observability.**

Streaming Orchestration objectives

Maxim Potekhin, Torre Wenaus

- **Unified data flow**
  - Move bulk and monitoring data from E0→E1, trigger prompt processing on arrival.
- **Streaming processing**
  - Process data in E1 continuously, not by run/job/file units.
- **Low-latency feedback**
  - Continuous detector and accelerator monitoring with automated workflows.
- **State-aware orchestration**
  - Use detector/data state machine to coordinate calibration and reconstruction.

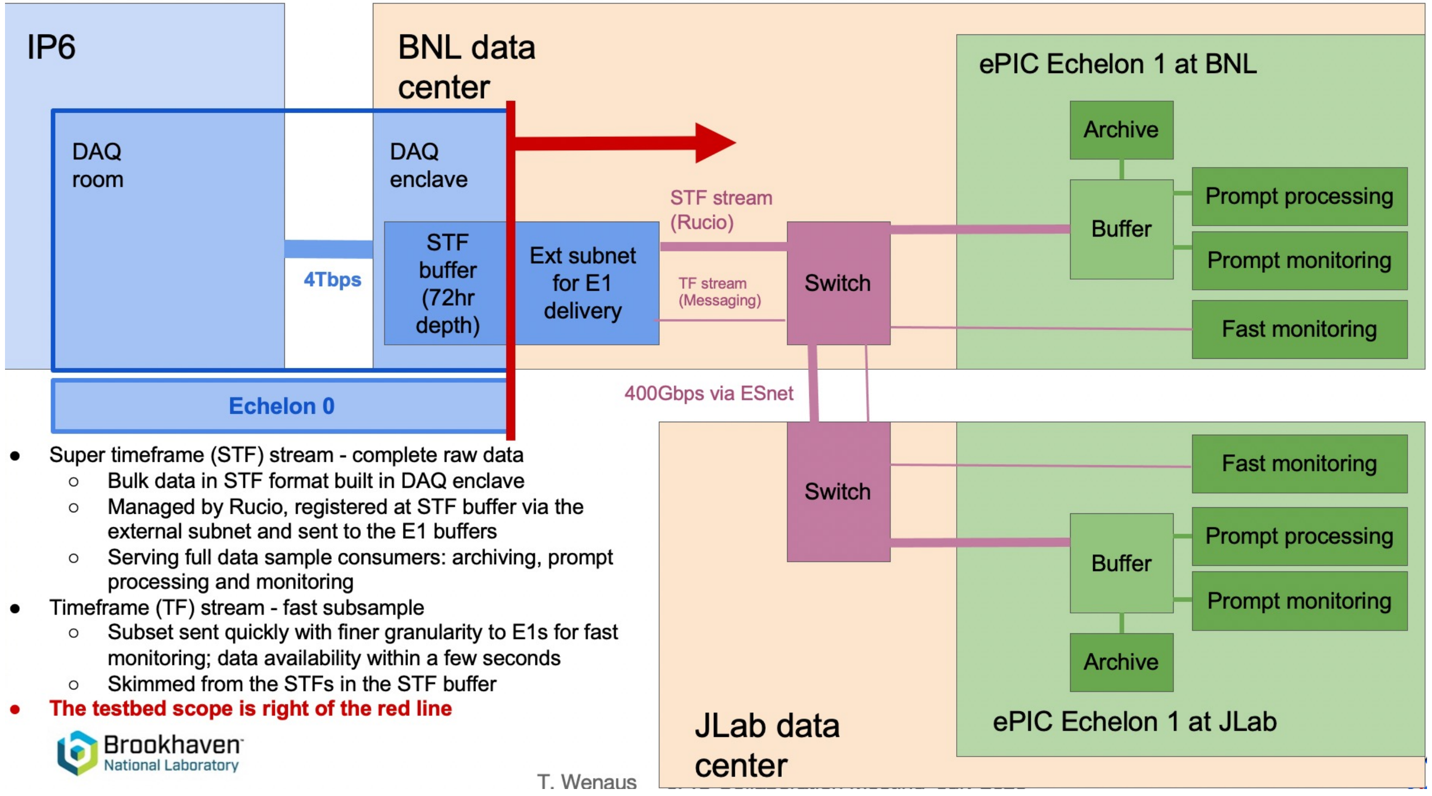

Work has begun in all the areas except

- E1 calibration orchestration
- E2 extension

Design and implementation

Testbed schematic

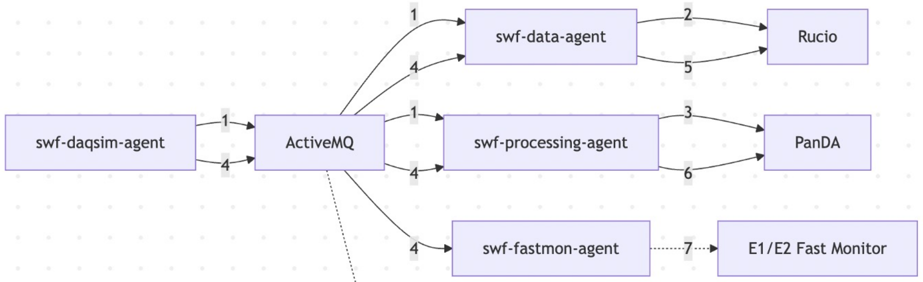

1. Run Start - daqsim-agent generates a run start broadcast message indicating a new datataking run is beginning
2. Dataset Creation - data-agent sees the run start message and has Rucio create a dataset for the run
3. Processing Task - processing-agent sees the run start message and establishes a PanDA processing task for the run
4. STF Available - daqsim-agent generates a broadcast message that a new STF data file is available
5. STF Transfer - data-agent sees the message and initiates Rucio registration and transfer of the STF file to E1 facilities
6. STF Processing - processing-agent sees the new STF file in the dataset and transferred to the E1 by Rucio, and initiates a PanDA job to process the STF
7. Fast Monitoring - fastmon-agent sees the broadcast message that a new STF data file is available and performs a partial read to inject a data sample into E1/E2 fast monitoring

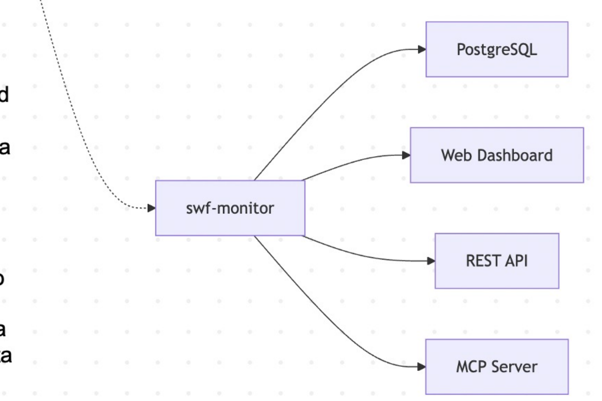

- Code and core services (PanDA, Rucio, etc.) are ready at BNL, and integration on the testbed has begun.
- **Next steps: integrate and operate components, enable full monitoring, and test prompt and concurrent STF/TF workflows. Then extend to calibration, fine-grained streaming, and multi-site operation (E1–E2 links).**

Monitor and database

- Django based testbed monitor manages a postgres database
- All system information aggregates there
- Monitor provides an authenticated REST interface to exchange data with agents
- Monitor listens in on system activit
and presents it
  - ActiveMQ messaging
  - Python logging
- Monitor will provide MCP service t
provide system info to LLMs, and later take direction
  - In progress
  - MCP services also being implemented in core PanDA, we will integrate them

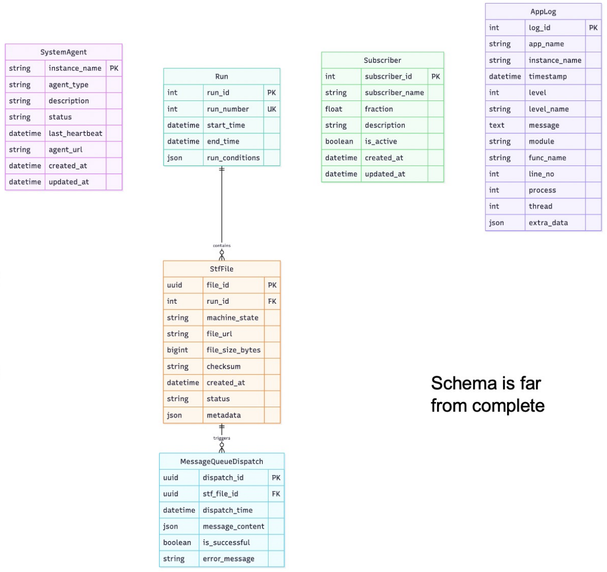

Schema is far from complete

State machine first version

- We came up with a 'minimal' state machine to describe the basics of the system

States
- no_beam
  - Collider not operating
- beam
  - Collider operating
- run
  - Physics running
- calib
  - Dedicated calibration period
- test
  - Testing, debugging
  - Any substates can be present during test

Substates
- not_ready
  - detector not ready for physics datataking
  - occurs during states: no_beam, beam, calib
- ready
  - collider and detector ready for physics, but not declared as good for physics
  - when declared good for physics, transitions from beam/ready to run/physics
  - occurs during states: beam
- physics
  - collider and detector declared good for physics
  - if collider or detector drop out of good for physics, state transitions out of 'run' to 'beam' or 'off'
  - occurs during states: run
- standby
  - collider and detector still good for physics, but standing by, not physics datataking (dead time!)
  - occurs during states: run
- lumi
  - detector, machine data that is input to luminosity calculations
  - occurs during states: beam, run
- eic
  - machine data, machine configuration
  - occurs during states: all
- epic
  - detector configuration, data
  - occurs during states: all
- daq
  - info, config transmitted from DAQ
  - occurs during states: all
- calib
  - a catch-all for a great many calib data types, we can start small
  - occurs during states: all (assuming there are cases where calib data is taken during beam on)

