



Poster
#40

Institute for Research and Innovation in Software for High Energy Physics (IRIS-HEP)

PI: Peter Elmer (Princeton), co-PIs: Brian Bockelman (Morgridge Institute), Gordon Watts (U.Washington) with UC-Berkeley, University of Chicago, University of Cincinnati, Cornell University, Indiana University, MIT, U.Michigan-Ann Arbor, U.Nebraska-Lincoln, New York University, Stanford University, UC-Santa Cruz, UC-San Diego, U.Illinois at Urbana-Champaign, U.Puerto Rico-Mayaguez and U.Wisconsin-Madison

<http://iris-hep.org>



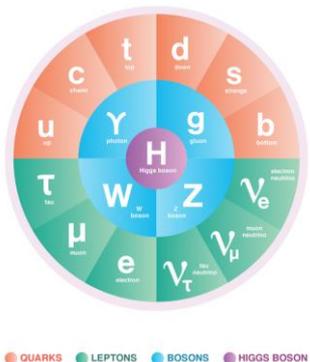
OAC-1836650

IRIS-HEP was funded as of 1 September, 2018



CSSI Meeting, Feb 14, 2020

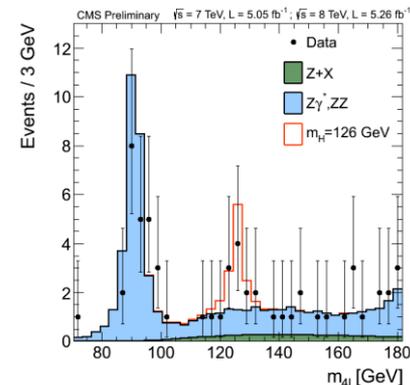
Science Driver: Discoveries beyond the Standard Model of Particle Physics



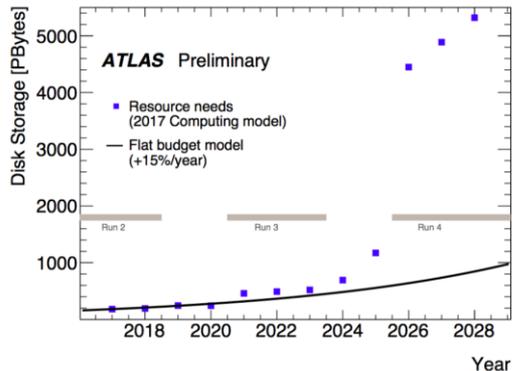
From “Building for Discovery - Strategic Plan for U.S. Particle Physics in the Global Context” - Report of the Particle Physics Project Prioritization Panel (P5):

- 1) Use the Higgs boson as a new tool for discovery
- 2) Pursue the physics associated with neutrino mass
- 3) Identify the new physics of dark matter
- 4) Understand cosmic acceleration: dark matter and inflation
- 5) Explore the unknown: new particles, interactions, and physical principles

Seattle
Snowmass 2021

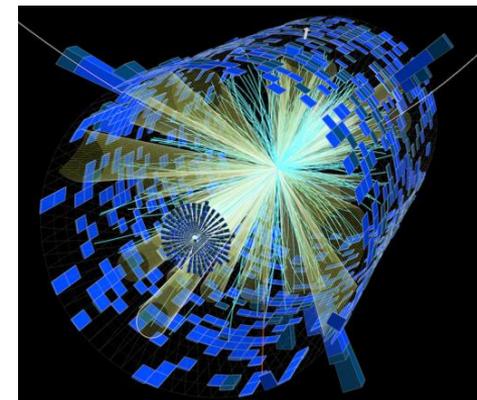


Computational and Data Science Challenges of the High Luminosity Large Hadron Collider (HL-LHC) and other HEP experiments in the 2020s



The HL-LHC will produce exabytes of science data per year, with increased complexity: an average of 200 overlapping proton-proton collisions per event.

