ePIC Software & Computing Priorities

Markus Diefenthaler (Jefferson Lab)

Common Themes From Priority Discussion

There is a need for better communication within the collaboration, across multiple different entities

Communication strategy:

1) **Unified Information Hub**: Ensure there is a single place that everyone should know to go to for software and computing information, and that place should provide a resolution to any request (including a way to request information to be added or more easily findable).
2) **Inclusive Meeting Participation**: Ensure that asynchronous participation in all WG meetings is possible, at little to no disadvantage over synchronous participation, to include the ability to formulate questions on meeting points in advance and the ability to receive responses afterwards.
3) **Effective Collaboration Communication**: Maintain regular, archived, and searchable communication to the collaboration, in a predictable format and at predictable intervals, in an effort to keep everyone informed of upcoming events and changes through a single channel.

Implementation of the communication strategy will be discussed in coming weeks.

Common Themes From Priority Discussion

We require additional engagement to sustain ongoing efforts—and even more to expand them.

Software & Computing Perspective: To ensure a sustainable effort, broader involvement is necessary. Within Software & Computing, we require additional engagement to sustain ongoing efforts—and even more to expand them.

Membership Survey

Promising: There is interest in ePIC Software & Computing. Reach out to use to get actively involved.

Concern: ePIC will be a compute-intensive experiment at the scale of ATLAS/CMS today. Do we have enough FTE to ensure we will not be compute-limited?