Brookhaven National Laboratory

Streaming Reconstruction

- **JANA2 updates**

Nathan Brei

  - Extend functionalities to cope with multilevel events to gain the flexibility

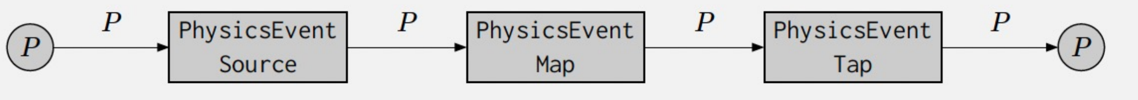

Many different information are interleaved with or bundled in timeframe

EICrecon timeframe splitting

Supporting Timeframe

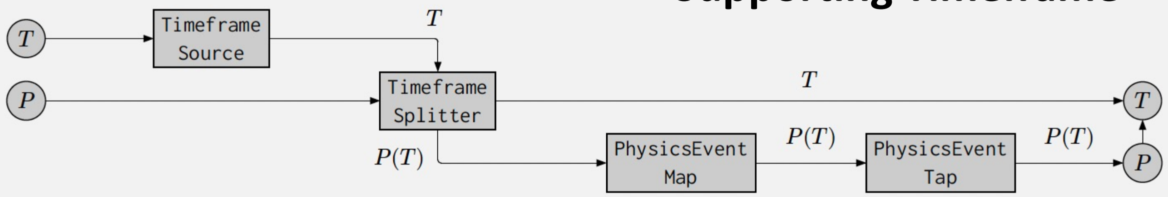

Multilevel sources with timeframe splitting

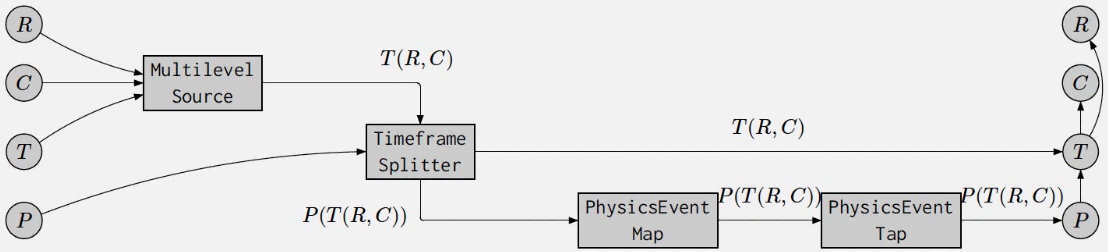

Introducing multilevel sources

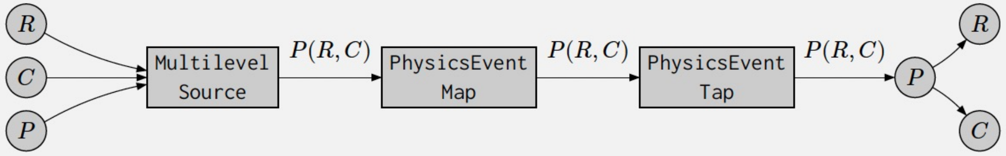

Supporting Run information, Slow Controls, etc

Timeframe sources with multilevel splitting

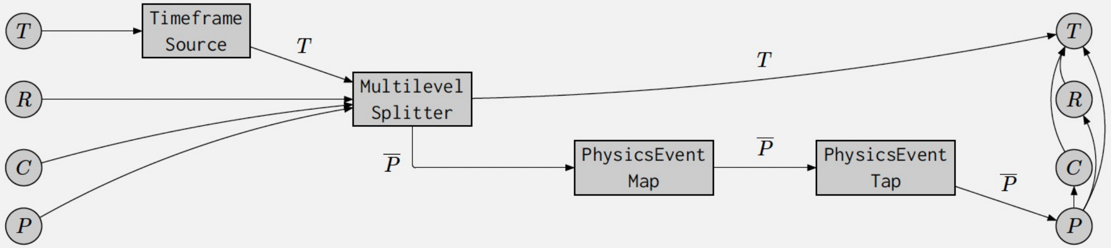

Streaming Reconstruction

Takuya Kumaoka