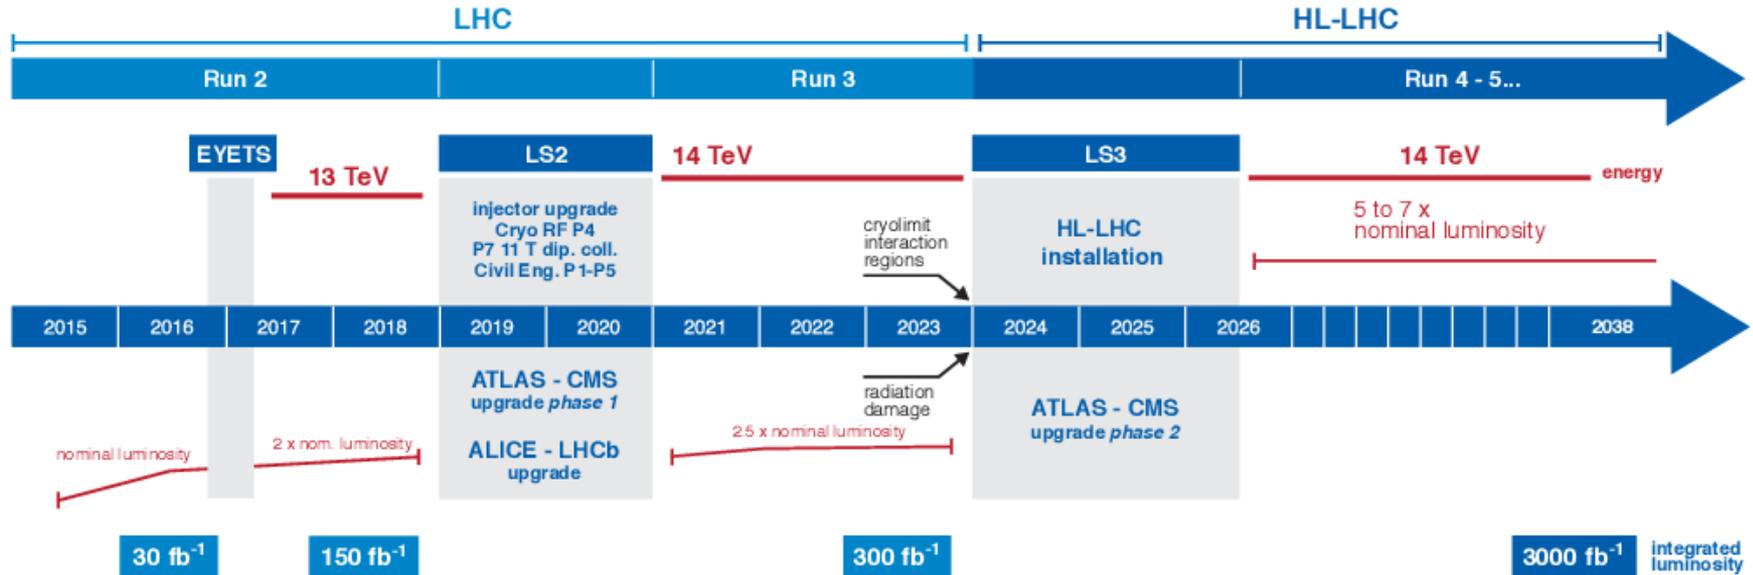
During the HL-LHC era, the ATLAS and CMS experiments will record ~10 times as much data from ~100 times as many collisions as were used to discover the Higgs boson (and at twice the energy).



