A Coordinated Ecosystem for HL-LHC Computing R&D

Workshop Closeout
2022-11-09

Workshop Overview

The coordinated ecosystem workshop aimed to answer the following three questions:

1. How does the **ensemble of US Software R&D efforts fit togethe**r to implement the HL-LHC Software/Computing roadmap and meet the challenges of the HL-LHC? Which areas are not covered by US R&D efforts and should have international coordination? Which areas present new challenges or new opportunities since the Community White Paper (CWP) process that was executed in 2017?
2. How do the US Software **R&D efforts collaborate with each other and with international efforts**? How do these efforts align with and leverage national exascale, national NSF OAC priorities and trends in the broader community?
3. How should the **US R&D efforts be structured and evolved** in the coming years in order to achieve our goals between now and the HL-LHC era?

Workshop Activities

Day 1 consisted of 13 presentations across the breadth of HL-LHC and related experiment and lab R&D projects.

- Each presentation highlighted their scope, how they intersect with the HL-LHC experiments, and their place in the R&D ecosystem.

Day 2 focused on specific breakout topics the community felt could have the largest coordination impacts:

- Storage & Network R&D
- Workforce Development
- AI/ML Coordination.
- Resources & Facilities
- Analysis Facilities & Facility Evolution.
- GPU Algorithm Development

A full summary will be prepared for the workshop outcomes report.

Workshop website with slides

Observation – Value in coordination activity

The group finds extraordinary value having this workshop to explicitly coordinate R&D efforts.

- Helps to play toward the expertise of different groups.
- This would be more valuable if it was sustained and regular.

While IRIS-HEP hosted this workshop, as a long-term activity, it should be independent of any given project

Today’s workshop focused on the HL-LHC science driver; however, participation was cross-cutting throughout HEP and particle physics.

- Experiments included DUNE, Dark matter, EIC, Vera Rubin.
- DOE Lab participation included BNL, LBNL, FNAL, ORNL, SLAC.
- NSF AI Institutes.
- … and more!

Potential to evolve this to a standalone workshop or coordination project.

- E.g., there was an explicit recommendation along these lines from the Snowmass process.

A planning process is underway in 2022 - IRIS-HEP Example

The last community planning effort was in 2017 and produced two documents:

- "Strategic Plan for a Scientific Software Innovation Institute (S2I2) for High Energy Physics"
- A Roadmap for HEP Software and Computing R&D for the 2020s (Community White Paper)

Over the past year a number of events have take place to update that planning:

- ✅ Analysis Ecosystems II workshop (23-25 May, Orsay)
- ✅ Connecting the Dots 2022 (31 May - 2 June, Princeton)
- ✅ Snowmass Community Summer Study Workshop (17-26 July, Seattle)
- ✅ Differential Programming Workshop (12-16 September, Crete)
- ✅ PyHEP Workshop (12-16 September, Virtual)
- ✅ IRIS-HEP Team Retreat (12-14 October, Princeton)
- A Coordinated Ecosystem for HL-LHC Computing R&D (7-9 November, Washington, DC)
- Software Citation and Recognition for HEP (22-23 November, Virtual)

Similar process is ongoing for HEP-CCE.

Many of these were organized or co-organized by the IRIS-HEP team.

IRIS-HEP aims to deliver an updated version of the Strategic Plan in Dec., 2022.