- **Development of algorithm for Event-Building from streamed data**

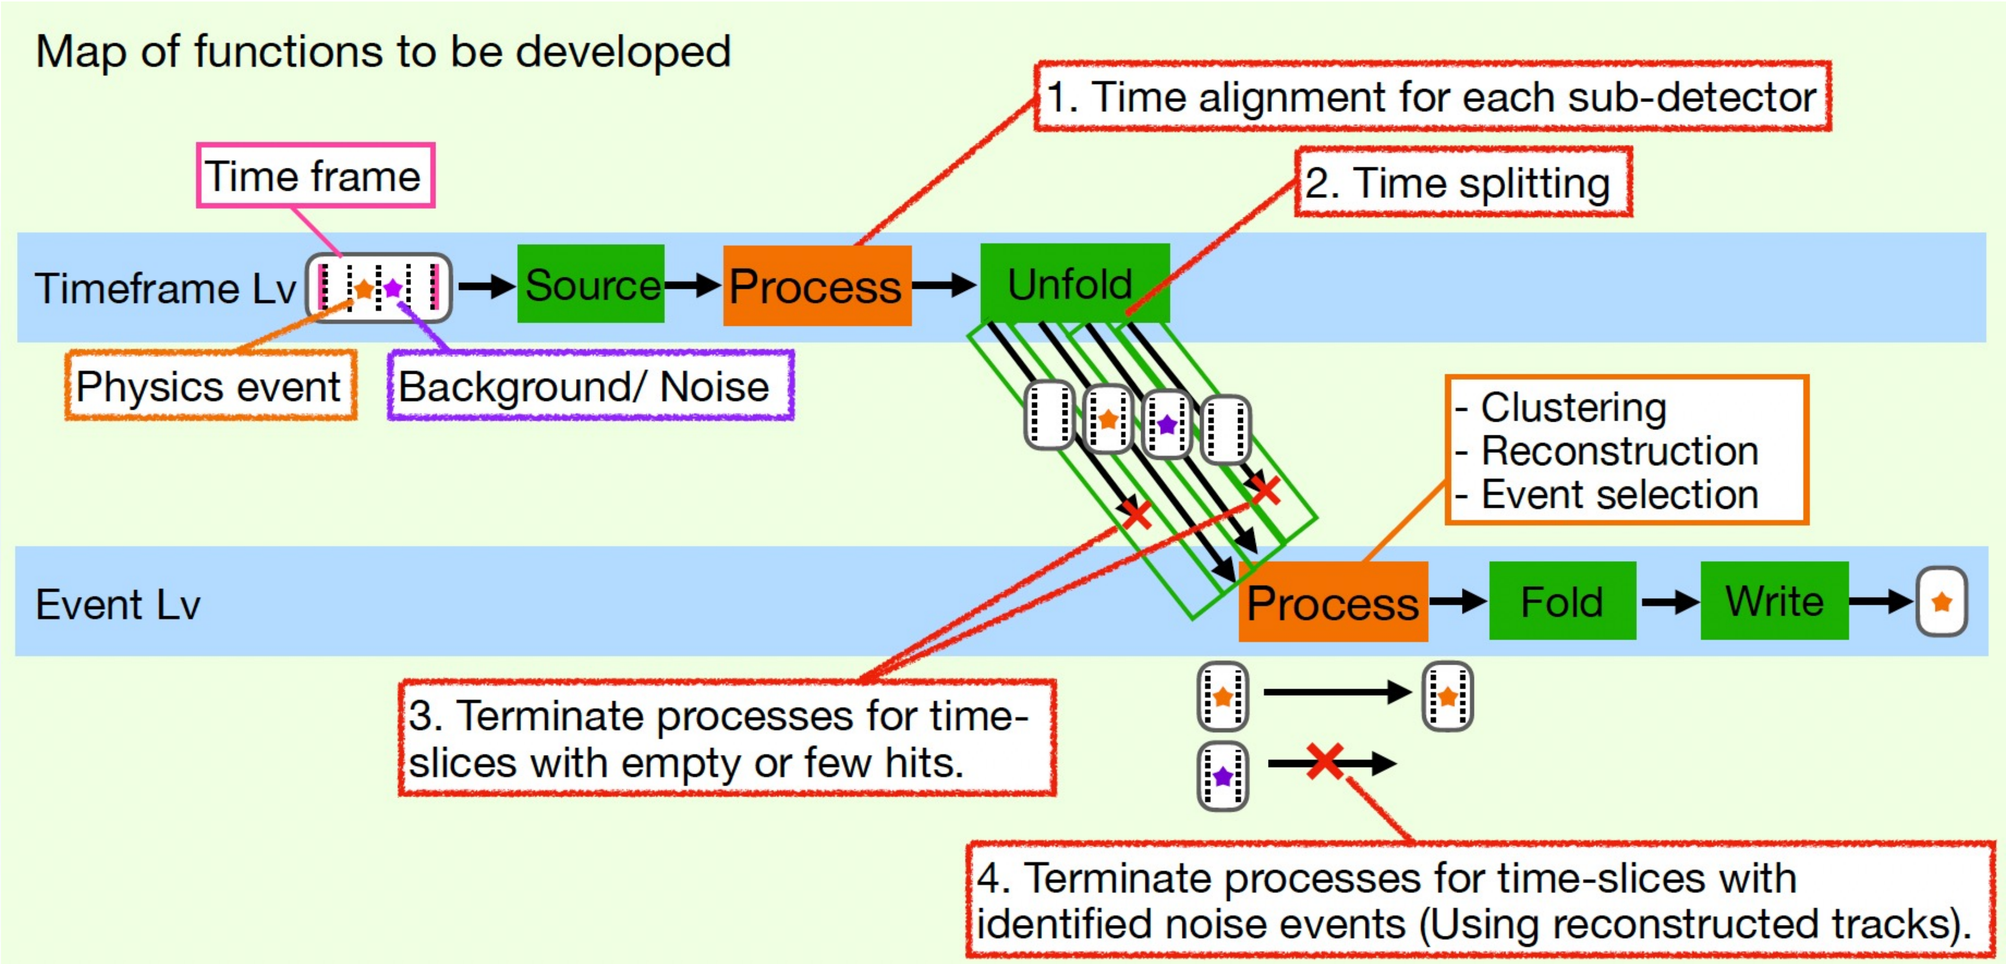

Streaming Reconstruction

- **Look for the coincidence**

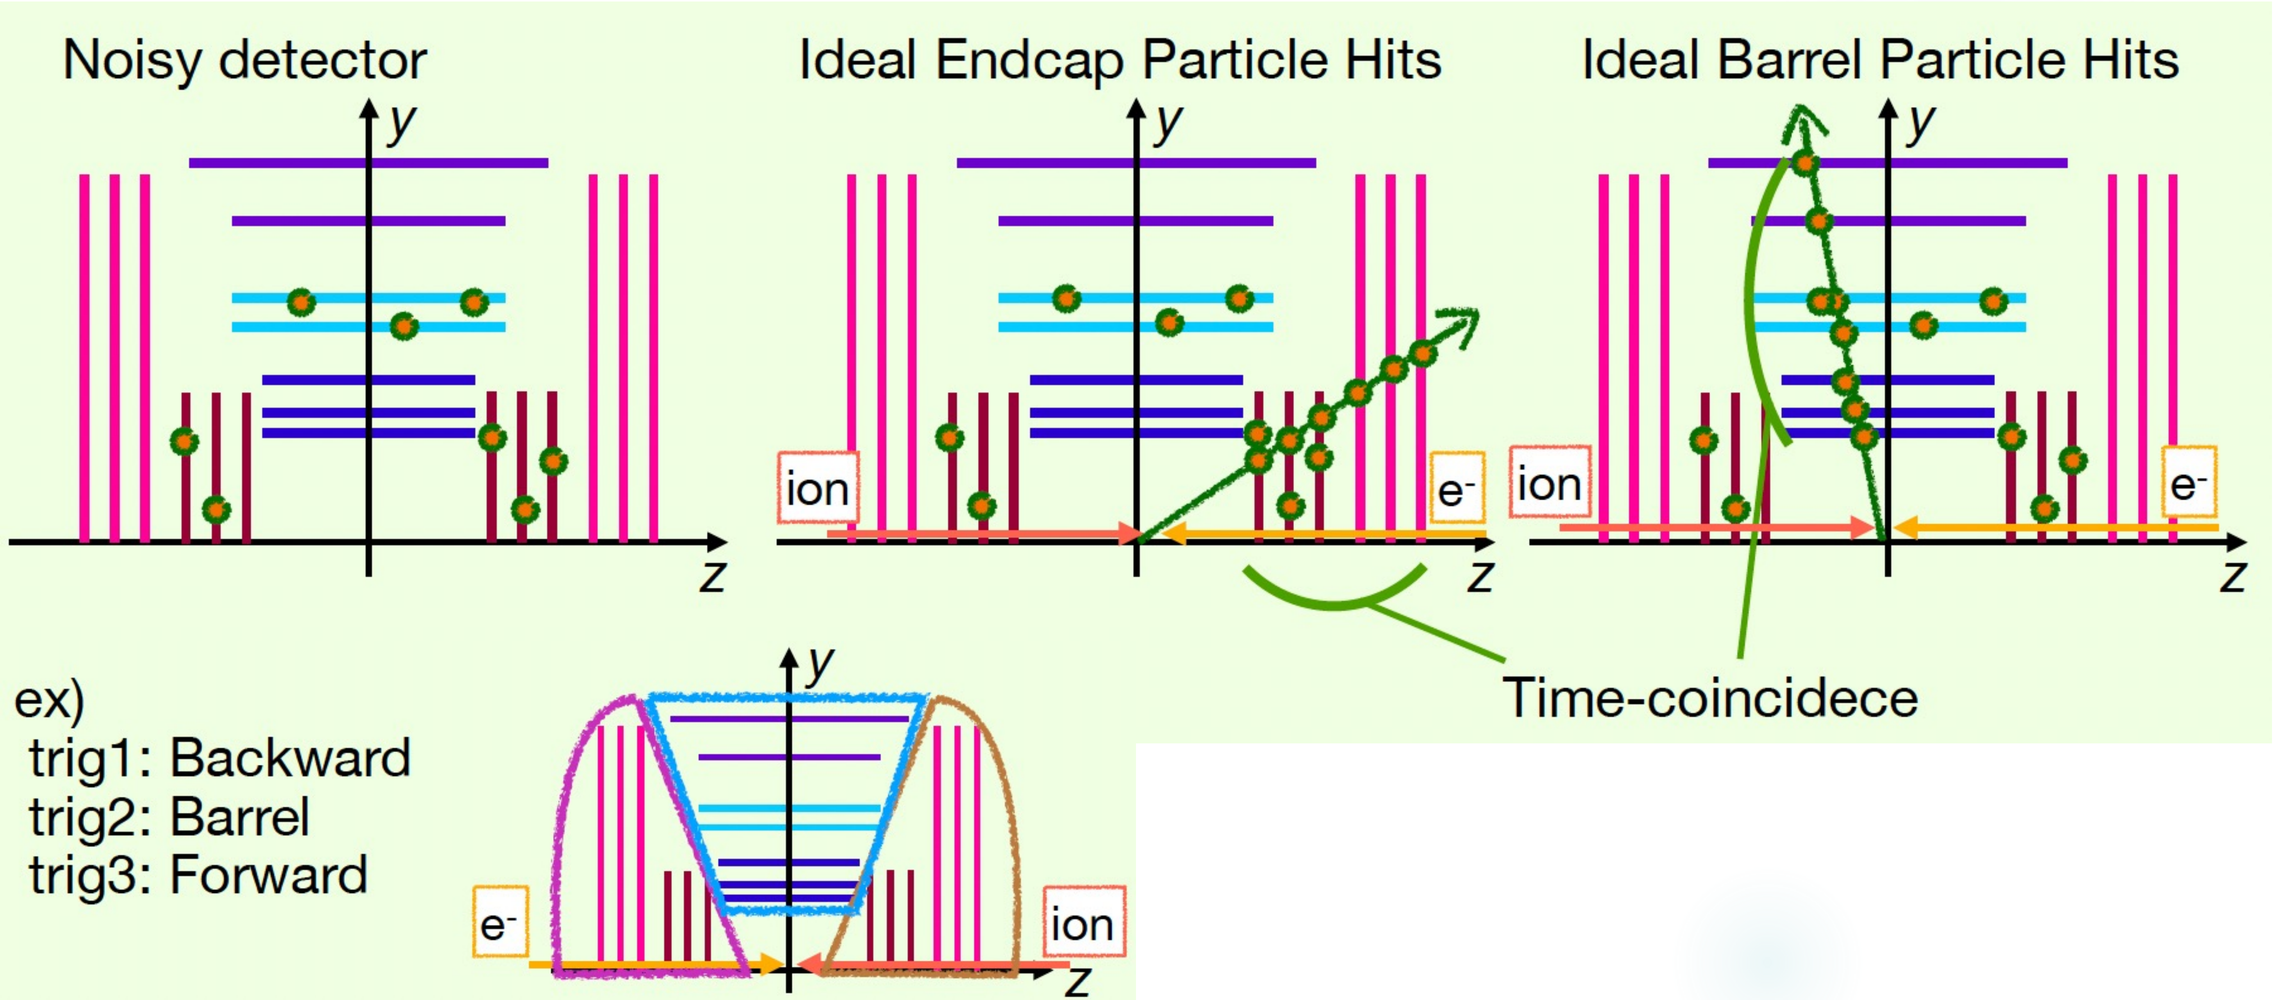

Streaming Reconstruction

Online selection

Input simulation file setting (1530 events) (**Root file page**)

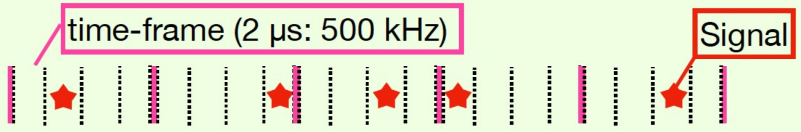

Symbol	Process	Description	Sampling Frequency (kHz)	Status Code Shift
🔴	signal	DIS NC 18x275 Q^2>1 (Deep inelastic scattering neutral current)	500	0
⚪	synrad	Synchrotron Radiation	14000	2000
🔵	ebrems	Electron bremsstrahlung radiation	316.94	3000
🟢	etouschek	Electron Touschek scattering (intrabeam scattering)	1.3	4000
🟡	ecoulomb	Electron Coulomb scattering processes	0.72	5000
🟣	p.b.gas	Proton beam gas interactions	22.5	6000