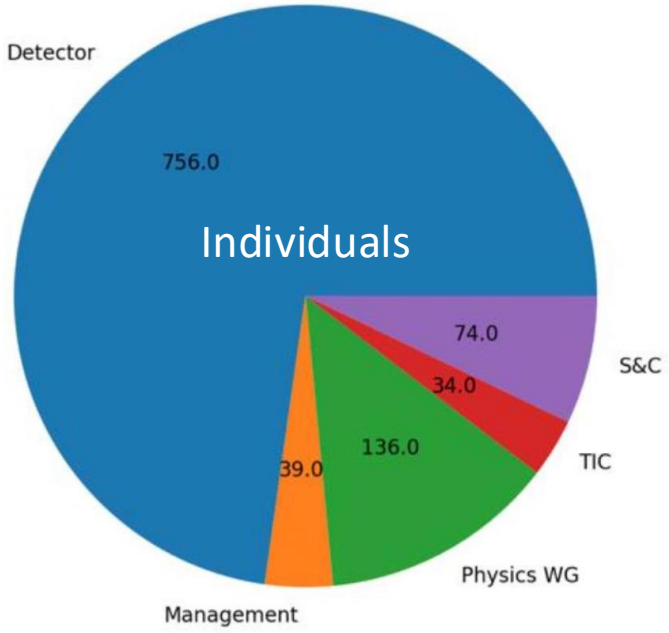

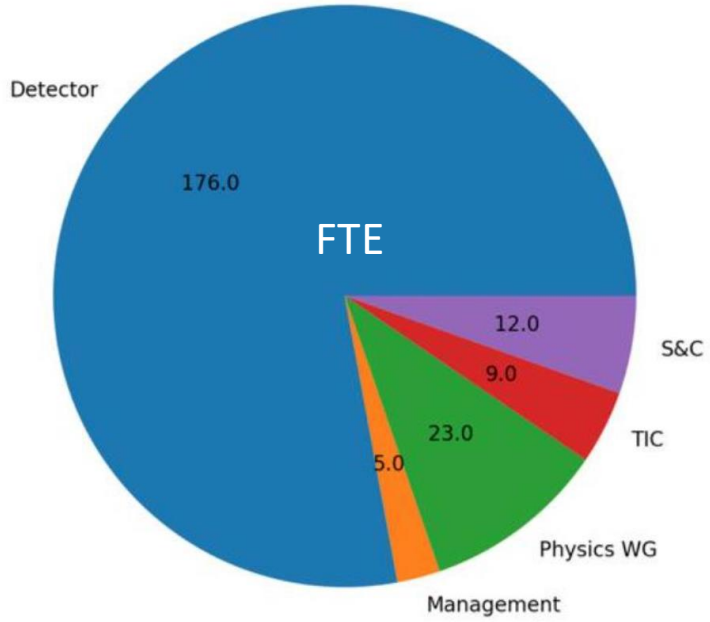

ECSAC Review

Is there a comprehensive and cost-effective short and long-term plan for the software and computing of the experiment?	**Yes**
• The pre detector technical design report (TDR) is scheduled to be delivered in 2025. Are the resources for software and computing sufficient to deliver the TDR?	**Yes**
• Is the design of the ePIC computing model and resource needs assessment adequate for this stage of the project?	**Yes**
• Is the ePIC computing and model flexible? Can it evolve and integrate new technologies in software and computing?	**Yes**
Are the plans for software and computing consistent and integrated with standard practices across nuclear physics and particle physics communities, especially given technical evolution over the next decade?	**Yes**
Are the ECSJI plans to integrate into the software and computing plans of the experiment sufficient?	**Yes**
Are the plans for the integrating international partners' contributions flexible and adequate at this stage of the project?	**Yes.**

Two Recommendation, to the Host Labs and ePIC:

- Provide a detailed plan and timeline before the next ECSAC meeting for creating dedicated effort to ePIC Software & Computing team. **Priority: Work now with ECSJI on plan and timeline, including synergy with other projects.**
- Investigate how U.S. universities can contribute to the software and computing needs of the experiment, and present a plan at the next ECSAC review.

Strategic Resource Allocation Plan

ECSAC: *"If you were given 4 dedicated experts, make a prioritized list of the tasks/areas where you would employ them for your short/medium/long term needs."*

Short/Medium-Term (next 3 years)

- Establish a dedicated effort in collaboration with Electronics & DAQ to develop integrated DAQ-computing workflows, working towards a full streaming DAQ chain test.
- Holistic full PID full reconstruction (lepton-hadron separation, lepton ID, hadron ID) implementation in the ePIC software stack utilizing the full capabilities of the integrated detector (PID, calo, tracking, etc.).
- Support AI/ML workflow integration in full simulation and reconstruction algorithms.
- ACTS expert for track seeding, track fitting, vertex finding algorithm development, tuning, and evaluation.

Long-Term (4+ years)

- Continued support for streaming DAQ workflows in collaboration with Electronics & DAQ
- Expert in fast simulations to reduce the computational cost of the simulation campaigns to interpret data
- Expert in hardware accelerators to develop collaboration expertise to speed up simulation and reconstruction and leverage HPC platforms
- Distributed computing expert to develop operations between Echelon-1/2 and support progressively scaled up challenges

Engagement With Physics & Detector Simulations WG (Chao Peng, Sakib Rahman)

- Development on detector geometries and digitization
- Background modeling and implementation
- Geometry visualization
- Manual creation of CAD drawings from the simulation models
- Automated process for the evaluation between simulation and engineering design

Consequences of Insufficient Workforce

- Till a more automated process is developed, we need to continue the process of manually creating CAD drawings from simulation geometry. Without additional workforce, this process can be tedious and hold up expertise that can better used for more advanced tasks.
- Background modeling and implementation requires efforts and additional workforce, without it the physics simulations would not be rigorous.

Engagement With Reconstruction WG (Derek Anderson, Shujie Li)

- **Example tasks available for new/interested collaboration members**:
  - **Opportunities with additional workforce**
    - Extensive to-do lists and many issues to address
      - More workforce means more parallelization in addressing them
    - Enhanced communication
      - More engaged in reconstruction → in touch with more of ePIC
      - Can identify wants and needs of the DSCs and PWGs faster
      - And distribute information *to* DSCs and PWGs faster
  - **Consequences of insufficient workforce**
    - Addressing to-do items and issues becomes increasingly serialized
      - Given limited time of workforce, only highest priority items will get addressed
    - Communication breakdown
      - Reconstruction touches on all DSCs and all PWGs, but too many meetings for just the regulars to attend
      - Can lead to lack of lateral communication between entities, possible divergence of priorities
    - Can lead to single point of failures in development:
      - Identified priorities dropped when developers leave positions

- **Example tasks available for new/interested collaboration members**:
  - Extensive list compiled in the same link as above!