Observation - Value in planning connected projects

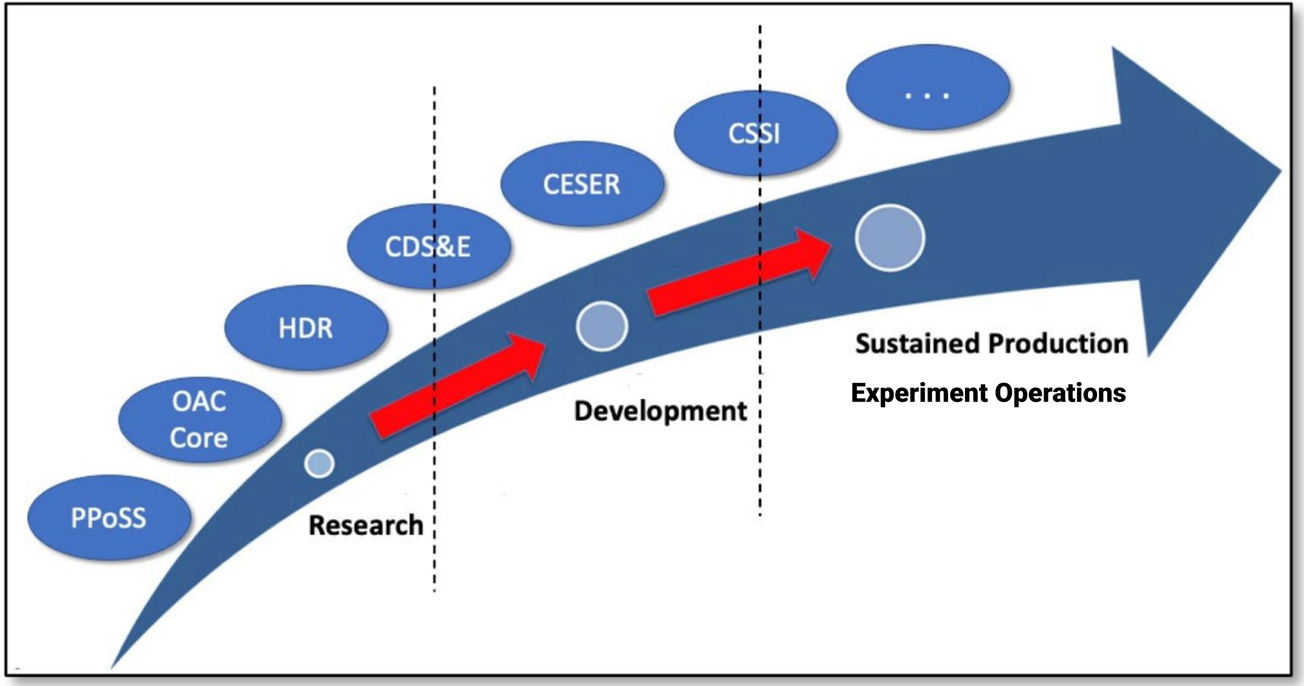

Figure 2: Data and Software CI pathways to production.

Example Projects at different stages and with different mandates in different R&D stages