Trigger ratio for the physics event region and background region

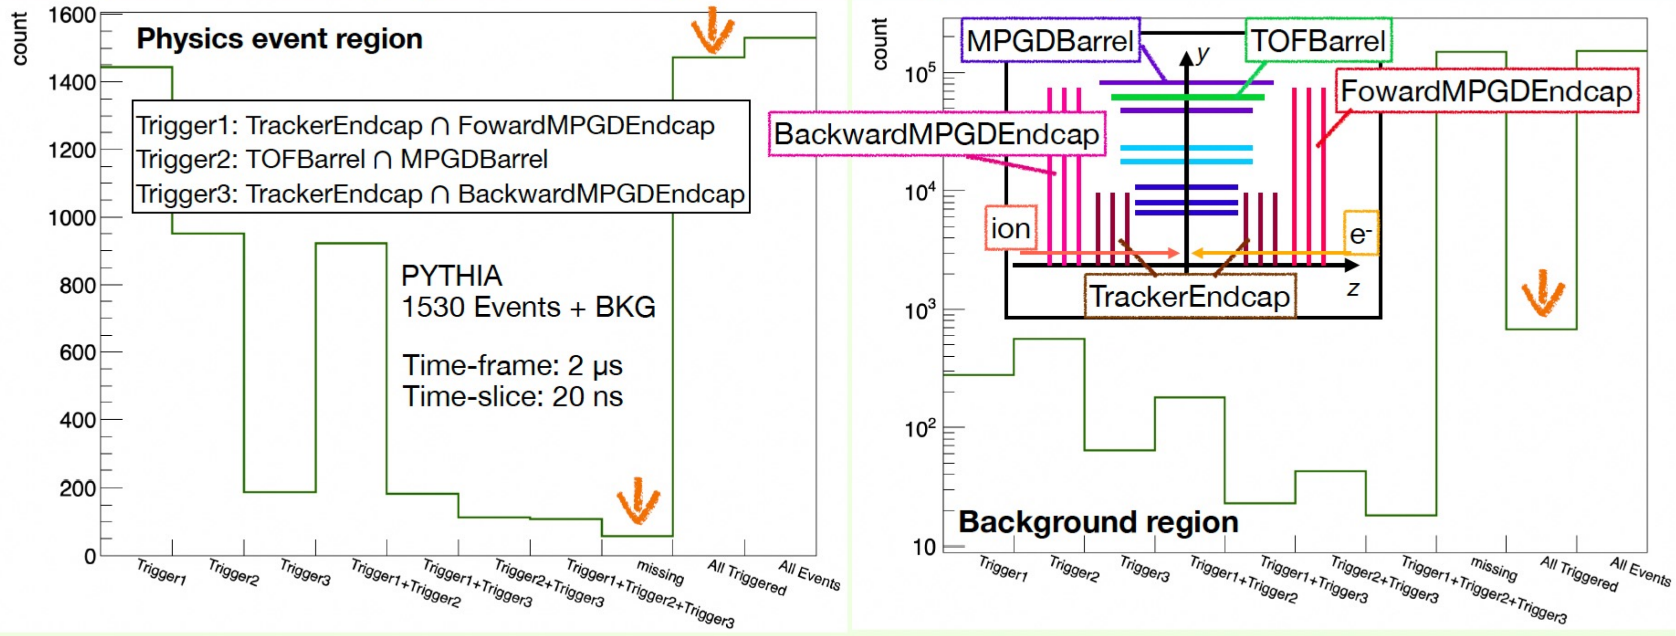

Physics event: The triggers can capture over 96% (1473/1530)
Background: The rejection ratio is about 99% 674/153000

Next step:

- **improving the efficiency of the baseline algorithm and testing it with a blind analysis.**
- **Include detector (timing) response and evaluate based on fully digitized information**
- **Full tracking and PID reconstruction for time-splited data**
- **Import tracking / PID reconstruction algorithms in streaming reconstruction**

Close collaboration with detectors, tracking, and PID teams will be crucial.

Alignement and Calibration

Marco Battaglieri

- **Autonomous alignment and calibration will be crucial to achieve a rapid turnaround to full reconstruction (2-3 weeks)**
  - **Several WG meeting last years focused on calibration procedure of subsystems (Jin's table)**

We will discuss with each subsystem soon and see if they have some updates on their calibration plans.
After understanding the interdependences, we will start prototyping workflows (ex, tracking and calorimeters as use cases).

Subsystem	Region	Pre-physics-operation calibrations (Cosmic, no-beam calibration, commissioning)	Task	Human intervention ?	Data Needed	Dependecy	T0 + 12hr	T0 + 24hr	T0 + 36hr	T0 + 48hr	T0 + 60hr	T0 + 72hr	T0 + 84hr	T0 + 96hr	Monitoring	Computing resource	Post-reconstruction calibrations (applied at analysis stages)	Comment	Subsystem
MAPS	Barrel+Disk	Threshold Scan / ALICE=20min Fake rate scan/noisy pixel masking	(See Alignment)															TIC meeting: https://indico.bnl.gov/event/21648/	MAPS
MPGD	Barrel+Disk	?	?																MPGD
bTOF, eTOF (ac-lgad)	Barrel/Forward	Bias voltage determination ASIC baseline, noise, threshold Clock sync Time walk calibration	Gain calibration TDC bin width determination Clock offset calibration Hit position dependency (intrinsic and c-by-c)	QA	High p tracks ~1hr of production data?	Tracking, pfRICH	Data Acc. Dependen	Dependen	Processin	Processing								SRO meeting https://indico.bnl.gov/event/21619/	bTOF, eTOF (ac-lgad)
Central Detector Tracker Alignment		Initial alignment	Alignment Check/Update (if needed)	QA	Prodcution data		Processing											SRO meeting https://indico.bnl.gov/event/21619/	Central Detector Tracker Alignme
pfRICH	Backward	Thresholds (noise dependent), dynamic range adjustments, timing offsets, synchronization Initial alignment	Alignment Check/Update (if needed) Time dependencies (Aerogel transparency, mirror reflectivity, Gas pressure)	?	Prodcution data		Data Acc.	Processing										TIC meeting: https://indico.bnl.gov/event/21648/	pfRICH
DIRC	Barrel	Laser data?	?	?														TIC meeting: https://indico.bnl.gov/event/21648/	DIRC
dRICH	Forward	Bunch timing offset scan Threshold scan Noise masking	Track based alignment	?	High p tracks ~1hr of of production data?	Tracking	Data Acc. Dependen	Processin	Processing									SRO meeting: https://indico.bnl.gov/event/22114/	dRICH
bEMC	Backward	Cosmic and LED for the initial gain balancing	DIS Electron Pi0->gg events energy scale	QA	DIS electron Pi0 di-photon resonance ~1 day of production data	Tracking	Data Acc. Dependen	Data Acc.	Processin	Processing					LED			SRO meeting: https://indico.bnl.gov/event/22412/ Carlos: aiming 1% precision Planning for LED flash during production run, processi	bEMC
AstroPix	Barrel																	TIC meeting: https://indico.bnl.gov/event/21648/	AstroPix
ScifiPb	Barrel		SiPM gain		?													TIC meeting: https://indico.bnl.gov/event/21648/	ScifiPb
fEMC	Forward		Pi0, eta->gg events energy scale				Data Acc.	Data Acc.	Processin	Processing							High energy cluster non-linearity	SRO meeting: https://indico.bnl.gov/event/22412/ Need pi0 filtered data for automated calibration AI driven calibration?	
		IV Scan	Second iteration pi0 (if needed)	QA	Pi0 di-photon resonance ~1 day of production data						Processing				LED				fEMC
bHCAL	Backward	LED	?															TIC meeting: https://indico.bnl.gov/event/21648/	bHCAL
cHCAL	Barrel	MIP calibration Gain calibration	(See hadronic e-scale calib)															SRO meeting: https://indico.bnl.gov/event/21785/	cHCAL
fHCAL	Forward																		fHCAL
fHCAL insert	Forward																		fHCAL insert
Hadronic energy scale calibration		?	Set full calo stack energy scale for hadroinc shower and jets	?	High energy hadronic showers and jets	Tracking h-PID	Data Acc. Dependen	Data Acc. Dependen	Data Acc. Dependen	?	?	?	?	?			Final energy scale calibration (if needed)	Comments from Oleg during SRO meeting: https://indi	Hadronic energy scale calibration
low Q2 Tagger	Far Backward	Alignment?																TIC meeting: https://indico.bnl.gov/event/22079/	low Q2 Tagger
low Q2 Tagger (CAL)	Far Backward																	TIC meeting: https://indico.bnl.gov/event/22079/	low Q2 Tagger (CAL)
Pair Spec Tracker	Far Backward																	TIC meeting: https://indico.bnl.gov/event/22079/	Pair Spec Tracker
Par Spec Cal	Far Backward																	TIC meeting: https://indico.bnl.gov/event/22079/	Par Spec Cal
Direct Photon Cal	Far Backward																	TIC meeting: https://indico.bnl.gov/event/22079/	Direct Photon Cal
B0 Tracking	Far Forward	Survey alignment/Cosmic	Alignment check		MIP		Processing											SRO/FF meeting https://indico.bnl.gov/event/22676/	B0 Tracking
B0 PbWO4	Far Forward	Survey alignment/Cosmic	SiPM gain		MIP/Gamma/Electrons		Processing								LED			SRO/FF meeting https://indico.bnl.gov/event/22676/	B0 PbWO4
Roman (Pots)	Far Forward					Acc. BPM Potential use of vertex of central detector	Data Acc. Dependen	Processing										SRO/FF meeting https://indico.bnl.gov/event/22676/	Roman (Pots)
Off Momentum	Far Forward	laser/survey alignment Low lumi running	beam position monitors/fill by fill correction		MIP rate distribution in RP		Data Acc. Dependen	Processing										SRO/FF meeting https://indico.bnl.gov/event/22676/	Off Momentum
ZDC PbWO4	Far Forward	Survey alignment, timing delay	SiPM/APD gain, timing	QA	Photon		Processing								LED			SRO/FF meeting https://indico.bnl.gov/event/22676/	ZDC PbWO4
ZDC Sampling	Far Forward	Survey alignment, timing delay	SiPM gain	QA	Single neutron		Processing								LED			SRO/FF meeting https://indico.bnl.gov/event/22676/	ZDC Sampling

