ePIC Software & Computing Meeting

ePIC Software & Computing Weeks

Regular ~~workshops~~ meetings to drive forward priority targets and provide an avenue for new collaboration members to actively engage. In-person component key but challenge.

12 pages of detailed notes
that enabled software progress, pushed the review preparations, and informed our planning.

Topics: Status and plans; software and simulations for TDR, tutorials; streaming computing; software projects with HEP.

Our Objectives: Monday

ePIC Software & Computing Update: SCC	40/S2-D01 - Salle Dirac, CERN	09:00 - 10:30
Break	CERN	10:30 - 11:00
ePIC Software & Computing Update: WGs	40/S2-D01 - Salle Dirac, CERN	11:00 - 12:30
Lunch Break	CERN	12:30 - 14:00
ePIC Software & Computing Update: WGs	40/S2-D01 - Salle Dirac, CERN	14:00 - 15:30
Break	CERN	15:30 - 16:00
ePIC Software & Computing Update: WGs	40/S2-D01 - Salle Dirac, CERN	16:00 - 17:30

Reviewing Status and Future Plans:

- May 1 marks the first anniversary of the convener roles being endorsed by the Collaboration Council.
- Over the past year, we have addressed many time-critical items.
- Now, as we approach the completion of the first year, it is crucial to consider the broader perspective.
- Therefore, we will dedicate the first day of the meeting to reviewing the status and future plans of each of the WGs.

Our Objectives: Tuesday

TDR Progress: Status and DAP	*40/S2-C01 - Salle Curie, CERN*	09:00 - 10:30
Break	*CERN*	10:30 - 11:00
TDR Progress: Data Model for TDR: Examples for Usage	*40/S2-C01 - Salle Curie, CERN*	11:00 - 12:30
Lunch Break	*CERN*	12:30 - 14:00
TDR Progress: Streaming Reconstruction Prototype	*2/R-030, CERN*	14:00 - 15:30
Break	*CERN*	15:30 - 16:00
TDR Progress: Calorimeter Reconstruction	*2/R-030, CERN*	16:00 - 17:00
TDR Progress: Topical ACTS Meeting on ePIC Tracking *Alexander J Pfleger, Andreas Salzburger*	*CERN*	17:00 - 18:00

TDR:

- We will review the progress of the software and simulation readiness for the TDR and work on critical tasks.
- **We will discuss**:
  - Current status.
  - Improve documentation of progress and completed tasks.
  - Finalize data model review.
  - Prototype for event reconstruction from timeframes.
  - Calorimeter reconstruction.
  - Track reconstruction, jointly as part of ACTS developer meeting.
- Priorities for discussion identified in April 10 meeting.

Our Objectives: Wednesday

Session	Speaker	Location	Time
Tutorial: Introduction to ePIC Software: Getting Started with Detector and Physics Studies for the TDR	*Holly Szumila-Vance*	*40/S2-D01 - Salle Dirac, CERN*	09:00 - 10:00
Tutorial: Working with Simulation Output	*Dr Stephen Kay*	*40/S2-D01 - Salle Dirac, CERN*	10:00 - 12:00
Lunch Break		*CERN*	12:00 - 13:30
Tutorial: Simulating Detectors and Their Readout	*Simon Gardner*	*40/S2-D01 - Salle Dirac, CERN*	13:30 - 15:30
Tutorial: Developing Reconstruction Algorithms	*Nathan Brei*	*40/S2-D01 - Salle Dirac, CERN*	15:30 - 17:30

Software Tutorials:

- The introductory session will guide collaboration members through the software aspects you need to familiarize yourself with to take the initial steps, whether they are part of a DSC, a PWG, or an SWG.
- The next three sessions will cover developing analyses based on simulation outputs and getting started with detector simulations as well as reconstruction algorithms.

- The tutorials will blend instructional materials with bite-sized tasks to apply the learned material.

- Scheduled in the CEST time zone, the tutorial is ideally suited for participants in Africa, Asia, and Europe.
- We need a lot of interest from Africa, Asia, and Europe.