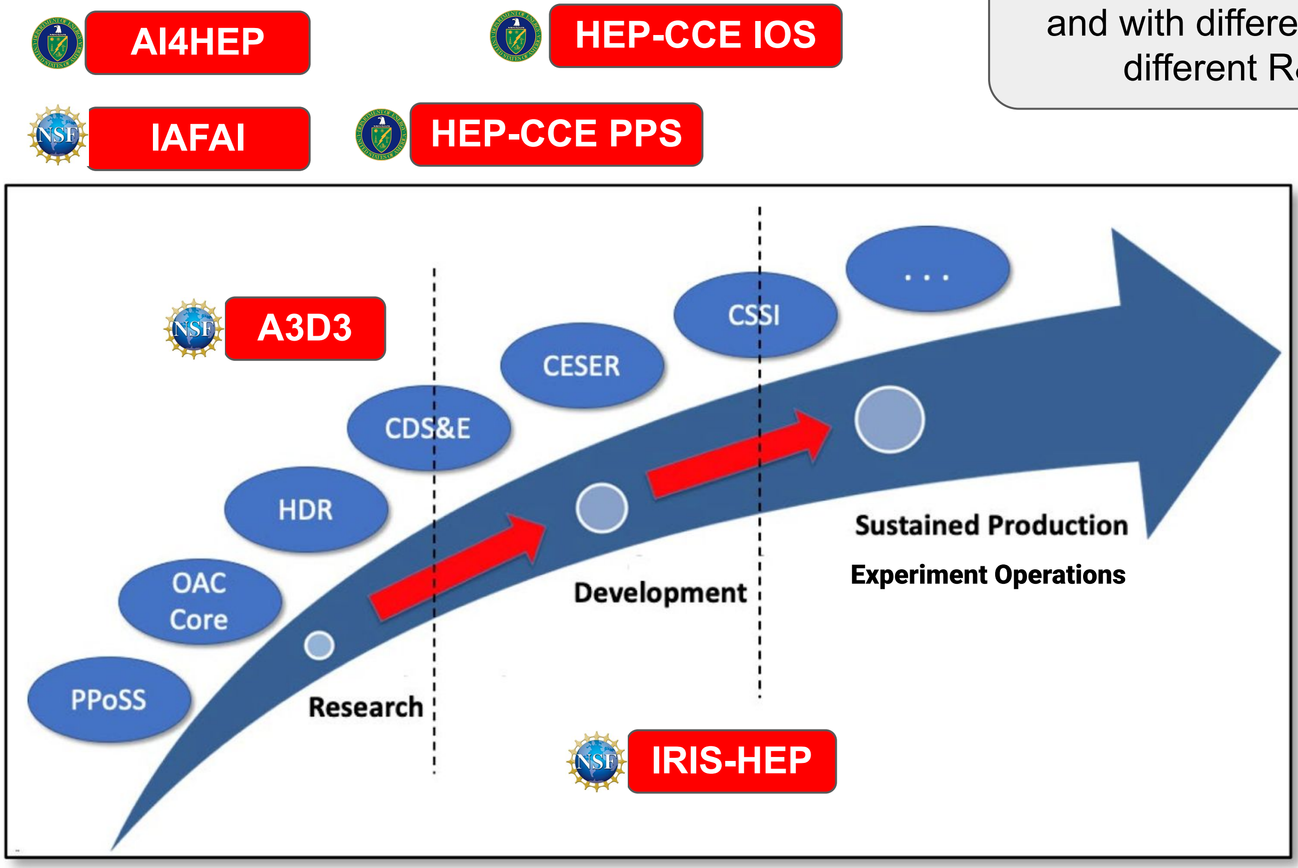

Figure 2: Data and Software CI pathways to production.

Software and Computing Challenges for the HL-LHC

C1. **Raw resource considerations** (CPU, storage) needed, with a flat funding model

C2. **Scalability** of the overall distributed computing cyberinfrastructure and its components

C3. **Analysis at scale** with the HL-LHC data volume/rate and complexity (user-centric)

C4. **Sustainability** of the system through the science lifetime of the HL-LHC

Map of R&D Activities throughout the ecosystem

- Physics Generators
- Machine Learning
- Visualization
- Detector Simulation
- DOMA
- Software Development, Deployment, Validation, Verification
- Software Trigger & Event Reconstruction
- Data-Flow Processing Framework
- Data and Software Preservation
- Data Analysis & Interpretation
- Networking, Storage, Infrastructure, Facilities
- Security
- Training

Building a spreadsheet for a more detailed look in the final report

Project Name	Short Project Description	CWP Area	Sponsor	Life cycle status
HEP-CCE PPS	portable parallel strategies	Processing Frameworks	DOE	Development
FAIROS-HEP	Resource Coordination Network on FAIR Data in HEP -> Living Publication	Data and Software Preservatio	NSF	Development
HEP-CCE IOS	I/O strategies on HPC and beyond	Processing Frameworks	DOE	Research
HEP-CCE Generators	Modernize event generators for heterogeneous architectures	Physics Generators	DOE	Research
Coffea	awkward-based analysis tools	Data Analysis & Interpretation	DOE	Development
Object Storage for Physics	Using key/value object store for physics analysis	DOMA	DOE+NSF	Research
Elastic Analysis Facility	Multi-experiment scalable AF	Data Analysis & Interpretation	DOE+NSF	Development
Particle Flow GPU adaptation	Adapt CMS particle flow code to run on GPUs	Software Trigger & Event Reco	DOE	Development
CREST	Fermilab's Computing Resources Evolution STrategy	DOMA	DOE	Research
End-to-End EM Recostruction	AI-based EM reconstruction (CMS)	Software Trigger & Event Reco	DOE	Development
Segment-Linking Tracking	New tracking algorithm for CMS Phase-2 Tracker	Software Trigger & Event Reco	NSF	Development
	Use SENSE paths for Rucio transfers	DOMA	DOE	Research
HEP-CCE workflows	Explore parsl and funcX to schedule complex workflows on HPC	DOMA	DOE	Research
AI Approaches to Simulation	improve GANs and other AI techniques (CMS)	Detector Simulation	DOE	Research
Framework LDRD (Fermilab)	Develop framework with user-defined processing levels	Processing Frameworks	DOE	Development
GPU Integration LDRD (Fermilab	Develop GPU-based algorithms for multidimensional integration	Physics Generators	DOE	Development
Data Lake/infrastructure studies	study cache/network/disk interplay (CMS)	DOMA	NSF	Research
WM for Heterogeneous Resource	(CMS)	Workflow Management (Other)	NSF	Development
Analysis Facility Scalability	back-end scaleout (CMS)	Data Analysis & Interpretation	NSF	Development
Vertexing on GPUs	rewrite CMS vertexing code for GPU	Software Trigger & Event Reco	DOE	Development
SONIC	remote execution on accelerators for AI inference/processing	Processing Frameworks	DOE	Development
WM for HPCs	(CMS)	Workflow Management (Other)	DOE	Development
CompHEP Traineeship Grants	Training for graduate students	Training	DOE	Development

