

Live Notes from CCE AHM June 2025

Please fill out the Poll at <https://forms.gle/h86fJdMmm8acUe1m6> if you haven't already.

Monday June 2

PAW session

Meifeng - Intro

- Salman: we should avoid fragmenting effort across the many workflow portability solutions, and avoid creating our own.

Tianle - Globus compute

- Dirk: level of support from HPC centers for GC? NERSC superfacility, OLCG s3m(?)
 - Doug: supported by multiple HPCs, runs in user space,
- Marco: how much effort to move from perlmutter to other sites?
 - Effort needed because of lack of outbound connectivity
 - Would be good to provide a unified HPC proxy to communicate with e.g. HTCondor

Ozgur - Dune 2x2 ND simulation portability

- Portability solution for ND sim kernels. Which one to use and how much work would it be?
- Hardcoded fs paths? Should they be there?
- Outbound connectivity preventing HPC use seems to be a theme...
- Why CUDA-only?
- Is the CUDA part computationally expensive? One can use the IaaS to run that algorithm remotely on a site that has NVIDIA GPUs.
- Amit:
 - Meghna Bhattacharya and I are the level 3 manager tasked for the accelerator integration in the DUNE software. We had recently done a survey that looked at how DUNE (ProtoDUNE, DUNE-ND, DUNE FD related projects) uses the accelerators currently. As discussed most of the development is around CUDA but perhaps because of the resources at disposal.

- Anyway, since Ozgur did the portability exercise of porting 2x2 from Perlmutter to Polaris, I guess it would be good for us if the porting exercise is/was/has been documented somewhere, like what were the things that were easy to port, what were the parts that had most work to be done and what kind of changes were done to port, what was the performance after porting in Polaris etc.
 - Bruno: We have some of this documentation here:
https://github.com/hep-cce/DUNE_2x2_sim_portability

Charles - miniapps

Meghna - DUNE use cases

- From art event loop to phlex functional-style parallel programming
- GausHitFinder
- Charles: compare with Gaudi Functional Algorithms experience about corner cases that are hard to port to functional programming
- Paolo: how phlex work connects with CCE activities like portable parallelization?
 - Peter: discussions ongoing

Esteban - automated numerical analysis for SYCL

- proof-of-principle P2r possible test case?
- Meifeng: could this work on other GPUs?
 - POCL should support CUDA
- Rob: similar efforts in the software verification tools

Atif - AI codegen

- Salman: important to follow industry. Expensive LLM now much more powerful than the one tested by Atif and team.
- XJ: vscode based models much better than only a couple of years ago
- Seth: what is HEP-specific about our code to need a new framework
 - RAG with HEP specific idioms
- Meifeng: need LLM that understand HPC code
- Tom: do we have access to the expensive models mentioned by Salman?

Rob - Transient Detector Pipeline

- Salman: how is this work relevant to LSST transient broker?
 - Through AMPEL

Saba - SOP introduction

Chris - RNTuple

- Philippe: ROOT team grateful for all input from CCE
- Paolo: what next for this activity
 - Scoring parallel reading
- Nick: what about future data models like SOA.
 - Philippe: what flavour of SOA, e.g. CERN group developing a new library that goes beyond C++.
- Philippe: plan to revamp TFile interface. Would be good if CCE could take a look

Rui - RNTuple for GNN4ITk

- Doug: How much data in training sets
 - XJ order of TB
- Some discussion on how to reduce conversion steps for GNN4ITk.
Ultimate RDO- (athena) → RNTuple – (pyTorch/TF DataLoader) → tensor

Wes K (C2P2 SOP summer student): integrating Darshan into ATLAS SPOT

Ami (C2P2 SOP summer student): SZ3

Seth: Celeritas

- NVIDIA A100 does not support ray tracing like gaming cards
 - Requires NVIDIA OptiX, tessellated geometry, etc

Taylor: ev gen

- NVIDIA Blackwell and later emulate F128, slow but faster than dropping back to CPU.