Our Objectives: Thursday

Session	Location	Time
Streaming Computing: Implementations	40/S2-D01 - Salle Dirac, CERN	09:00 - 10:30
Break	CERN	10:30 - 11:00
Streaming Computing: Distributed Computing Model	40/S2-D01 - Salle Dirac, CERN	11:00 - 12:30
Lunch Break	CERN	12:30 - 14:00
Streaming Computing: EICrecon Roadmap	40/S2-D01 - Salle Dirac, CERN	14:00 - 15:30
Break	CERN	15:30 - 16:00
Streaming Computing: Metadata	CERN	16:00 - 17:30

Streaming Computing:

- **We will discuss**:
  - (Near) Real-time processing for streaming data. Discussion with Allen.
  - **Distributed computing**: Scientific data and workflow and workload management. Reports from Rucio and PanDA teams.
  - **EICrecon Roadmap**: How should EICrecon evolve? What are the incremental steps needed to achieve this evolution?
  - **Related to that**: Configuration management, calibrations and conditions databases.

Our Objectives: Friday

Session	Location	Time
Common Projects: HEP Software Foundation, Stacks	40/S2-D01 - Salle Dirac, CERN	09:00 - 10:30
Break	CERN	10:30 - 11:00
AI-Empowered Research	40/S2-D01 - Salle Dirac, CERN	11:00 - 12:30
Lunch Break	CERN	12:30 - 14:00
Wrap Up	CERN	14:00 - 15:30

Collaboration, AI, Wrap Up:

- Common scientific projects and how to enhance collaboration. Discussion with HSF and Stacks / key4hep.
- ePIC AI Strategy. Input from CERN EP-SFT.
- **Wrap Up**: Ensure that we depart from CERN with a clear understanding of our next steps and schedule.

What Else?

- **Reminder**: This is our meeting. We can and likely will revise the agenda to prioritize progress over adherence to a strict schedule.

- Please take pictures for reports and presentations.

- Additionally, as requested by Shujie, let us explore options for a tour.

Software for ePIC

After a presentation on “Breakthroughs in Detector Technology”, Ian Shipsey (Oxford) was asked about the role of software.

"***Software is the soul of the detector,***” Ian Shipsey replied in a poetic way and emphasized the **importance of great software for great science**. He added that we need to **work together**, on a global scale and with other fields, to achieve this goal.

Our Philosophy

- We focus on **modern scientific software & computing practices** to ensure the **long-term success of the EIC scientific program** throughout all CD milestones.
  - Strong emphasis on modular, orthogonal tools.
  - Integration with HTC/HPC, CI workflows, and enable use of standard data science toolkits.
- We **leverage cutting edge sustainable community software** where appropriate, **avoiding the "not invented here" syndrome**.
  - Can build our software on top of a mature, well-supported, and actively developed software stack by using modern community tools, e.g. from CERN, the HPC community, and the data science community.
  - Actively collaborate with external software projects, while externalizing some support burden to external projects.
- We embrace these practices today to avoid starting our journey to EIC with technical debt.
- **We are writing software for the future, not the lowest common denominator of the past!**