Areas of Highlight

We'll highlight 4 areas discussed through the workshop:

1. Scalable Storage and Networking (C2)
2. Translational AI (C1, C3, C4)
3. Facilities Evolution and Integration (C1, C2, C3, C4)
4. Algorithm R&D (C1, C4)

While not comprehensive, each area is critical for making progress against the HL-LHC challenges and has existing R&D efforts that need to coordinate across the community.

Areas of Highlight - Data Challenge

Thanks to ESNet HEP Requirements Review, we have good models of the WAN activities expected during HL-LHC.

- IRIS-HEP and WLCG have kicked off a series of data challenges to show our services & systems can scale up to the needs for HL-LHC.
    - DC21 -> 10% of HL-LHC scale. DC23 -> 30% of HL-LHC scale.
- These data challenges provide specific integration and evaluation checkpoints for R&D activities.
    - DC21 -> HTTP-TPC
    - DC23 -> Packet marking, Engineered network paths, tokens-based authorization.
- Additional “mini-challenges” as prep for the DCxx are being discussed in the community, and require coordination.
- At the site-level, provide a mechanism for evaluating progress of facility storage R&D.

From

500Gbps

Today

To

9,600Gbps

For HL-LHC

Data Challenge - Active R&D

- **Leveraging engineered network paths** (ESNet SENSE, USCMS Ops, IRIS-HEP DOMA): Dynamically bring up dedicated network paths (providing some sort of guarantee) associated with a flow.
- **Packet marking** (IRIS-HEP OSG-LHC, CERN, ESNet): Either through IPv6 headers or “firefly packets”, inform network monitoring of identity of the flow (CMS? ATLAS? Priority? Data type?).
- **Storage architecture tailored for workflows** (BNL LDRD): Instead of having a single hardware/software/service solution for all HL-LHC workflows, can we utilize some differentiation to optimize cost? Example: tiered storage at BNL - less IO-intensive workflows on lower-QOS disk? Example: At T2s Random Access media for high performance, HDDs for cost effective volume.
- **Demonstrate scaling of existing solutions** (IRIS-HEP DOMA). Can we take an existing set of production software (XRootD + FTS + Rucio) and run it at HL-LHC scale on testbed hardware?
- **Technology evaluations across the community** (cross-cutting): As always, continuous evaluation of new storage solutions like Ceph (and/or including vendor solutions - e.g., VAST, WekaIO, …).

Data Challenge - Coordination Activities

Existing important coordination activities:

- Global cross-cutting Data Challenge activities toward **DC23**.
- WLCG Data Organization, Management, and Access (**DOMA**) working groups.
- Monthly **ESNet** Meetings.

Potential new activities:

- **Dedicated US-DC23 planning workshop**: Identify the US resources to use for DC23, develop a more concrete set of goals for the US resources, and plan any intermediate mini-challenges.
    - Important to include DUNE here as they will also leverage the shared WAN resources.
- **Unknown / Unclear needs**:
    - Better engage with DOE facilities around the Integrated Research Infrastructure activities (or other ASCR projects)?
    - Better engage with DOE ASCR to find other projects with complex WAN data flow needs.