Engagement With Production WG (Sakib Rahman, Thomas Britton)

Consequences of insufficient workforce

- When things run fine little involvement is required
  - However, when problems arise and both Sakib and Thomas are away it can cause significant disruptions to the monthly campaigns
- This risk can be mitigated by training more individuals in managing the infrastructure and giving more ownership to other working groups

What more WG members could do:

- More members could help develop more automated systems for running and monitoring workflows
  - Or free up other members to do so
- More members engaged with the process could better manage the needs of other working groups by acting as a Liason.
- Those with other commitments, such as students, may not be able to fully take over mission critical tasks

Engagement With Validation WG (Torri Jeske, Dmitry Kalinkin)

Contribute Benchmarks

Include detector/physics studies as benchmarks

Learn the ropes, get engaged

Integration of benchmarks into campaign simulation workflow

For rapid QA of the full dataset

Maintain and improve workflows

Improve orchestration using Snakemake

Come up with performant ways to test and benchmarks things

Image Browser

Participate in user testing

Suggest and/or implement new frontend features

Include physics plots

Consequences of insufficient workforce

- Lack people who to receive practical training means there will be lesser continuity in Validation effort.
- We may not be able to keep up with on-going maintenance of existing Continuous Integration. Without those systems running, regressions and bugs will likely make it into simulation campaigns.

Engagement With User Learning WG (Stephen Kay, Holly Szumila-Vance)

Plenty of ways to contribute and get involved:

- Tell us what you want/need to know about
- FAQ development
- Volunteer to teach a tutorial
- Participate in Mattermost issues/discussions

- **Opportunities with additional workforce**
  - Broader range of tutorials (and more frequent)
  - More perspectives/skillsets
  - Rapid updating of FAQs/tutorials

- **Consequences of insufficient workforce**
  - Less able to respond to needs of users
  - Tutorials less frequent/relevant
  - Not able to pursue new initiatives such as manuals

Engagement With Streaming Computing WG (M. Battaglieri, J. Huang, J. Landgraf)

Publication on "ePIC Streaming Computing Model": We **welcome your feedback now via email** and aim to submit the manuscript to *Computing and Software for Big Science*.

The ePIC Streaming Computing Model
Version 2, Fall 2024

Marco Battaglieri[1], Wouter Deconinck[2], Markus Diefenthaler[3], Jin Huang[4], Sylvester Joosten[5], Dmitry Kalinkin[6], Jeffery 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.
[6]University of Kentucky, Lexington, KY, USA.

Abstract

This second version of the ePIC Streaming Computing Model Report provides a 2024 view of the computing model, updating the October 2023 report with new material including an early estimate of computing resource requirements; software developments supporting detector and physics studies, the integration of ML, and a robust production activity; the evolving plan for infrastructure, dataflows, and workflows from Echelon 0 to Echelon 1; and a more developed timeline of high-level milestones. This regularly updated report provides a common understanding within the ePIC Collaboration on the streaming computing model, and serves as input to ePIC Software & Computing reviews and to the EIC Resource Review Board. A later version will be submitted for publication to share our work and plans with the community. **New and substantially rewritten material in Version 2 is dark green. The present draft is preliminary and incomplete and is yet to be circulated in ePIC for review.**

https://doi.org/10.5281/zenodo.14675920

1

- **The next steps will be discussed at the parallel SRO session at the collaboration meeting next week**
  - Support of test beam measurements
  - Streaming Orchestration
  - Streaming Challenges in (Echelon 0 & Echelon 1)
  - Distributed Data Challenges (Echelon 1 & Echelon 2)
  - Analysis (Echelon 3, in conjunction with services to be provided by Echelon 1 & Echelon 2)
  - Autonomous Experimentation and Control (Alignment, Calibration, Validation)