Timeline



LHC / HL-LHC Plan



Institute Conceptualization and Community White Paper Process

S2I2-HEP



Snowmass

CTDR

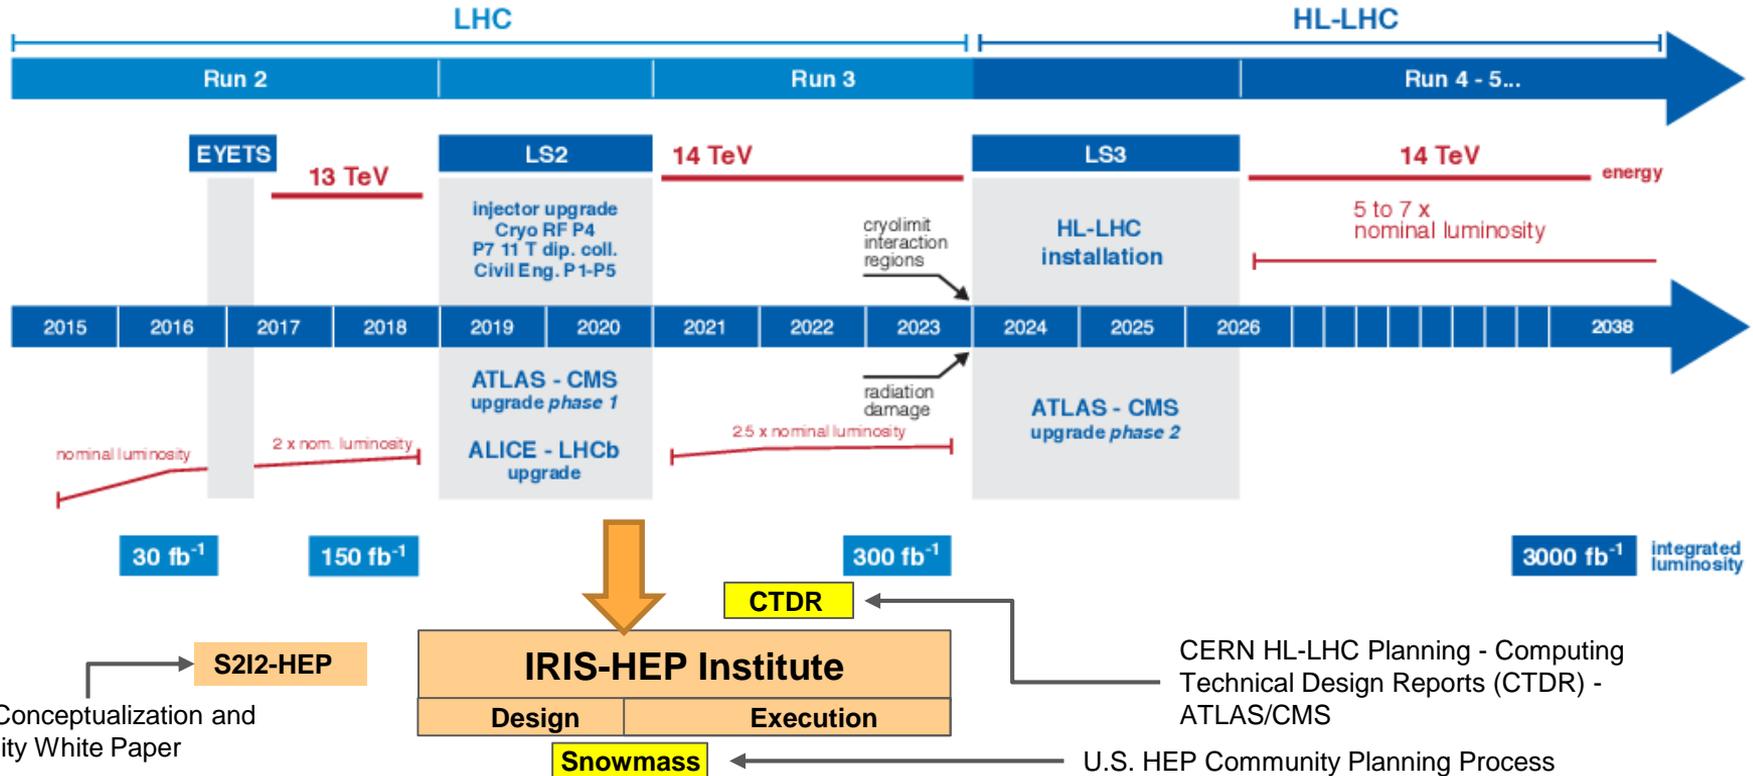
CERN HL-LHC Planning - Computing Technical Design Reports (CTDR) - ATLAS/CMS