ePIC Software Stack: A Modular Simulation, Reconstruction, and Analysis Toolkit

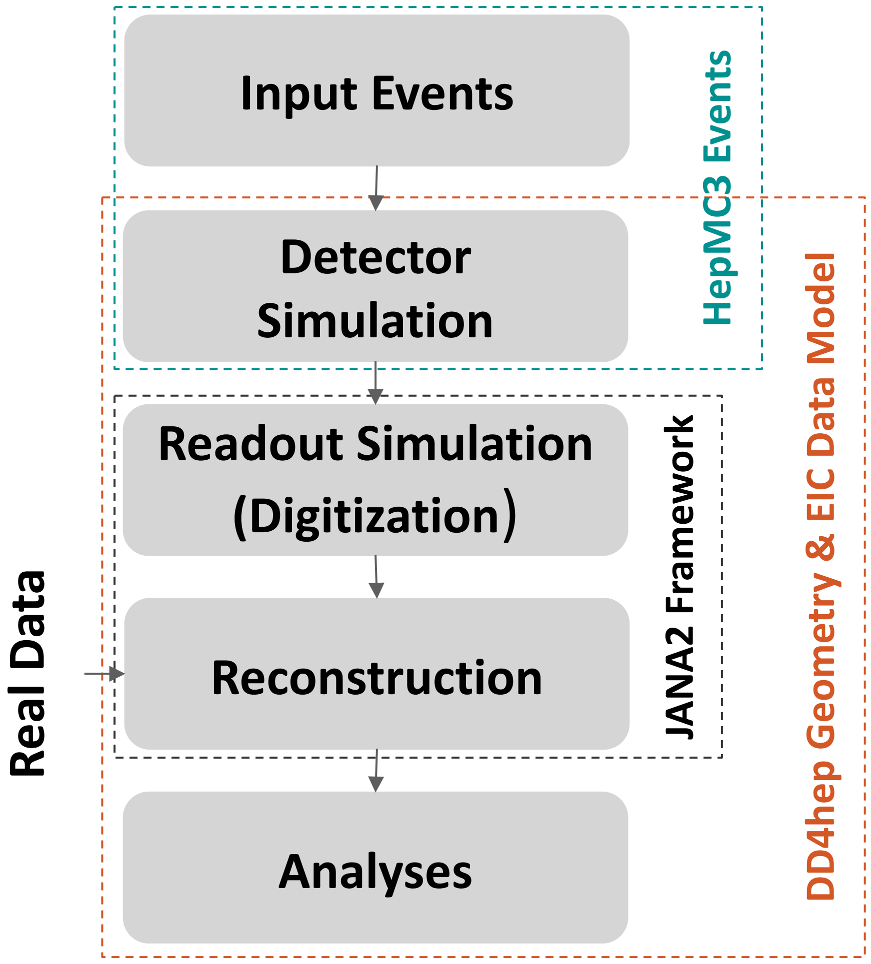

Input events from **MC event generators** or particle guns, with optional physics background merging.

Geant4 simulations with **DD4hep** for geometry description and exchange, output data in the **EIC Data Model** (EDM4hep + EDM4eic, described in Podio).

Algorithms to transform the GEANT4 hits to mimic real detector readout, including background stacking, “pileup”, DAQ frames

Realistic reconstruction algorithms starting from raw detector output (from digitization or real data).

User analyses in plain C++/ROOT or Python/uproot, facilitated by using a flat data model.

Continuous integration for detector and physics benchmarks and monthly production campaigns ensure a production-ready software stack at any time.

Open, Collaborative Development

Enable Access Without Restrictions

- ePIC collaboration members include over 170 institutions worldwide.
- **Data publicly available** at host labs.
- **Simple, flat data structures** (i.e. could be a csv), stored as ROOT files:
    - Straightforward analysis in ROOT or Python without the need for data structure libraries.
    - Encourage collaboration with computer, data, and other scientists outside of NP and HEP.
- **Public software repositories** on GitHub.

Encourage Upstream Contributions

- Requirements of well-formed HepMC as input has resulted in real improvements to multiple MCEGs used by EIC community.
- Various upstream contributions to software packages, e.g., ACTS, DD4hep, Spack, uproot.

Encourage Social Coding

- CI platform provides the incentive for developers to commit code frequently: achieving data management and analysis preservation goals.
- Pull request reviews to ensure higher quality code and build developer skills.

Cultural change for NP: How to see reviews as asset?

An Important Snapshot: The ePIC Software & Computing Review in 2023

- Convened by and reporting to the host labs: Haiyan Gao (BNL) and David Dean (Jefferson Lab).
- An excellent **review panel**, which will continue as the **EIC Computing and Software Advisory Panel**:
  - Frank Wuerthwein (chair, UCSD), Mohammad Al-Turany (GSI), David Brown (LBNL), Simone Campana (CERN), Pere Mato (CERN), Christoph Paus (MIT), Heidi Schellman (OSU).
- We addressed five charge questions posed in the context of being 10 years away from data:
  - Briefly summarized below with the panel's conclusions; see Indico and the closeout for the full story.

- **Is there a comprehensive and cost effective long term S&C plan?**
  - *Yes. Impressive organization and plan. ePIC should verify software and simulation readiness for the TDR by May 2024. A long term computing needs assessment was not presented; present one in a year.*
- **Are there adequate plans for integrating international partners?**
  - *Yes. The opportunities are significant (Canada, Italy, UK thus far), the collaboration is doing all the right things to leverage them.*
- **Are S&C plans integrated with HEP/NP community developments?**
  - *Yes. ePIC is leveraging tools and services widely adopted and supported in the HEP/NP community. ePIC should contribute to supporting key software. Document dependencies and ePIC contributions.*
- **Are there sufficient S&C resources to deliver the TDR?**
  - *Yes. Computing resources available from OSG (80%) and labs (20%) look sufficient. Software development plan is credible to meet the critical TDR milestones.*
- **Do BNL/JLab joint institute plans integrate sufficiently with experiment S&C?**
  - *Yes. Organization resembles those of the LHC which have worked well. The sharing of responsibilities between experiment and institute, particularly for software, still being discussed. Ongoing evaluation will be part of future reviews.*

Priorities

TDR: Software and simulations for detector and physics studies. Prototype for reconstruction of physics events from streaming data. Monday, Tuesday and Friday sessions.

Onboarding and engagement. Tutorials on Wednesday.