Use cases : Backward ECal

Calibration with π^0 *Carlos Muñoz Camacho*

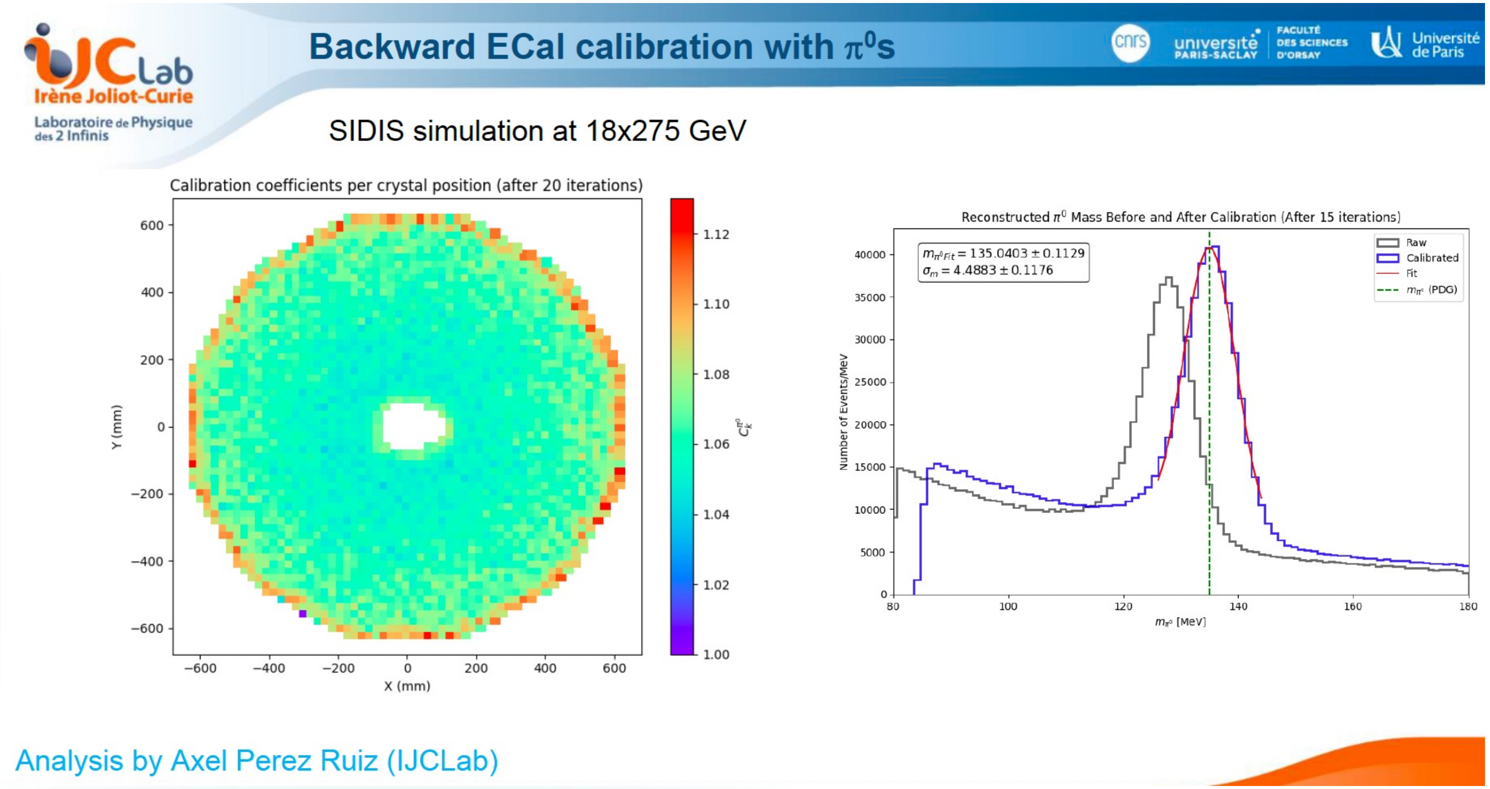

Methods successfully used with EMCals at JLab
Based on NIM A566 (2006) 366

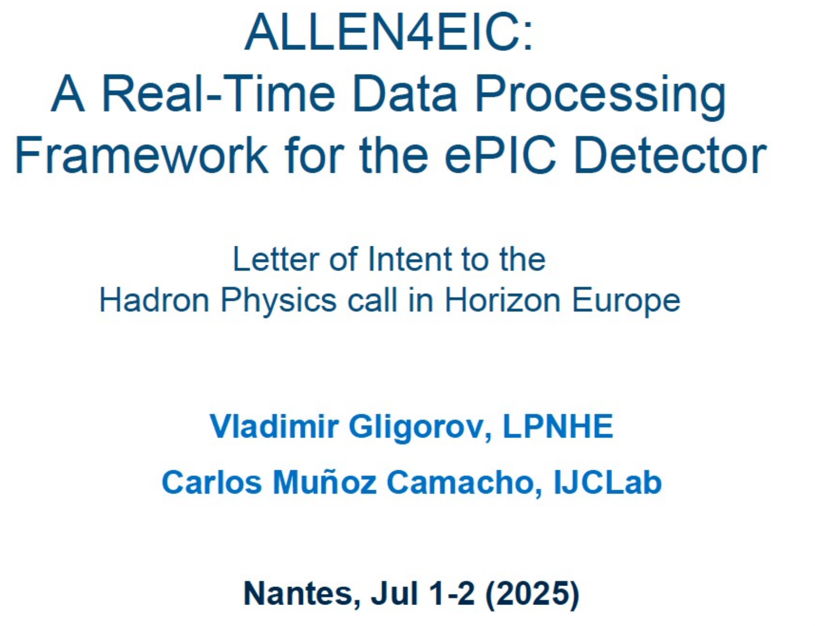

Allen Framework

- Developed for GPU-based high-level triggering at LHCb.
- Processes 30 million events/sec using commodity GPUs.
- Cross-platform (x86, CUDA, ARM); Python + C++/CUDA.
- Includes monitoring, high energy-efficiency.

We will investigate more on how we can advance Allen framework for ePIC: Adapt ePIC geometry, data format, streaming architecture, routines as in EICrecon

Use case: Tracking

Joe Osborn

- Still much offline development needed for alignment work in EICRecon
  - Transformation hooks in EICRecon
  - Calibration file format and storage (fit into broader calibration database structure)
  - Development (in Acts) of residual and derivative calculations (ongoing work, some examples exist)
  - Linking output into Millepede2 (or something else?)
  - Putting the workflow together

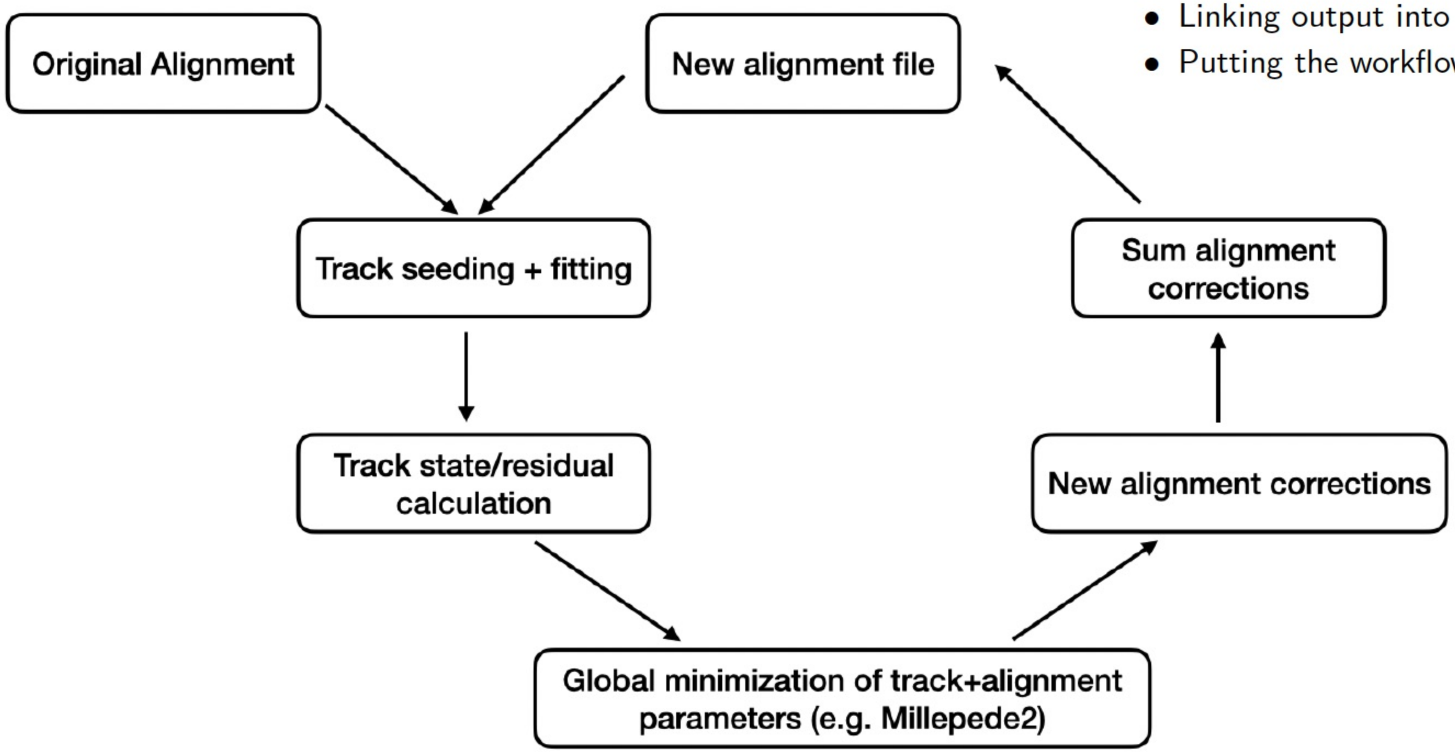

Our next steps:

- **Discussion with central tracking team**
- **Prototyping of workflows (close relation to streaming reconsturction)**
- **Estimation of required N_{evt} or N_{track}**

AI assisted calibration/control

Torri Jeske

Autonomous monitoring and control workflows, which significantly reduces calibration resources

Dynamic HV Control

RoboCDC **adjusts the anode high-voltage settings** in response to environmental conditions, **stabilizing the chamber gain** and reducing our dependence on offline calibrations.