U.S. HEP Community Planning Process

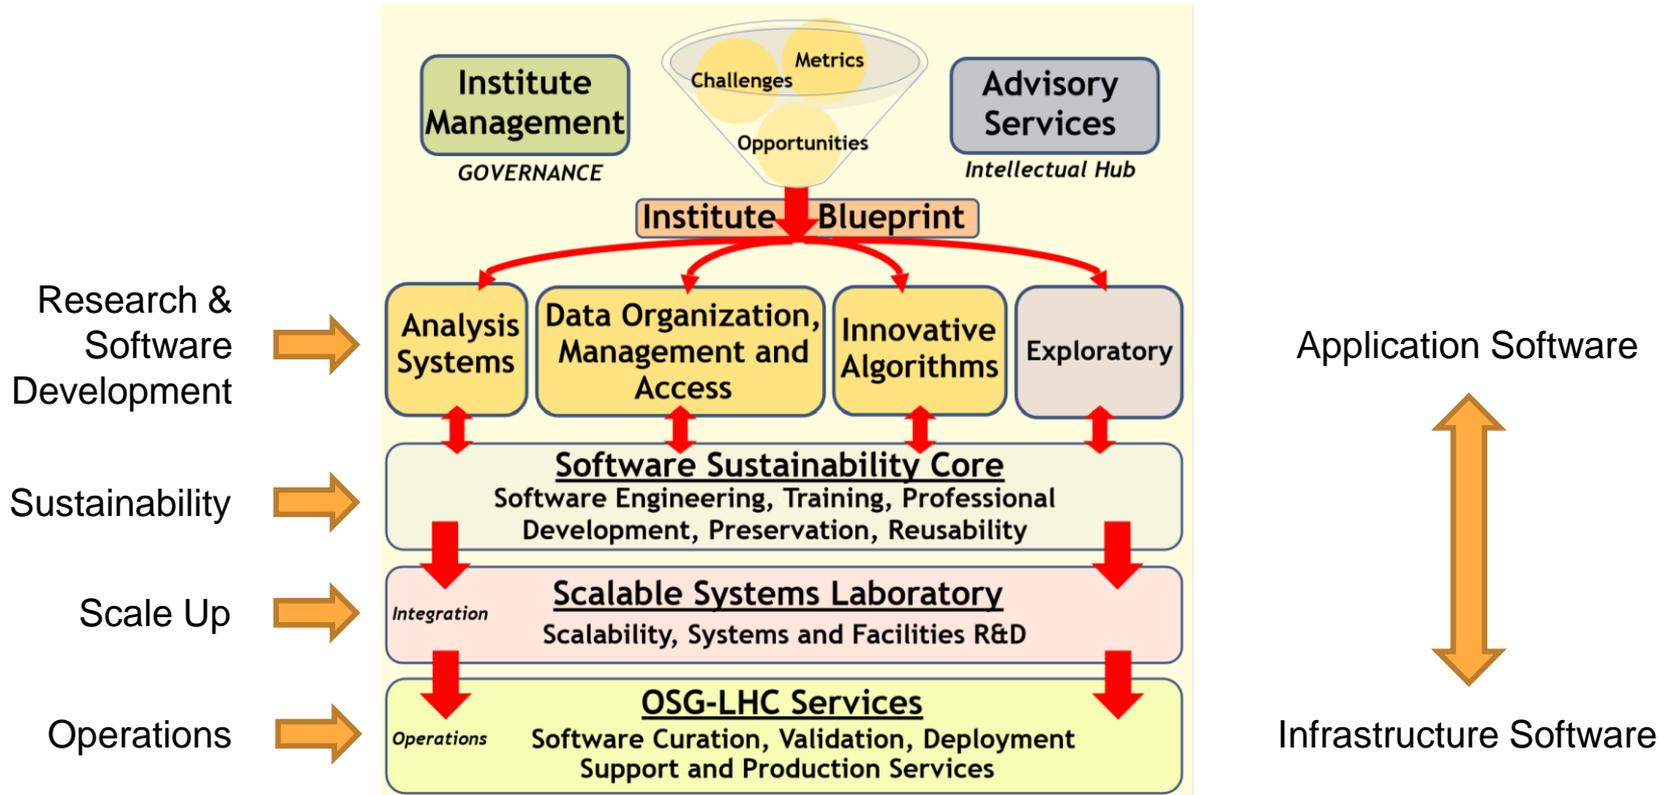
Timeline



LHC / HL-LHC Plan



Structure And Focus Areas



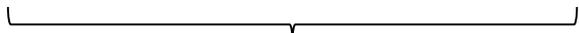
Our Audience: LHC Physicists and LHC Facility Operations Groups

Analysis Systems

Develop sustainable analysis tools to extend the physics reach of the HL-LHC experiments.

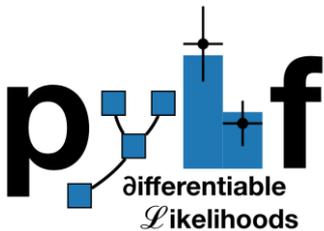
- create greater functionality to enable new techniques,
- reducing time-to-insight and physics,
- lowering the barriers for smaller teams, and
- streamlining analysis preservation, reproducibility, and reuse.

Experiment's Production System



Data Query, histogramming, plotting, statistical models, fitting, archiving, reproducibility, publication

All software is open source



Statistical Modeling Language and Tool

→ Limit Extraction

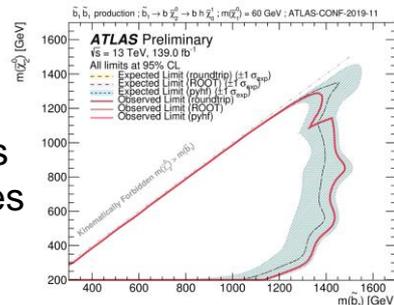
Rewritten from C++ in Python to use TensorFlow or PyTorch as back end.

GPU acceleration comes for "free"



C++: 10+ hours
pyhf: 30 minutes

Just released and being incorporated into Analyses Now