Computing Model. Thursday.

ePIC Software & Computing Organization

Spokesperson's Office

Detector | **Software & Computing** | Physics

Guiding Principles:

- *Diversity, Equity, and Inclusion*
- *Statement of Software Principles*
- *Sustainability.*

Software & Computing Coordinator
Markus Diefenthaler (Jefferson Lab)

Cross-cutting Working Group:

- *Data and Analysis Preservation*

Deputy Coordinator (Operations)
Wouter Deconinck (U. Manitoba)

Operation Working Groups:

- Production
- User Learning
- Validation

Deputy Coordinator (Development)
Sylvester Joosten (ANL)

Development Working Groups:

- Physics and Detector Simulation
- Reconstruction
- *Analysis Tools*

Deputy Coordinator (Infrastructure)
Torre Wenous (BNL)

Infrastructure Working Groups:

- Streaming Computing Model
- *Multi-Architecture Computing*
- *Distributed Computing*

Changes to Organization

- **Deputy Software & Computing Coordinator for Development**:
  Sylvester → Dmitry (Kalinkin)

- **Validation WG**: Successor for Dmitry?
- **User Learning WG**: Successor for Kolja?

- **Analysis WG**: Kick off?

Discuss changes over course of meeting.

Onboarding

Landing Page https://eic.github.io/documentation/landingpage.html

Get started

ePIC Tutorials

HEP Software Training Center

FAQ

Welcome to the **ePIC Landing Page**!

Our mailing list: eic-projdet-compsw-l@lists.bnl.gov

Subscribe here: https://lists.bnl.gov/mailman/listinfo/eic-projdet-compsw-l

Discuss plan for improving and enhancing the landing page, essential for onboarding new users.

Compute-Detector Integration to Accelerate Science

- **Problem** Data for physics analyses and the resulting publications available after O(1year) due to complexity of NP experiments (and their organization).
  - Alignment and calibration of detector as well as reconstruction and validation of events time-consuming.
- **Goal** **Rapid turnaround of 2-3 weeks for data for physics analyses**.
  - Timeline driven by alignment and calibrations.
- **Solution** Compute-detector integration using:

Streaming readout for continuous data flow of the full detector information.

AI for autonomous alignment and calibration as well as reconstruction and validation for rapid processing.

Heterogeneous computing for acceleration.

Alignment and Calibration of the ePIC Detector

[Spreadsheet]