Areas of Highlight - Translational AI

HEP has a long history of using AI / ML and has been extremely active in embracing and contributing to modern AI (mainly driven by deep learning).

- Modern AI has proven to **transformative**: they are effective with our complex low-level sensor data.
- Need for **translational AI:** techniques developed by AI community do not work out of the box for our problems.
- Excellent examples of **use-inspired research:** context-specific challenges drive foundational AI advances
    - Several papers co-authored by AI experts in academia and industry published in AI venues and multi-disciplinary journals
- One of HEP's most compelling examples of workforce development and broader impact

Exploratory AI R&D often done outside of experiments on fast simulation etc. (faster iteration, facilitates innovation).

- Significant effort required to integrate into the experiments production / operations (**this inefficiency is a gap**).
- Modern AI techniques **are being used in published analyses** and are displacing previous approaches
- **Widely accepted** that modern AI will be a part of multiple components of HL-LHC software and computing (trigger, simulation, reconstruction, and analysis)

Coordination is needed to guide the evolution from this period of rapid R&D, prototyping, and bespoke solutions for deployment to a more mature / established set of practices for ML in various contexts (e.g. trigger, reconstruction, simulation, analysis)

Translational AI - Existing R&D

- NSF AI Institutes (e.g. IAIFI)
- NSF HDR Institutes (e.g. A3D3)
- ML R&D components of IRIS-HEP (and previously DIANA-HEP)
- DOE AI 4 HEP awards, etc. (lab / university / industry partnerships)
- Misc AI/ML focused awards FAIR4HEP, EAGER
- A component of many base grants
- Ops programs support various AI / ML efforts and projects such as SONIC
- Lab efforts: e.g. Fermilab's Computational Science and Artificial Intelligence Directorate (CSAID), LBNL Scientific Data Division, etc.
- International: ELLIS, Punch4NFDI, etc.

This list is not complete, see talks from Monday

Translational AI - Coordination Activities

Existing important coordination activities:

- Living review, contributions to Snowmass CompF3, CERN IML
- Workshops: ML4Jets, ML4PS @ NeurIPS, Hammers&Nails, Aspen, IPAM, Dagstuhl, MIAPbP
- IRIS-HEP, FAIR4HEP, A3D3 (HDR Hub), FAIROS-HEP, multiple training activities

Potential new activities:

1. **Retraining challenge**: retrain / fine tune established ML components being used in production when run conditions or calibration change
   a. Would help focus effort on streamlining training > production; automated workflows, meta-data, versioning, provenance
2. **Connect benchmarking infrastructure to deployment infrastructure** used in experiments to reduce the gap between R&D and production
   a. Distinguish between trigger, reconstruction, simulation, and analysis context
   b. Coordinate running of “ML benchmarks / challenges for LHC” that are popular within the CS ML community.
3. Develop **realistic Open Simulation** (building on top success with tracking) and expand **accessibility** of Open Data to support AI research
4. Utilize **centralized testbed & expertise** dedicated for distributed research for AI on FPGAs / accelerators / highly constrained computing environments (connect to other projects in this area).

5: Workshop to explore impact of move to a “foundation model” paradigm (larger, multi-purpose model)

Areas of Highlight - Facilities Evolution and Integration

The HL-LHC experiments have high-level numbers for disk, CPU, and tape requirements.

- **Not all resources are equal**. CPU resources are *potentially* fungible with HPC CPU/accelerator resources (many challenges). Disk and custodial storage is not.
- What are the consequences of non-homogeneous sites?
- LHC community needs to provide sites with better guidance for system requirements (particularly storage throughput / IOPS) and design a build-out plan for T1s and T2s.
- Missing from the resource curves are **sensitivity analyses** - e.g., what's the budget impact of a 20% increase in tape needs?
- BNL T1 presented a budget exercise using historical cost curves - helps illuminate impact of various scenarios.

From E. Lancon's presentation on the BNL T1.