Built into SciKit-HEP, a suite of packages that are being adopted by the community



Salvatore Rappoccio @srrappoc... 13h

Replying to @claranellist

Successfully reduced a workflow that used to take O(2 days) to a jupyter notebook that takes O(2 minutes).



3 2 5 ...



Salvatore Rappoccio

@srrappoccio

Replying to @srrappoccio @claranellist

This improvement brought to you by @iris_hep software.

"Processed 18455107 events in
75.30 s = 245074.18 Hz"

Our previous workflow topped out at 1000 Hz.

uproot
awkward array } DIANAHEP
coffee } And IRIS-HEP



DOMA (Data Organization, Management, Access)

Fundamental R&D related to the central challenges of organizing, managing, and providing access to exabytes of data from process systems of various kinds.

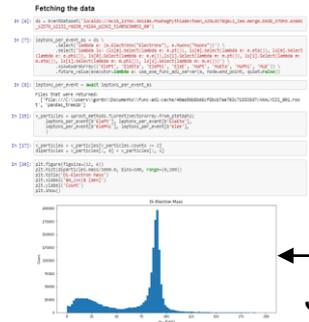
- Data Organization: Improve how HEP data is serialized and stored.
- Data Access: Develop capabilities to deliver filtered and transformed event streams to users and analysis systems.
- Data Management: Improve and deploy distributed storage infrastructure spanning multiple physical sites. Improve inter-site transfer protocols and authorization.