			Steady State calibrations: aim to produce final reconstruction-ready calibration within few days of physics data taking in a continous process														Post-reconstruction calibrations (applied at analysis stages)		
Subsystem	Region	Pre-physics-operation calibrations (Cosmic, no-beam calibration, commissioning)	Task	Human intervention?	Data Needed	Dependecy	T0 + 12	T0 + 24	T0 + 36	T0 + 48	T0 + 60	T0 + 72	T0 + 84	T0 + 96	Monitoring	Computing resource		Comment	Subsystem
MAPS	Barrel+Disk	Threshold Scan Fake rate scan/noisy pixel masking	(See Alignment)															SRO meeting: https://indico.bnl.gov/even Threshold Scan require substaintial DAC	MAPS
MPGD	Barrel+Disk	?	?															TIC meeting: https://indico.bnl.gov/event	MPGD
bTOF, eTOF (ac-	Barrel/Forw	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 Ac Depend	Depend	Process	Processing								SRO meeting https://indico.bnl.gov/even	bTOF, eTOF (ac-lgad)
Central Detector Tracker Ali		Initial alignment	Alignment Check/Update (if needed)	QA	Prodcution data		Processing											SRO meeting https://indico.bnl.gov/even	Central Detector Tracker
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 Ac	Processing										TIC meeting: https://indico.bnl.gov/event	pfRICH
DIRC	Barrel	Laser data?	?	?														TIC meeting: https://indico.bnl.gov/event	DIRC
dRICH	Forward	Bunch timing offset scan Threshold scan Noise masking	Track based alignment	?	High p tracks ~1hr of of production data?	Tracking	Data Ac Depend	Process	Processing									SRO meeting: https://indico.bnl.gov/ever	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 Ac Depend	Data Ac	Process	Processing					LED			SRO meeting: https://indico.bnl.gov/ever Carlos: aiming 1% precision Planning for LED flash during productior	bEMC
AstroPix	Barrel																	TIC meeting: https://indico.bnl.gov/event	AstroPix
ScifiPb	Barrel		SiPM gain		?													TIC meeting: https://indico.bnl.gov/event	ScifiPb
fEMC	Forward	IV Scan	Pi0, eta->gg events energy scale Second iteration pi0 (if needed)	QA	Pi0 di-photon resonance ~1 day of production data		Data Ac	Data Ac	Process	Processing	Processing				LED		High energy cluster non-linearity	SRO meeting: https://indico.bnl.gov/ever Need pi0 filtered data for automated cali AI driven calibration?	fEMC
bHCAL	Backward	LED	?															TIC meeting: https://indico.bnl.gov/event	bHCAL
cHCAL	Barrel	MIP calibration Gain calibration	(See hadronic e-scale calib)															SRO meeting: https://indico.bnl.gov/ever	cHCAL
fHCAL	Forward																		fHCAL
fHCAL insert	Forward																		fHCAL insert
Hadronic energy scale calib		?	Set full calo stack energy scale for hadroinc shower and jets	?	High energy hadronic showers and jets	Tracking h-PID	Data Ac Depend	Data Ac Depend	Data Ac Depend	?	?	?	?	?			Final energy scale calibration (if needed)	Comments from Oleg during SRO meeti	Hadronic energy scale c
low Q2 Tagger	Far Backwa	Alignment?																TIC meeting: https://indico.bnl.gov/event	low Q2 Tagger
low Q2 Tagger ((	Far Backwa																	TIC meeting: https://indico.bnl.gov/event	low Q2 Tagger (CAL)
Pair Spec Track	Far Backwa																	TIC meeting: https://indico.bnl.gov/event	Pair Spec Tracker
Par Spec Cal	Far Backwa																	TIC meeting: https://indico.bnl.gov/event	Par Spec Cal
Direct Photon C	Far Backwa																	TIC meeting: https://indico.bnl.gov/event	Direct Photon Cal
B0 Tracking	Far Forwar	Survey alignment/Cosmic	Alignment check		MIP		Processing											SRO/FF meeting https://indico.bnl.gov/e	B0 Tracking
B0 PbWO4	Far Forwar	Survey alignment/Cosmic	SiPM gain		MIP/Gamma/Electrons		Processing								LED			SRO/FF meeting https://indico.bnl.gov/e	B0 PbWO4
Roman (Pots)	Far Forwar					Acc. BPM Potential use of vertex of central detector	Data Ac Depend	Processing										SRO/FF meeting https://indico.bnl.gov/e	Roman (Pots)
Off Momentum	Far Forwar	laser/survey alignment Low lumi running	beam position monitors/fill by fill correction		MIP rate distribution in RP		Data Ac Depend	Processing										SRO/FF meeting https://indico.bnl.gov/e	Off Momentum
ZDC PbWO4	Far Forwar	Survey alignment, timing delay	SiPM/APD gain, timing	QA	Photon		Processing								LED			SRO/FF meeting https://indico.bnl.gov/e	ZDC PbWO4
ZDC Sampling	Far Forwar	Survey alignment, timing delay	SiPM gain	QA	Single neutron		Processing								LED			SRO/FF meeting https://indico.bnl.gov/e	ZDC Sampling

Discussion on April 16 on automated calibrations.