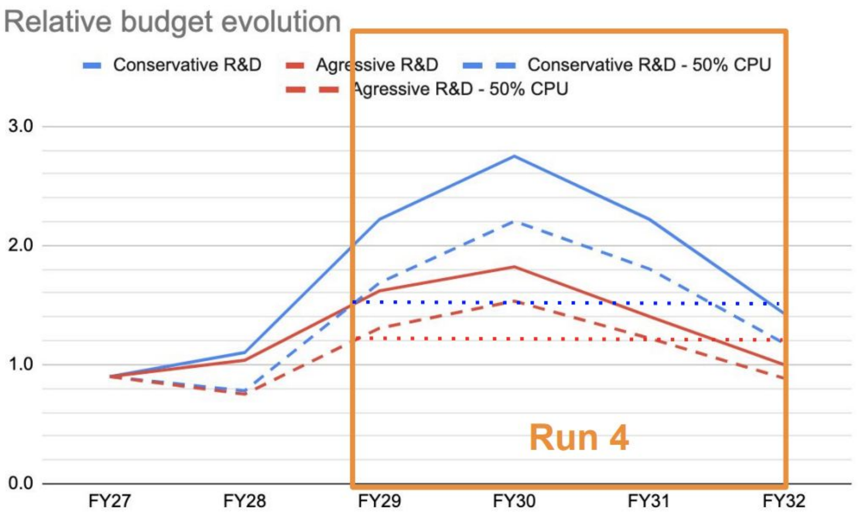

If software allows **offloading 50% for CPU** requirement to other facilities (like HPCs)

- Conservative R&D Scenario - 50% CPU **1.5 x Flat**
- Aggressive R&D - 50% CPU : **1.2 x Flat**
- In both cases, **disk costs >50% of total**.

Facilities Evolution and Integration - Active R&D

- **Storage evolution & specialization** (BNL LDRD, FNAL core, US-CMS Ops/R&D): Changes or specializations in storage systems enabling
- **Composable, flexible facility services** (IRIS-HEP SSL): Using new industry technologies such as Kubernetes, enabling new deployment paradigms.
- **Analysis Facilities (**cross-cutting; *Ops, IRIS-HEP AS, FNAL, BNL): Specialized facilities & services for increased dataset size, new capabilities, and new techniques expected for HL-LHC.
- **Numerous HPC integration exercises** (USATLAS Ops, USCMS Ops, IRIS-HEP OSG-LHC): Leverage computing available at the wider range of DOE ASCR and NSF OAC resources.
  - N.B. - Coincidentally, the Joint Blueprint on Cloud/HPC activities initial draft came out this week.
- **Integrating services into the network** (ESNet, FAB/FABRIC, IRIS-HEP DOMA): Running future services (caches, data delivery, SENSE) at network locations instead of sites).

Facilities Evolution and Integration - Coordination Activities

Existing important coordination activities:

- Analysis facility work is coordinated through existing experiment Ops Programs, WLCG, and IRIS-HEP (US-centric) efforts.
- Cloud/HPC integration is led through the Ops programs.

Potential new activities:

- Develop a facilities working / consulting group to coordinate, esp. considering impact of newer analysis ideas.
- Host a workshop to develop guidance for missing site-level requirements and begin build-out plan to HL-LHC.
- Gather HEP requirements for LCFs in the style of the ESNet requirements review.
- Evolve the raw capacity plots to better indicate R&D impact.
    - Can be US-internal to help guide R&D prioritization.
    - E.g. - use historical purchasing power evolution and show cost sensitivity (have bands showing +/- $1M investment in hardware)

Areas to Highlight - Physics Algorithms

Algorithmic improvements are central to reducing CPU need and increasing science reach

- Algorithmic adaptation/re-write to run on accelerators
  - Madgraph and Sherpa event generators; Celeritas for detector simulation; CMS Particle Flow and vertexing, etc.
- Adopting successful approaches from other fields
  - Celeritas has grown from ECP
- New approaches aimed at dramatic speed increases
  - mkFit, ATLAS FastChain/FastCaloSim, Segment linking tracking,
- AI-based algorithms
  - GNN Tracking, CMS HGCal Reconstruction, Exa.TrkX, Atlas FastCaloSim