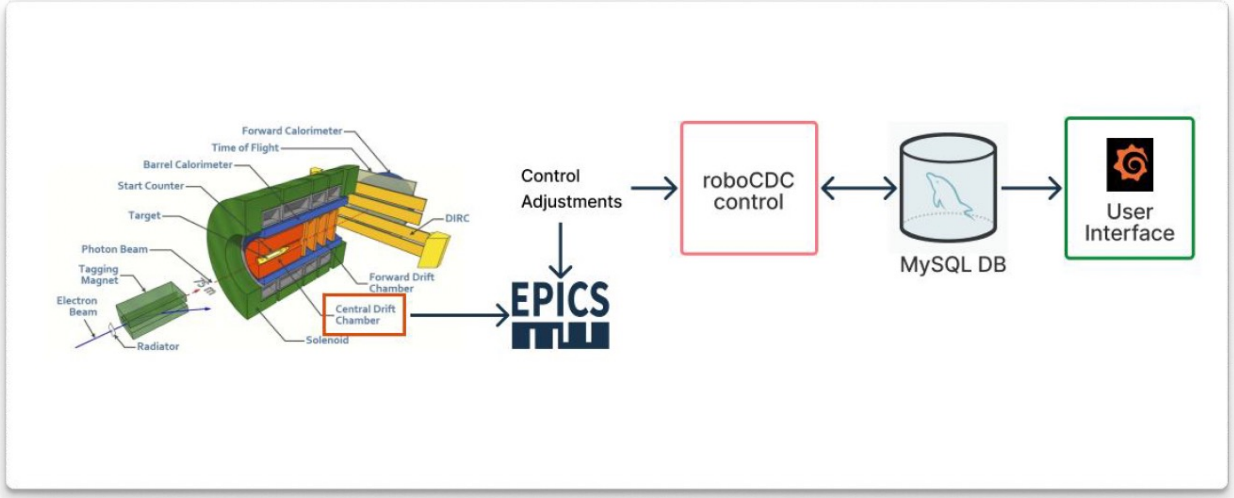

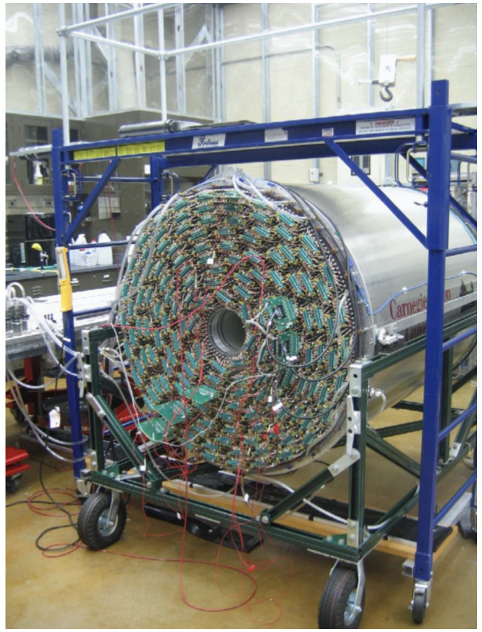

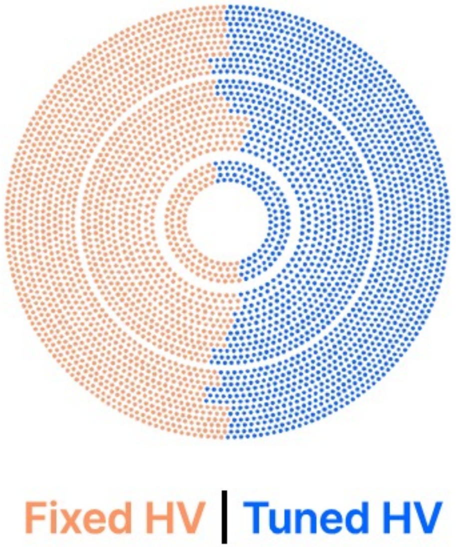

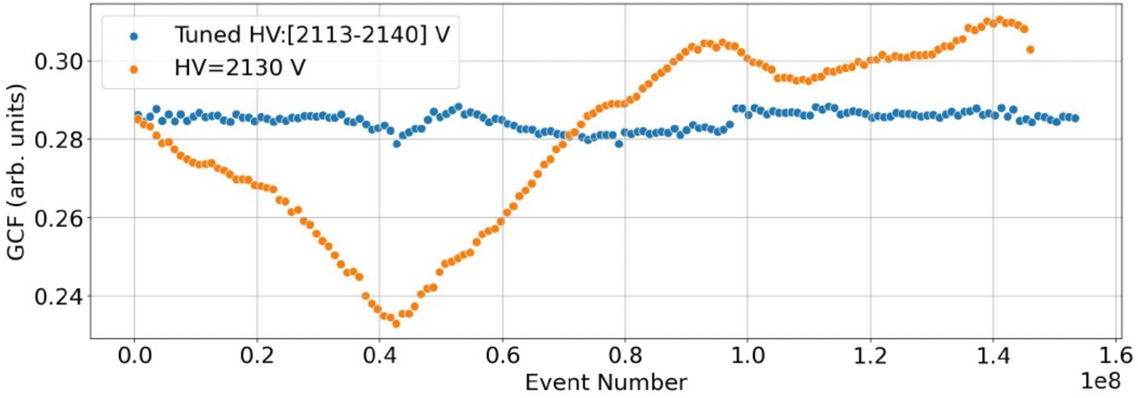

Summary and Outlook

- **We started integrating basic components (orchestration, reconstruction, etc) for SRO.**
- **Next steps:**
  - **Streaming Orchestration**
    - Integrate data flow control, monitoring, and prompt processing workflows
    - Move toward fully automated, low-latency workflows
  - **Streaming Reconstruction**
    - Implement realistic detector response and develop event identification and reconstruction
    - Communication with central tracking and PID team to integrate real-time tracking and PID
    - Optimize GPU/CPU hybrid processing for scalability at EIC
  - **Alignment and Calibration workflows**
    - Start contacting with subsystems to clarify requirements
    - Develop real-time calibration pipelines tied directly to streaming data
    - Use detector/data state models to trigger updates dynamically
- **We are laying the groundwork —let’s push forward together to build and realize a fully streaming and computing model for the EIC!**

Summary and Outlook

- **Upcoming events related to the development of the full computing model, from E0 through E2**
  - 09/17 Echelon 2: Roles and Resources; Ongoing Activities @ Software & Computing meeting
  - 10/08 Mini-Workshop at JLab on Echelon 0-1 Data Flow and Processing.

Backup slides

Streaming DAQ and Computing

Collider Characteristics

- 1260 Bunches arriving at 98.5Mhz (10.15ns bunch separation)
- 1.015us abort gaps (100 bunches)
- $\sqrt{s} \Rightarrow$ 20 – 141 GeV
- $\mathcal{L}_{max} \Rightarrow 10^{34}$ cm^{-2} s^{-1}
- Electron, proton, and light nuclei beams can be polarized
- Each bunch can have different polarization states
  - DAQ must tag data to specific bunch crossings
  - Need to track luminosity for each bunch crossing