Current View of the ePIC Computing Model

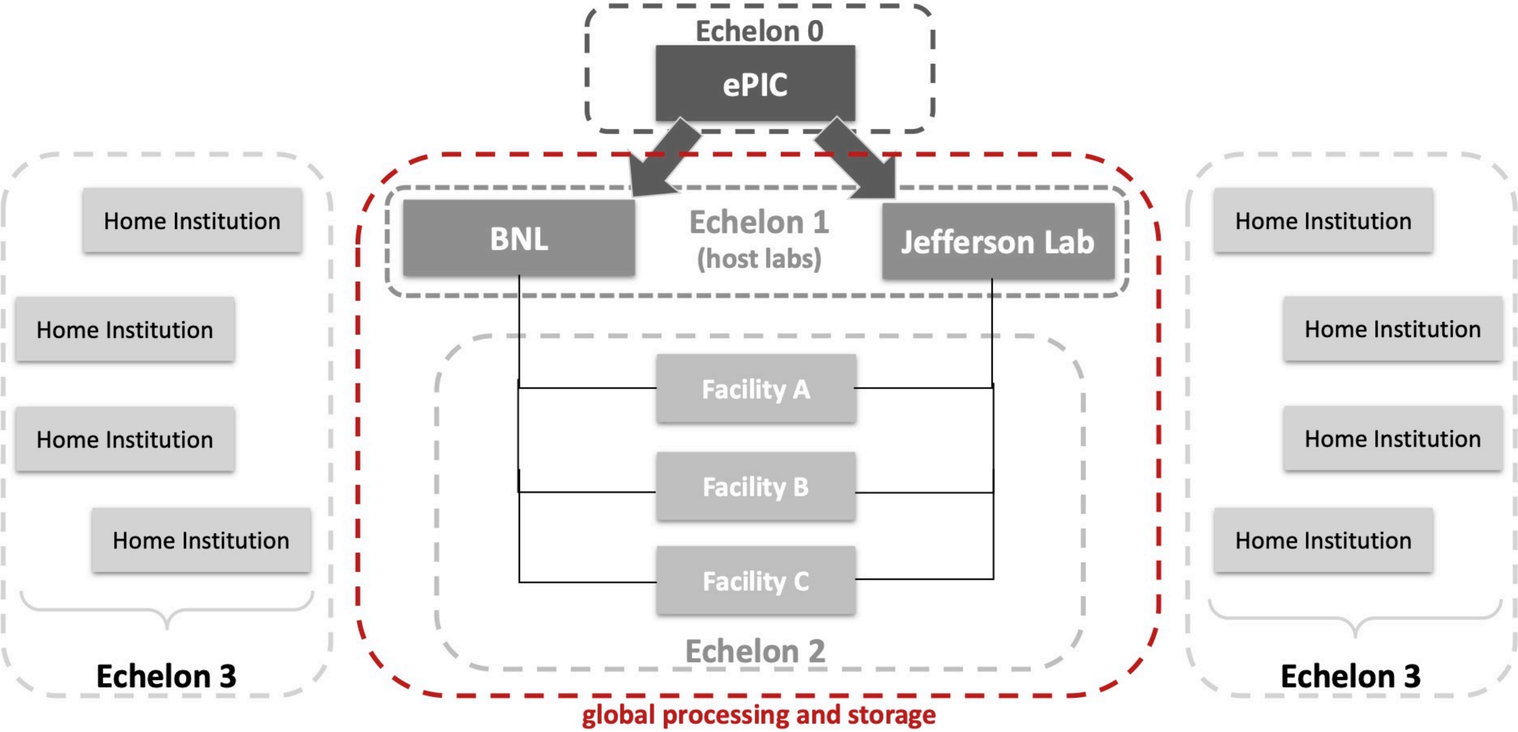

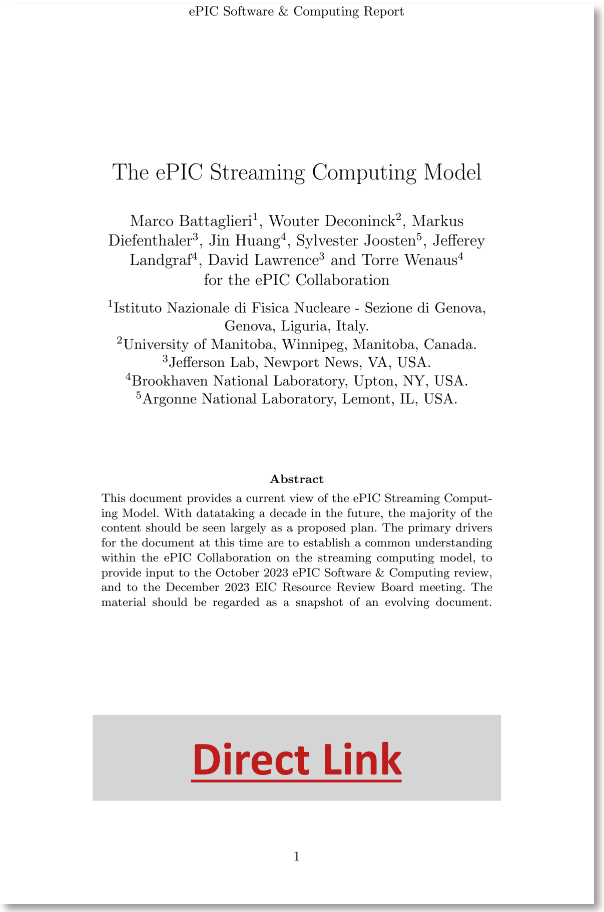
ePIC Software & Computing Report

The ePIC Streaming Computing Model

Marco Battaglieri[1], Wouter Deconinck[2], Markus Diefenthaler[3], Jin Huang[4], Sylvester Joosten[5], Jefferey Landgraf[4], David Lawrence[3] and Torre Wenaus[4] for the ePIC Collaboration

[1]Istituto Nazionale di Fisica Nucleare - Sezione di Genova, Genova, Liguria, Italy.
[2]University of Manitoba, Winnipeg, Manitoba, Canada.
[3]Jefferson Lab, Newport News, VA, USA.
[4]Brookhaven National Laboratory, Upton, NY, USA.
[5]Argonne National Laboratory, Lemont, IL, USA.