Transitions from R&D to production use in experimental software environments are starting. Each will take substantial time and effort

- Combinations of approaches can be used to capitalize on the strengths of different projects

Physics Algorithms - Examples of existing Tracking R&D

Multiple approaches to HL-LHC tracking have demonstrated improvements over the state-of-the-art approaches

- Reengineering existing algorithms
- New domain-specific approaches
- Machine learning approaches

mkFit Vectorization

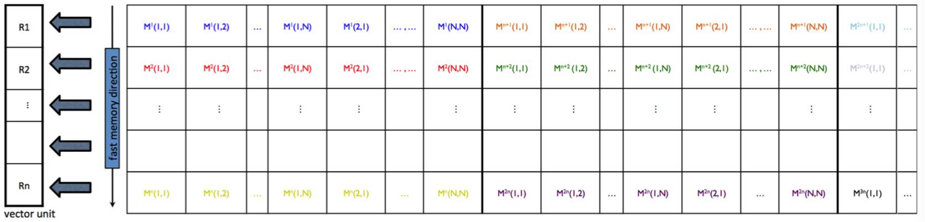

ExaTrkX **ATLAS ITk Performance Plots**

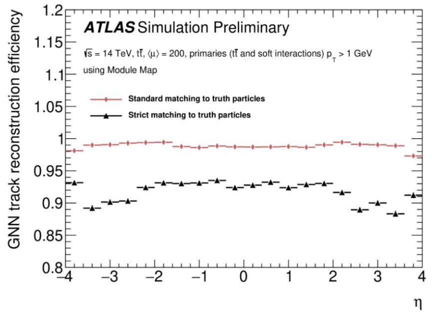

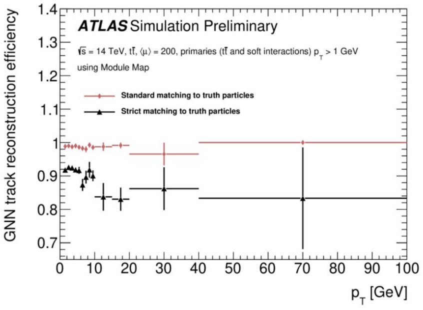

Segment Linking approach

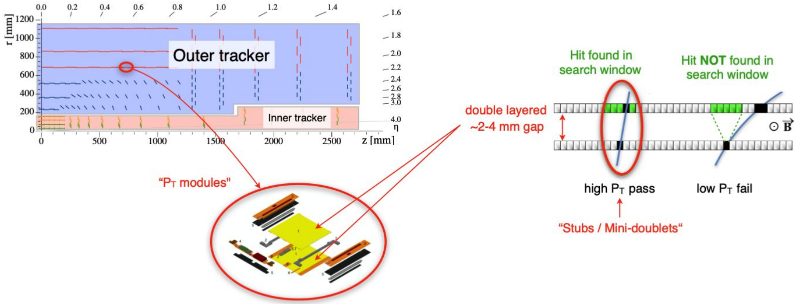

Physics Algorithms - Coordination Activities

Projects are currently at different stages in the software lifecycle. An intentional approach is necessary to effectively transition early stage R&D to integration and deployment

- c.f. ATLAS CDR/roadmap and CMS strategic plan

Integrated teams critical to successfully research, prototype and integrate high-quality and sustainable physics algorithms for HL-LHC experiments

- **GPU expertise becoming critical and must be embedded in research teams**
- **Portability of code will be a key issue**

Coordination between R&D projects and between R&D teams and experiments is critical especially as projects reach the point of selection and integration by experiments

- Ensure proper metrics are established to make decisions
- Decision making for R&D directions that can be stopped or reinforced

Workforce development is key for these efforts.

Workforce Development

Multiple programs underway aiming at different career stages:

- Summer programs: US-CMS PURSUE program, US-ATLAS SUPER program, IRIS-HEP Fellows program
- IRIS-HEP/HSF Training activities (materials and events) - 1600 students and 50 educators in the past few years
- Summer schools: IAIFI, CoDaS-HEP
- US-ATLAS and US-CMS postdoc R&D programs
- Upcoming: DOE CompHEP Traineeship awards/projects
- International: Fellows/mentoring in NSF-funded HSF-India project, Google summer of code, CERN summer students, etc.