Physics Performance

- Maximum DIS rate ~500kHz
- Large number of Channels
- Low occupancy

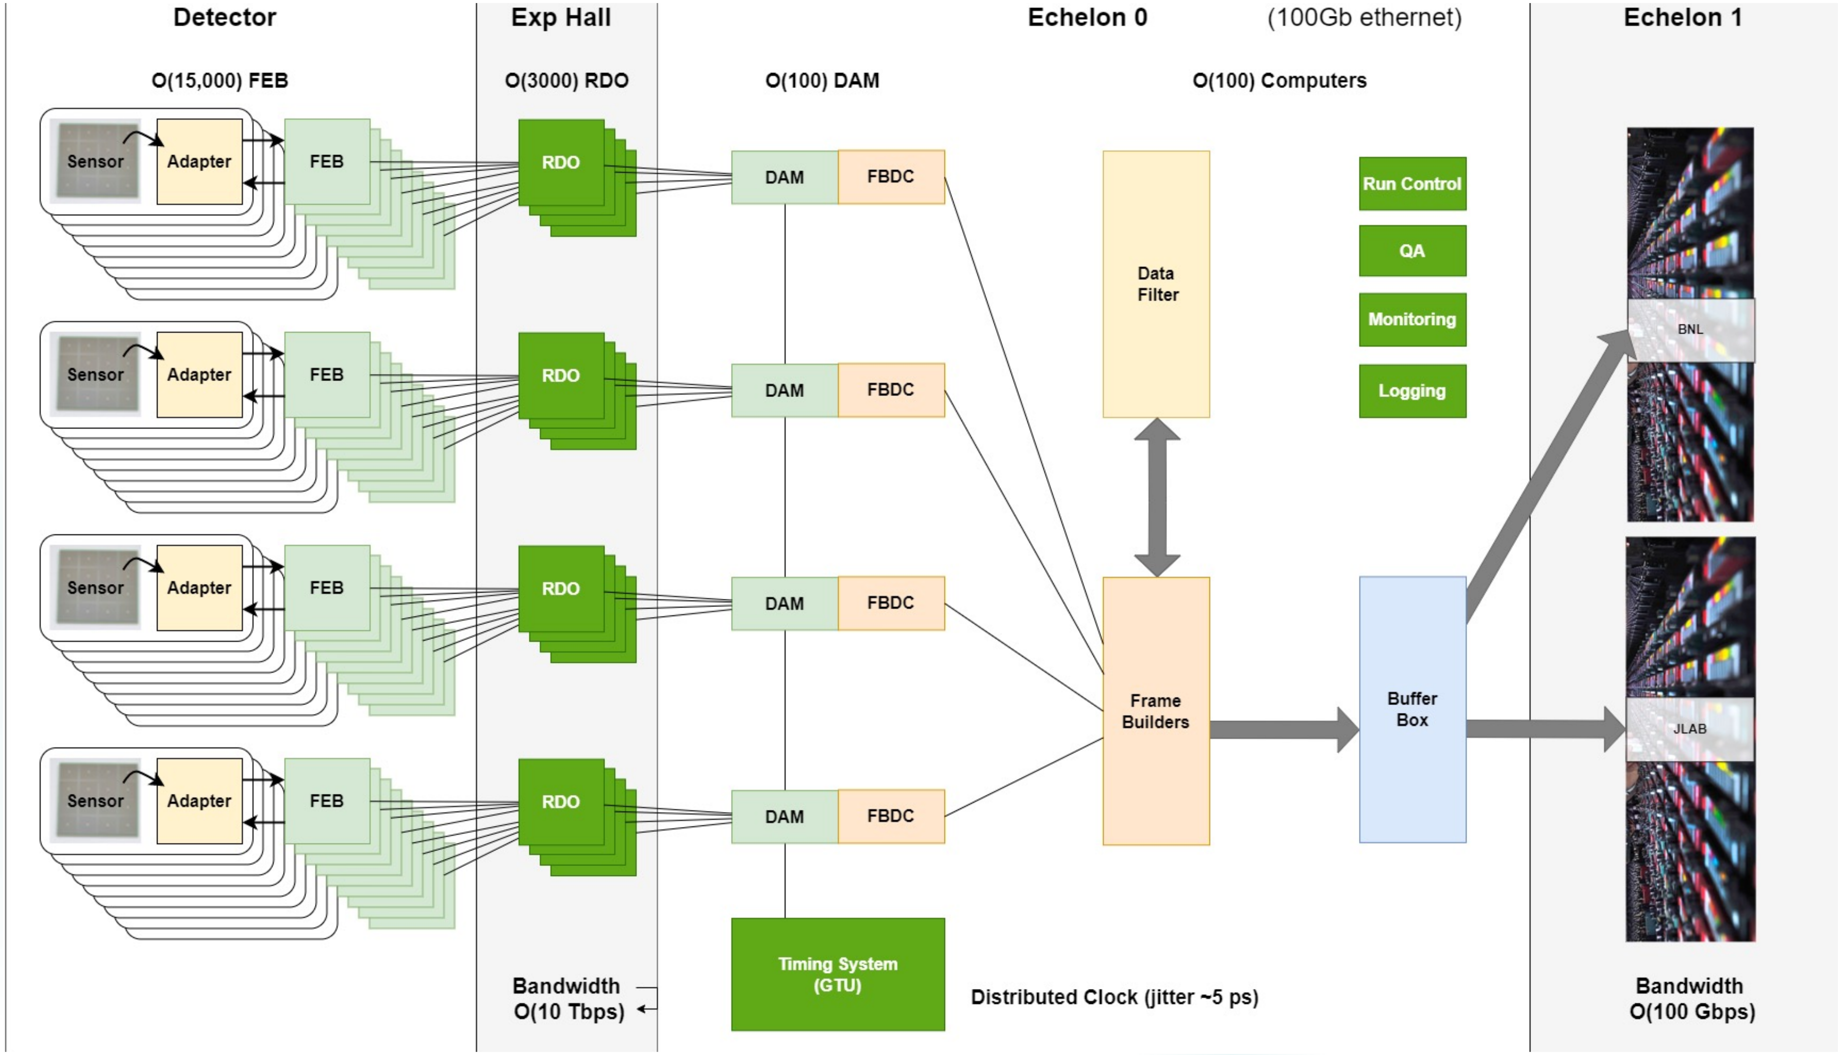

Streaming Computing model

https://zenodo.org/records/14675920

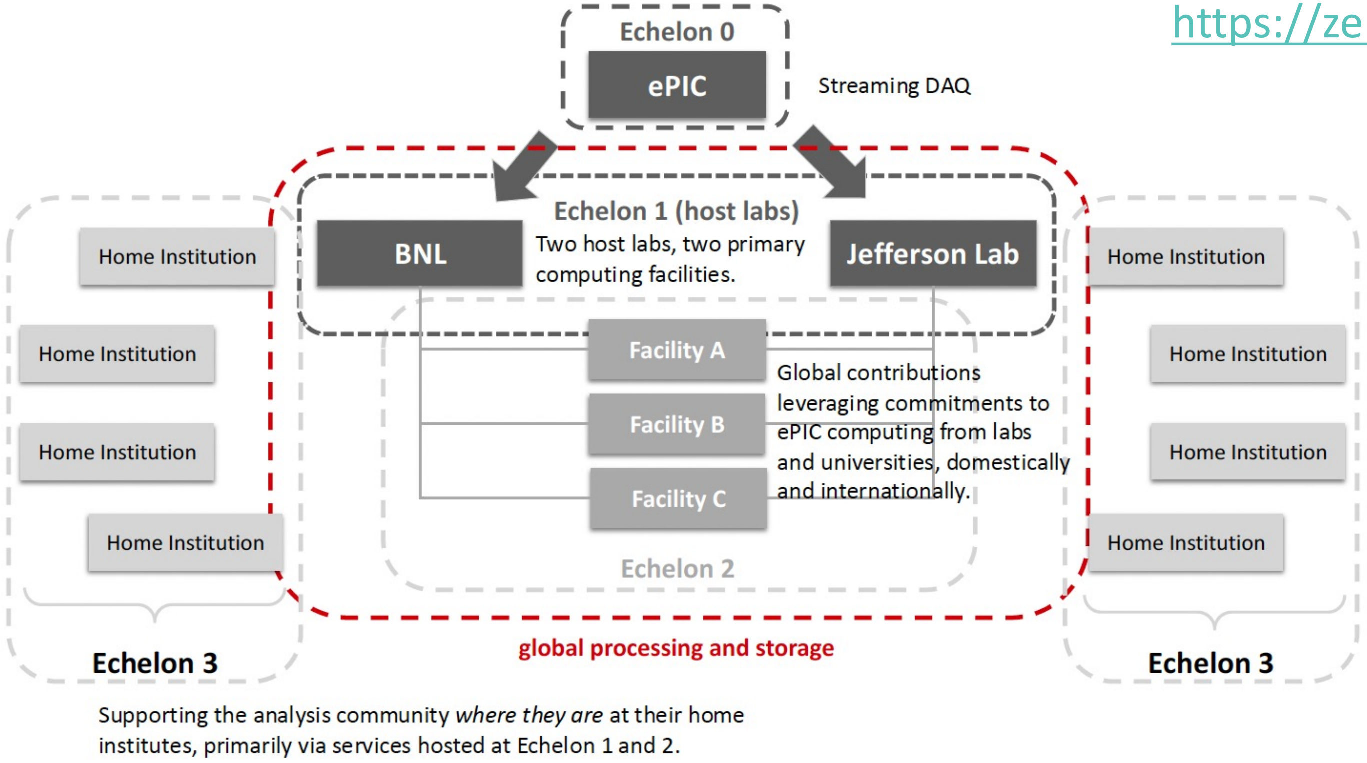

Use Case	Echelon 0	Echelon 1	Echelon 2	Echelon 3
Streaming Data Storage and Monitoring	✓	✓		
Alignment and Calibration		✓	✓	
Prompt Reconstruction		✓		
First Full Reconstruction		✓	✓	
Reprocessing		✓	✓	
Simulation		✓	✓	
Physics Analysis		✓	✓	✓
AI Modeling and Digital Twin		✓	✓	

- Echelon 1 sites perform the low-latency streaming workflows consuming the data stream from Echelon 0
- Echelon 2 sites fully participate in use cases and accelerate them:
- EIC International Computing Organization (EICO) formed

Design and implementation

Testbed schematic

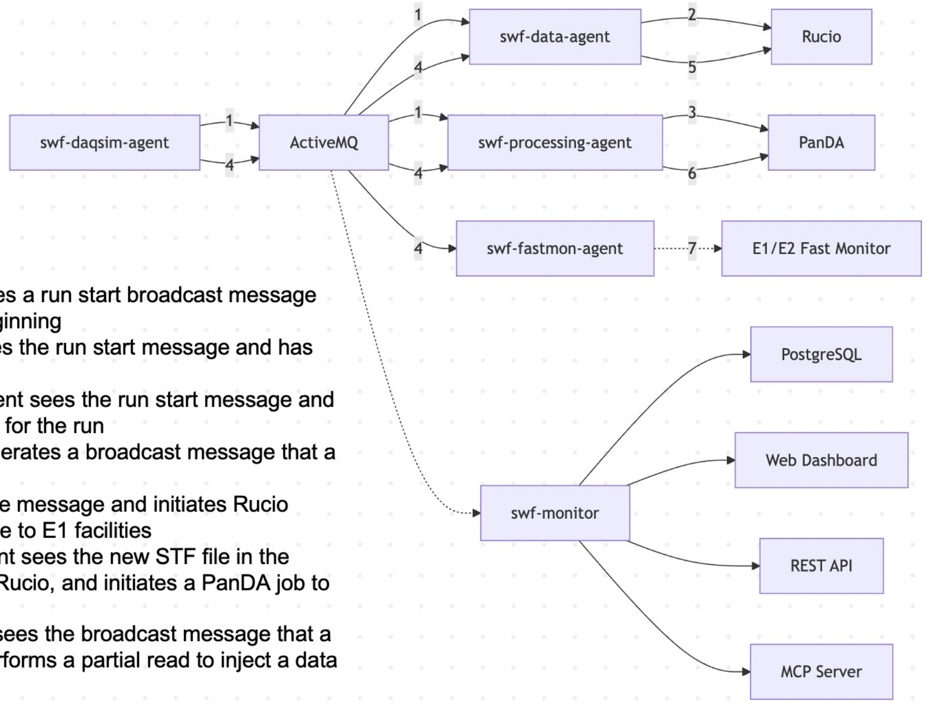

1. Run Start - daqsim-agent generates a run start broadcast message indicating a new datataking run is beginning
2. Dataset Creation - data-agent sees the run start message and has Rucio create a dataset for the run
3. Processing Task - processing-agent sees the run start message and establishes a PanDA processing task for the run
4. STF Available - daqsim-agent generates a broadcast message that a new STF data file is available
5. STF Transfer - data-agent sees the message and initiates Rucio registration and transfer of the STF file to E1 facilities
6. STF Processing - processing-agent sees the new STF file in the dataset and transferred to the E1 by Rucio, and initiates a PanDA job to process the STF
7. Fast Monitoring - fastmon-agent sees the broadcast message that a new STF data file is available and performs a partial read to inject a data sample into E1/E2 fast monitoring

Monitor and database

- Django based testbed monitor manages a postgres database
- All system information aggregates there
- Monitor provides an authenticated REST interface to exchange data with agents
- Monitor listens in on system activit and presents it
  - ActiveMQ messaging
  - Python logging
- Monitor will provide MCP service t provide system info to LLMs, and later take direction
  - In progress
  - MCP services also being implemented in core PanDA, we will integrate them