Abstract

This document provides a current view of the ePIC Streaming Computing Model. With datataking a decade in the future, the majority of the content should be seen largely as a proposed plan. The primary drivers for the document at this time are to establish a common understanding within the ePIC Collaboration on the streaming computing model, to provide input to the October 2023 ePIC Software & Computing review, and to the December 2023 EIC Resource Review Board meeting. The material should be regarded as a snapshot of an evolving document.

1

Direct Link

Initial version of a plan set to develop over the next decade.

Priority Target for the TDR

Prototype of event reconstruction from realistic frames:

Purpose: Demonstrate that we can reconstruct events from Streaming DAQ.

Purpose: Estimation of streaming reconstruction time for compute resource planning.

Prototype of Event Reconstruction from Realistic Frames

Key Tasks: We limit the scope of the first study to the track reconstruction only. The key is to demonstrate we can correlate hits in a realistic time frame to the various events in the time window of the MAPS.

- **Electronics and DAQ WG**
  - Define the *realistic* time frames. We need a consensus on the meaning of *realistic*.
- **Simulation WG**:
  - Prepare simulation productions, using detailed information on FEEs for tracking detectors, utilizing the full, wide MAPS integration window for tracking purposes.
  - Implement and utilize the frame-building infrastructure post-Geant4 and post-digitization.
- **Reconstruction WG**:
  - Integrate Jana2's built-in workflow for supporting frames in and events out in EICrecon.
  - Adapt the reconstruction process to work with frames, making it frames-aware.
  - Demonstrate tracking from realistic frames.

Key Task: Realistic Time Frames

- **Electronics and DAQ WG**
  - Define the *realistic* time frames. We need a consensus on the meaning of *realistic*. **Meeting on March 21.**

- Timing distributions in simulations realistic. Very long tails due to slow neutrons.
- **Time frames definition**:
  - DIS, Proton Beam Gas, Electron Beam Gas (eventual Synchrotron radiation) should be randomly distributed according to desired beam energies and intensities.
  - Hits from simulated physics collisions events should be distributed in time according to:
    - T_{hit} = nominal event T_0 + $_{\text{Thit\_sim}}$
  - The hits for the MAPS detectors need two special characteristics:
    - The time should be set to the beginning of the 2us period it occurs within.
    - The hit should be replicated in either the prior or following 2us period depending on whether it is closer to the end of the hit period or closer to the beginning.
  - Add noise only to the MAPS and dRICH. Assume that reconstruction, e.g., cluster finding, in other detectors will suppress nearly all of the noise.

Key Task: Time Frames in Simulations

- **Simulation WG**:
  - Prepare simulation productions, using detailed information on FEEs for tracking detectors, utilizing the full, wide MAPS integration window for tracking purposes.
  - Implement and utilize the frame-building infrastructure post-Geant4 and post-digitization.

Another Meeting on March 21.

- **Time frames**:
  - First step in digitization is to allow multiple hits per cellID, with appropriate time buckets (work in progress).
  - **After the prototype phase**: Design a flexible method for merging signal and background events post-Geant4 simulation and post-digitization, not at the physics collision level.

Key Task: Reconstruct Events from Time Frames

- **Reconstruction WG**:
  - Integrate Jana2's built-in workflow for supporting frames in and events out in EICrecon.
  - Adapt the reconstruction process to work with frames, making it frames-aware.
  - Demonstrate tracking from realistic frames.

Meetings on March 26 and April 2.

Timeslices topology

Nathan Brei (CST) presented major update to JANA2 for processing of timeframes, events, and subevents using JOmniFactory. The update is available for tests and needs to be validated.