Elliot - Orange geometry

- What are surface primitives?
 - A handful, surfaces, tori, ...

Julien - Celeritas in ATLAS

Hayden - Celeritas Optical Physics & Calvision

Gabor - Opticks

- Tom: why CPU performance does not scale linearly in slide 4 Is this coming from NUMA? Plus HT
- Meifeng: will G4 ever run on GPU?
 - Seth: not natively, but we will progressively integrate GPU models.

Walter - SML intro

- Foundation models for HEP being attempted in Europe, CMS, SLAC. CCE should play a role in providing scaling and resource

Alina - Distributed training for GNN

- Can we learn about distributed training for GNNs from what has been done to train transformers?
 - Not sure, will investigate
- Are there common features in HEP graphs that may be used to optimize GNN training?
 - Probably not. Medium size graphs, many of them

Xiangyang - IaaS

- How many events in // to hide remote offload latency
 - Async call is key, number of events model dependent
- Ken: DUNE has approach with C&S in same node
- Meifeng: how to handle different requests to same server
 - Thinking to launch dynamically servers specific to workloads

Dennis - HPO for NSBI

- XJ: very nice project could reduce time to HPO result by weeks if we can package whole optimization to a single HPC job

Tuesday, June 3

Dirk - CMS

- Importance of quarterly/monthly leadership meetings with experiments
 - Could be in the context of monthly tech leads meeting
- Doug: data placement crucial
- Walter: agents for operations
- Taylor: fraction of production code on GPU?
 - Offline none running. HLT only
- Meifeng: at what stage would you like to test the code generators
 - Chris: any time!
- Charles: will alpaka decision be reviewed
 - Matti: not likely before Run 4. It works. Support mostly from small number of CERN people

Gordon - IRIS-HEP

- 160 Undergrad students got DIANA fellowship. Would be great if CCE people would act as mentors. Would fund any student project matching IRIS mission
- CCE participation in US ATLAS/US CMS Ops training
- Analysis on HPC: Perlmutter >30 active ATLAS users 17 used >100 node hours
- Bridge between UChicago AF and NERSC. Lincoln's work with Doug on access to Perlmutter batch queues. Also IaaS UC/NERSC.
- Invite IRIS-HEP experts to AGC to meetings on IaaS to make them aware of what is happening
- Peter: Using RNTuple potential area of collaboration. Writing RNTuple that can be read efficiently in analysis
 - Gordon: Benchmark RNTuple dataset. IRIS could help defining these benchmarks
- Nick: analysis on GPU in CMS. Decompressing RNTuple on GPU, fill histograms with cuPY, could use collaboration.
 - Histogramming with cupy promising candidate for a library
- XJ: standardizing kernel implementations
 - CCE expertise on running on GPU could help with e.g awkward-GPU kernels.

Ken - DUNE

- Charles: How different DUNE RNTuple requirements?
 - Peter: very different data sizes. LHC events small. Phlex non-event based different access patterns.
- Paolo: ARM port easy compared to GPU port
- Matti: but data races due to different memory architecture are a problem. ROOT 6.36 first release that addressed all known ones

Verena - ATLAS

- Add CCE AHM to HSF calendar
- [SOFIE](#) (CERN/Manchester)
- Meifeng: is there still an interest in portability layers like kokkos alpaka sycl etc?
 - Not yet. Focus is on writing modern C++. ATLAS GPU tutorial important to provide basics of CUDA. Participants expressed desire of more advanced CUDA material.
- Stefan: What is the issue with Pepper/madgraph4gpu workflows?
 - Questions about how to evaluate generators on GPUs, e.g. which resources to test it on
 - What should be the next step for the generators community?
 - ATLAS generator group very interested in testing Sherpa/PEPPER. Good time to start discussion on how to test them in ATLAS workflows. Will have a discussion during next software week in July
 - **Paolo will coordinate with Stefan/Taylor**
- Taylor: any interest in distributed training?
 - Ease of use is crucial. GNX and GNN4ITk being trained on perlmutter. Order of magnitude speedup. Same for NSBI on SMU cluster 1 week → 1 hour