Workforce Development Coordination Opportunities

- We have all of the elements of a **training pipeline** from undergraduate through postdocs, including opportunities for "cohort building"
- That pipeline can be connected to the R&D projects in **integrated teams** with computing professionals (e.g. GPU expertise)
- Significant opportunities to connect to growing **University "Research Software Engineer" groups** with relevant expertise (often from beyond HEP) - already leveraged in some of our R&D projects
  - This also provides career opportunities for the students/postdocs
- We have a growing **central repository of shared training curriculum materials** (HSF Training area) - basic material is now well covered, and several projects will be developing more advanced material (IRIS-HEP/HSF, DOE CompHEP Traineeships)
- The workshop also uncovered a strong commitment between the groups to coordinate on DEI questions (e.g. managing Code of Conduct violations), broadening participation, mentoring of mentors, effective training and evaluation

Gaps & Opportunities

Several items (not comprehensive) in the prior focus areas are not covered by any project:

- More **coordination effort for DC23** is needed to successfully include all the desired projects.
  - Data cache infrastructure and policy
  - Databases and software deployment
  - AI training and inference workflows
- Translational AI / **AI coordination is all new compared to the 2019 workshop**. No clear ownership at this point.
- Exploratory AI R&D, and others e.g. data services at analysis facilities, often done outside of experiments. Leaves significant **gap to integrate AI R&D** into the experiment's production / ops.
- Continued **integration of HPC into the LHC** computing environment for production and "data reduction" provides opportunities for collaboration with ASCR and LCF, as well as including HPC capacities into the resource planning for the HL-LHC

IRIS-HEP and HEP-CCE are both executing strategic planning exercises which incorporate these opportunities as part of future projects.

Conclusions

In the 2019 workshop, we spent significant effort defining how the projects in the field interacted.

- In 2022, these interactions are now well-established, effective, and delivering value toward HL-LHC.
- This round, we were able to dig deeper into the actual project coordination - highlighting a few areas the group felt important.

Plan is to integrate all these inputs into a more formal closeout report, targeted for December 2022

Thank you everyone for your participation!

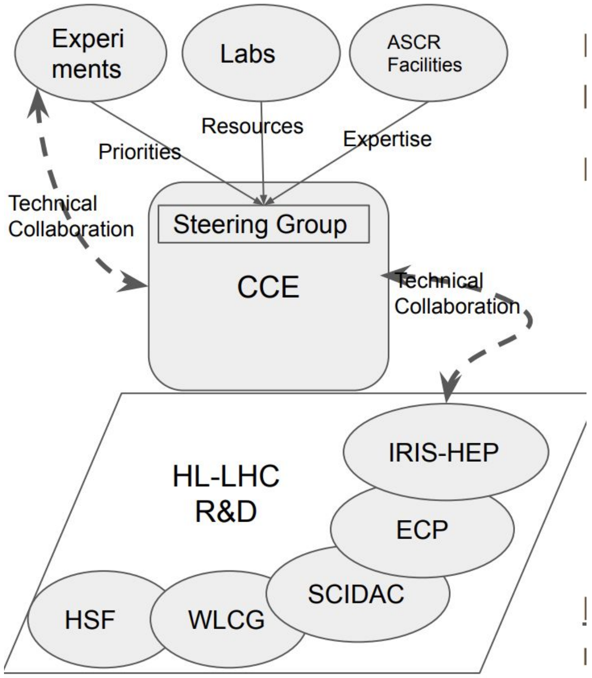

Figure reproduced from the 2019 closeout slides.

Looking forward

2017

2019

2022

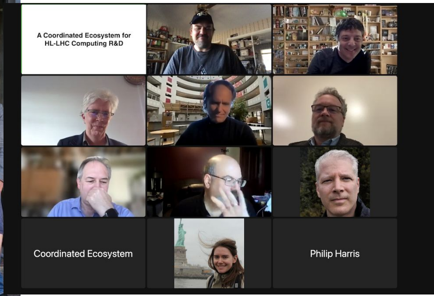

????

Thanks!