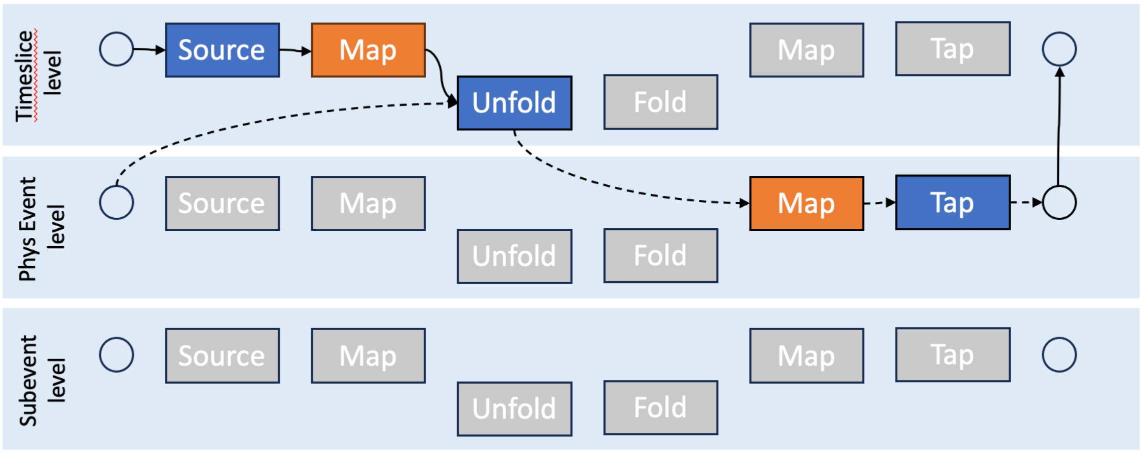

- Source
  - JEventSource::GetEvent()
- Map
  - JOmniFactory::Process()
  - JEventProcessor::Preprocess()
  - JEventSource::Preprocess()
  - JEventUnfolder::Preprocess()
  - JEventFolder::Preprocess()
- Tap
  - JEventProcessor::Process()
- Unfold
  - JEventUnfolder::Unfold()
- Fold
  - JEventFolder::Fold()

Timeslice
Event
Subevent

Parallel
Sequential

Triggered discussion on meaning of runs and events in streaming context on April 9.

Current Estimate for Compute Resources

Streaming DAQ sends data in **1ms time frames**.

Each time frame corresponds to 10MB of data.

Based on our current detector readout design and when running at peak luminosity and in standard operating conditions.

40% of data bunch crossing related, 60% background.

In a year, we will record 15.5 billion frames.

Assuming a 50% up-time for 6 months.

Number of expected events (assuming a 50% up-time for 6 months):

- The event rate at peak luminosity is 500kHz, which gives roughly 4 x 10^{12} events.
    - Lower at start of operations, where the luminosity will be lower (but relatively speaking background rate is expected to be higher).
- The expected number of physics events of interest for one year of running at peak luminosity is ~ 10^{10}.
    - The actual physics events is only a very small fraction of the total physics bunch crossings.

Number of simulation events:

- We expect to simulate 10x events for each event of interest, yielding O(10^{11}) simulated events.

While considerable (~ 60k core years on today's hardware), this should be a realistic target in a decade.

Core-seconds for simulation and reconstruction (on a typical modern machine):

- Our current simulations including background take ~17s for simulation and ~ 2-3s for reconstruction, per event.
    - Simulation and reconstruction on event level only.
- **Unknown**: How much this will change once changing to streaming data processing?
- **Priority target for TDR**: Prototype of event reconstruction from realistic frames.

We need reliable estimate for our compute resources by the next review in October.

Milestones

Milestones for NOW

- **Prototype of event reconstruction from realistic frames**:
  - Demonstrate that we can reconstruct events from Streaming DAQ.
  - Estimation of streaming reconstruction time for compute resource planning.

Milestone for review

- Quantitative computing model.
- Related to that: Publication on computing model.
- Work with international partners to integrate computing resources.

Milestones During Detector Construction Phase

Develop a plan with the Electronics and DAQ WG to align on shared priorities.

- Provisioning DAQ and software sufficient for test beams, which can serve as small scale real-world testbeds for the developing DAQ and software.
- Streaming challenges exercising the streaming workflows from DAQ through offline reconstruction, and the Echelon 0 and Echelon 1 computing and connectivity.
- Data challenges exercising scaling and capability tests as distributed ePIC computing resources at substantial scale reach the floor, including exercising the functional roles of the Echelon tiers, particularly Echelon 2, the globally distributed resources essential to meeting ePIC's computing requirements.
- Analysis challenges exercising autonomous alignment and calibrations.
- Analysis challenges exercising end-to-end workflows from (simulated) raw data to exercising the analysis model.

ePIC Software & Computing at CERN

I am looking forward to exciting discussions, developments, and tutorials this week.

We will conclude the meeting with a list of concrete action items.

Additionally, we will create a meeting schedule for the next two months or possibly even beyond.

Thank you so much to Yasemin Altinbilek and Pere Mato for making the meeting at CERN possible.

DALL-E, could you please draw a picture of exciting discussions at CERN?