- **Consequences of insufficient workforce**
  - Not *YET* at point where the SRO prototyping activities are on critical path
  - The SRO working group acts as an interface between Electronics & DAQ, the various software working groups, the background group, and detector groups.
    - Participation and communication from these groups critical
  - Refinement of the computing model and Roles of diverse institutions
    - Input from and participation of world-wide computing facilities is critical

Engagement With Wider Community

- ePIC Software is being used for efforts on 2nd Detector and SoLID.
- We are an active part of the software & computing community in NHEP:

Slide from Dmitry

Giving back to the community

- Contributed fixes for
  - **ACTS**: improvements in tracking and infrastructure, with plans for more development
  - **DD4hep**: general bugs, HepMC3 ROOT and gzip support, new readout segmentation types, …
  - **EDM4hep**: build infrastructure
  - **Geant4**: addressed bugs UI, performance optimizations
  - **HepMC3**: addressed bugs in ROOT/XRootD IO, build infrastructure
  - **PODIO**: API improvements, build infrastructure, memory access semantics, `streaming over network`
  - **ROOT**: addressed bugs in UI, regressions in ROOT file handling
  - **Snakemake**: adressed general bugs
  - **Spack**: maintenance work on the package set
  - **UpROOT**: S3 object storage support (for legacy test setup at BNL/SDCC)
- EICrecon provides an ever expanding set of reusable LGPL3-licensed algorithms
- For HEP-NP projects, we encourage our users to submit their questions and problems to ePIC Helpdesk first

Review of ePIC Computing & Software, 2024, Dmitry Kalinkin 6

- ePIC leadership in software & computing has been recognized by HEP, e.g. CERN, HEP Software Foundation (HSF), and US-FCC.
- It is a priority to continue this engagement and collaborate with the wider NHEP community on common software and standard practices on computing, something we have benefitted immensely.

Common Themes from Priority Discussion

Need a better way to grow/onboard analysis efforts

- Support PWGs in adding analysis information to the landing page (where appropriate) and updating the list of existing tutorials.

- Continue to prioritize reconstruction efforts in collaboration with the PWGs:
  - **Current priorities with PWGs**: Secondary Vertexing, Particle Flow, PID, Event Kinematics (physics objects).

- Leverage joint meetings to further discuss representing physics objects in simulation outputs and integrating analysis methods and tools into our software stack.
  - Next joint meeting will be on February 5.

- Maintain our commitment to holding weekly AC/SCC meetings.

Common Themes from Priority Discussion

Need to define a set of 2025 milestones for progress in simulations

- This includes:
  - **Fidelity to the engineering design**: Need to discuss with EIC Project scope of co-design between engineering and simulations and goals for 2025
  - **Simulation campaigns**
    - Continue to define needs with PWGs
    - Work more closely with DSCs on needs, e.g., specialized geometries.
  - **Inclusion of backgrounds (**see slide 6: Requires more workforce to sustain effort, also OPA background**)**
  - **Reconstruction and common tools (**see slide 7 for status and next steps)
    - Continue to work with ACs on setting reconstruction priorities, using joint meetings for broad input.
    - Co-lead discussion on how to represent physics objects in simulation outputs.
    - Discussion will define scope and goals on integrating analysis methods and tools into our software stack.

- **Work with DSCs to set software goals**
  - Goals could include digitization (simulation of electronics), noise models, reconstruction, simulation needs, test-beam support.
  - Support would include support for integration, simulation campaigns, tutorials, e.g., on digitization, event display.

- **Define how we work towards simulation reconstruction in timeframes**
  - Urgency in defining streaming computing projects for 2025 to address review recommendations and related funding opportunities. Streaming reconstruction will be part of these discussions.