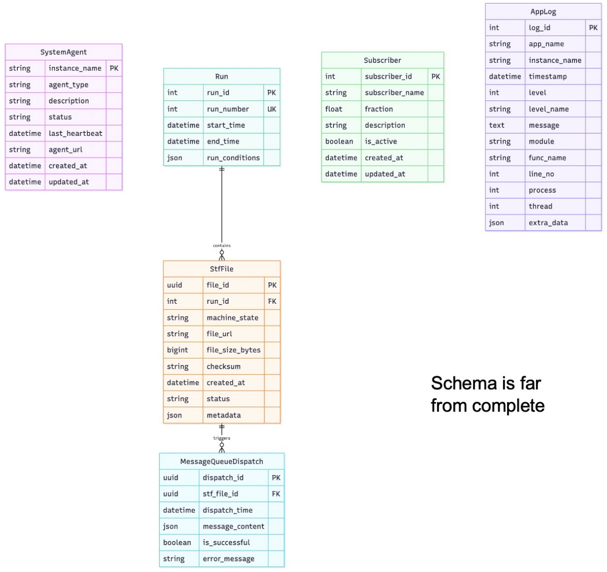

Schema is far from complete

State machine first version

- We came up with a ‘minimal’ state machine to describe the basics of the system

States
- no_beam
  - Collider not operating
- beam
  - Collider operating
- run
  - Physics running
- calib
  - Dedicated calibration period
- test
  - Testing, debugging
  - Any substates can be present during test

Substates
- not_ready
  - detector not ready for physics datataking
  - occurs during states: no_beam, beam, calib
- ready
  - collider and detector ready for physics, but not declared as good for physics
  - when declared good for physics, transitions from beam/ready to run/physics
  - occurs during states: beam
- physics
  - collider and detector declared good for physics
  - if collider or detector drop out of good for physics, state transitions out of ‘run’ to ‘beam’ or ‘off’
  - occurs during states: run
- standby
  - collider and detector still good for physics, but standing by, not physics datataking (dead time!)
  - occurs during states: run
- lumi
  - detector, machine data that is input to luminosity calculations
  - occurs during states: beam, run
- eic
  - machine data, machine configuration
  - occurs during states: all
- epic
  - detector configuration, data
  - occurs during states: all
- daq
  - info, config transmitted from DAQ
  - occurs during states: all
- calib
  - a catch-all for a great many calib data types, we can start small
  - occurs during states: all (assuming there are cases where calib data is taken during beam on)

- Code is in GitHub in the BNL NPPS group organization
- Needed service infrastructure has been established at BNL (PanDA, Rucio, Postgres, ActiveMQ etc)
- **We are presently working on gathering the pieces together on the testbed server and starting integration of agents and services into an operating whole.**

Next steps for Streaming Orchestration

- **Near-term plans**
  - Integrate the components and operate them as a system on our dev/deploy platform
  - Complete the visibility of the full system in the monitor
  - Start creating and exercising processing scenarios starting from basic prompt processing
  - Introduce concurrent STF workflows for bulk data processing and TF workflows for fast monitoring
  - Address more complex workflows, like calibration,
  - Address fine grained streaming processing below the STF level
  - Extend the testbed… (E1-E2 connection to Manitoba, Japan)

Streaming Reconstruction

- **Timeframe splitter**

Streaming Reconstruction

- **time-slice size dependence**

- All time slice width could capture over 90% physics events.
- The background reduction rates for all setting are not significantly changed.

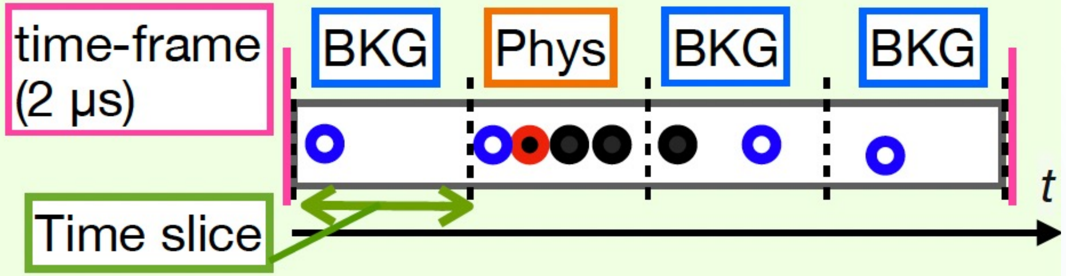

※1 This simulation uses a 500 kHz physics rate; however, in realistic simulation (83 kHz), it is expected that the number of backgrounds increase fivefold.
→ **Next: Need more strong constrain.**

※2 This test does not consider the track reconstruction quality.
→ **Next: Evaluate how the time window affects the track reconstruction quality.**

trigger/time slice width [ns]	5	10	20	200
Physics trigger rate (trigger/all events)	1389/1530 = 90.8%	1441/1530 = 94.2%	1473/1530 = 96.3%	1512/1530 = 98.8%
BKG trigger rate (trigger/all events)	1700/612000 = 0.3%	1100/306000 = 0.4%	674/153000 = 0.4%	173/15300 = 1.1%

Use case : Backward ECal

Carlos Muñoz Camacho

Calibration with π^0

Methods successfully used with EMCals at JLab
Based on NIM A566 (2006) 366

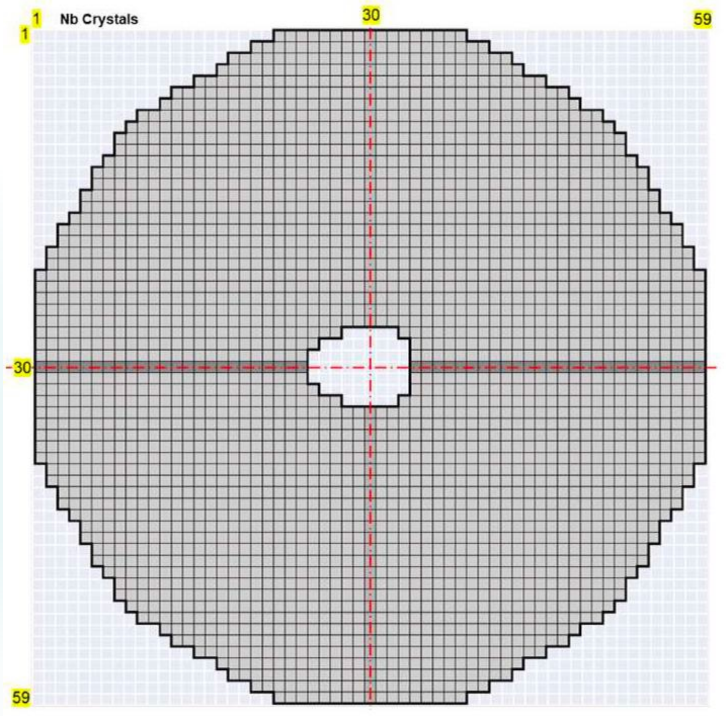

PWO crystals (2x2 cm^2) with SiPM
(16 3x3 mm^{2}or 4 6x6 mm^2 per crystal)

IJCLab Irène Joliot-Curie Laboratoire de Physique des 2 Infinis

Backward ECal calibration with π^0s

SIDIS simulation at 18x275 GeV

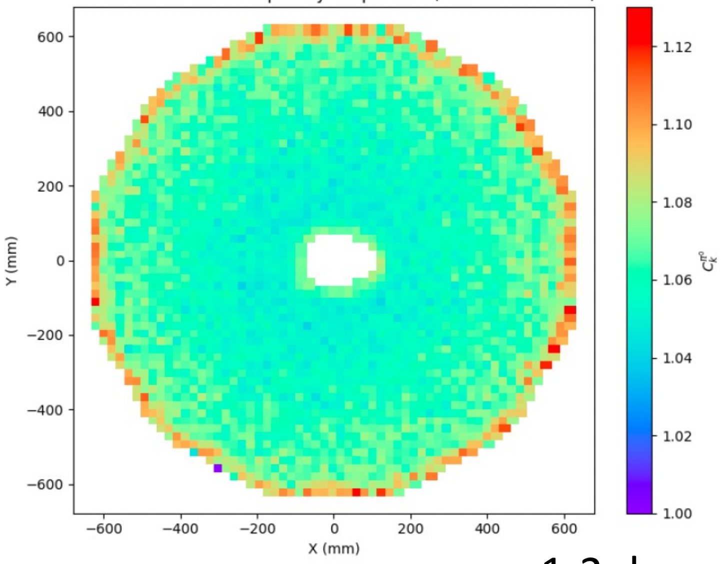

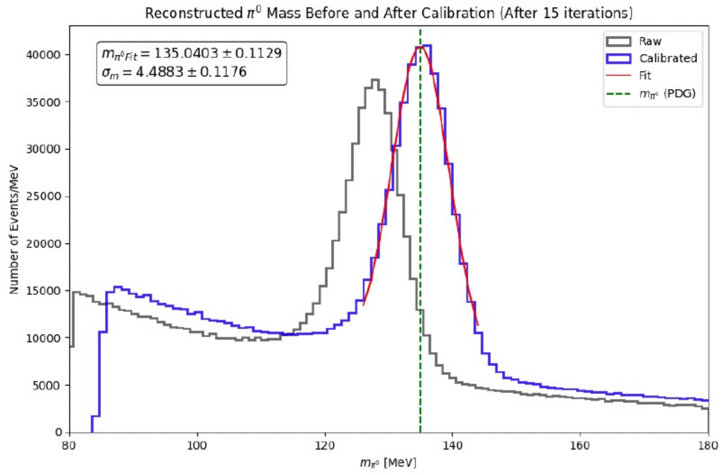

1-2 days of data should be sufficient

Analysis by Axel Perez Ruiz (IJCLab)

Use case : Backward ECal

ALLEN4EIC:
A Real-Time Data Processing Framework for the ePIC Detector

Letter of Intent to the
Hadron Physics call in Horizon Europe

Vladimir Gligorov, LPNHE
Carlos Muñoz Camacho, IJCLab

Nantes, Jul 1-2 (2025)

ALLEN 4EIC

Allen Framework: Proven at the LHCb Experiment

- Developed for GPU-based high-level triggering at LHCb.
- Processes 30 million events/sec using commodity GPUs.
- Cross-platform (x86, CUDA, ARM); Python + C++/CUDA.
- Includes monitoring, high energy-efficiency.

Allen replaces hardware triggers using GPUs and high-performance software. It's portable, open-source, and energy-efficient. Perfect candidate for EIC-style real-time computing.

We will investigate more on how we can advance Allen framework for ePIC:
Adapt ePIC geometry, data format, streaming architecture, routines as in EICrecon