Ben - Rucio/Globus

- Doug: CCE should put together FTS, LHC experiments, and globus team and converge on an authentication scheme.
- Marco: can Rucio authenticate users to Globus?
 - Separate token mechanism

Maria Elena - LZ

- Charles: other issues besides reliability with SPIN?
 - Authentication, policy → significant coding effort
 - Salman: globus keeps their databases on the cloud - on prem cloud would be ideal for HPC centers

Salman - plans for parallel sessions

- Much more solid connection with experiments, tight feedback
 - List CCE activities, experiment contacts, working groups
 - Celeritas good example
 - Prioritize (in case of cuts)
 - Experiment interest/adoption
- Ramp up connection with LCFs
 - One or two solid ideas for IRI pathfinder

- NERSC 10

=====

Joint SML/SOP Session - AI-ready Data

Training as a Service: nascent efforts,

HPO on HPCs: DeepHyper@ANL
XJ propose ATLAS FTAG as a dataset to test

RDO→pyTorch
Doug: could use ServiceX to do data transform

Uproot can turn a RNTuple into a numpy array in memory without passing from disk

Peter: we have the effort and expertise, we should start. Rui slides provide the requirements in terms of what we need from ROOT

Joint SML/PAW Session - IaaS

Kevin: scheduling needs of GPU and CPU very different. CPU jobs requests change very fast.
SuperSonic dynamic scheduling based on latency in responding
XJ: CCE can help find solution to maximize throughput

More notes:

<https://docs.google.com/document/d/1pZSER-qCablkasB5qtQh4m4TZrowGCxeQAWxoHEQahc/edit?usp=sharing>

- IaaS good candidate for IRI pathfinder. SML estimates we are missing ~1 FTE ← PAW?

=====

SML Parallel Session

-

=====
=====

PAW Parallel Session

-

=====
=====

SIM Parallel Session

-

=====
=====

SOP Parallel Session

Wednesday, June 4

Plenary closeout

- Ss

Tech Leads Meeting

- Input from experiments on organization:
 - Have regular tech leads meetings with experiment representatives. Monthly or Quarterly?
 - Increase participation in experiment software weeks
 - For example: invite CCE to discuss possible ATLAS evaluation of Sherpa/PEPPER during ATLAS July software week
 - Present celloAI to experiments (e.g. ATLAS in July, discuss with chATLAS developers)
 - “Mini-blueprint” meetings with IRIS-HEP. For example on RNTuple reading performance and usability for Analysis workflows

PAW Plans for FY 2026 and Beyond

- Verena: include some miniapps as HEPscore GPU benchmarks
- Work with NERSC through NESAP N10
- Work with ANL 4 team (Taylor!)

SIM Plans

- Need experiment contact person to consult on requirements, framework, data access, workflows

SML Plans

- Verena: NSBI toolkit Kyle Cranmer group as part of IRIS-HEP
 - Aware of activity, plan to scale up the existing and future NSBI code. Also ROOT interested in NSBI
- Training/educational activities

SOP Plans

Tech Leads Meeting

- IaaS: 1 FTE short. FNAL effort coming after reorganization. NERSC 10 Pilot. Talk with Tom Uram at ANL
- Data transformation for AI (from ROOT)
- SML AI-assistant effort, add Atif from BNL, Steve can help with GPU portability
- Rucio/Globus – Doug/Ben/Marco + experiments (CMS, DUNE, Rubin/LSST, LZ, –)
- IRIS-HEP interaction on training/mentorships
- Regular meetings with experiments (quarterly)
- CERN involvement and HEP-CCE visibility, increased engagement in meetings
- Pepper interest from ATLAS
- HEP-CCE documentation for using HPC systems, AI tools, etc.
- Web presence