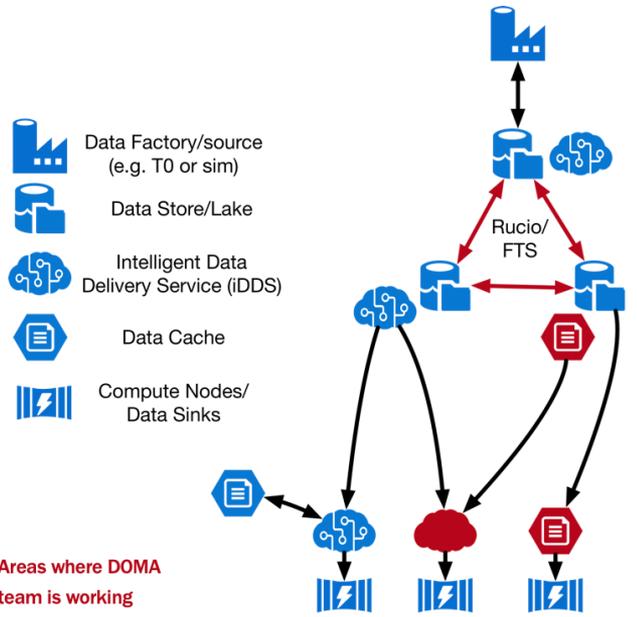
ServiceX / Intelligent Data Delivery

Low-latency delivery of numpy-friendly data transformed from experiment custom formats enabling the use of community supported data science tools.

(joint effort with Analysis Systems)



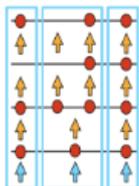
Jupyter Notebook



Innovative Algorithms – Trigger & Reconstruction

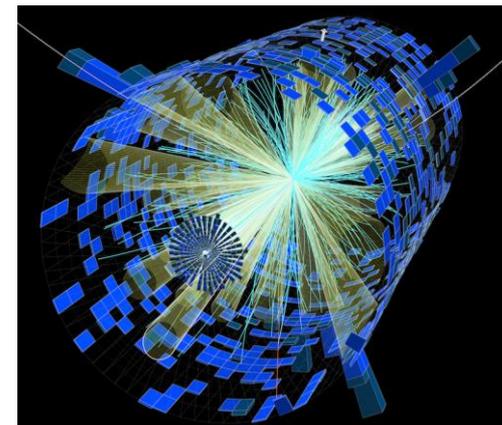
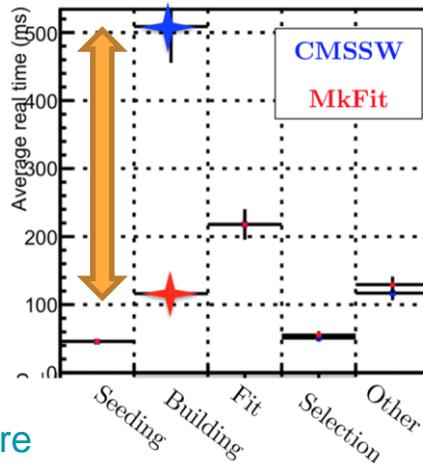
Algorithms for real-time processing of detector data in the software trigger and offline reconstruction are critical components of HEP's computing challenge.

- How to redesign tracking algorithms for HL-LHC?
- How to make use of major advances in machine learning (ML)?



mkFit – Parallel Track Fitting

- Develop track finding/fitting implementations that work efficiently on many-core architectures (vectorized and parallelized algorithms):
- 4x faster track building w/ similar physics performance in realistic benchmark comparisons



Pileup in the HL-LHC will increase combinatorics dramatically

Now being integrated into CMS production software

Will supply tracking enhancements for ~3500 physicists

~300 have attended various small trainings we've run or sponsored



Software Sustainability Core

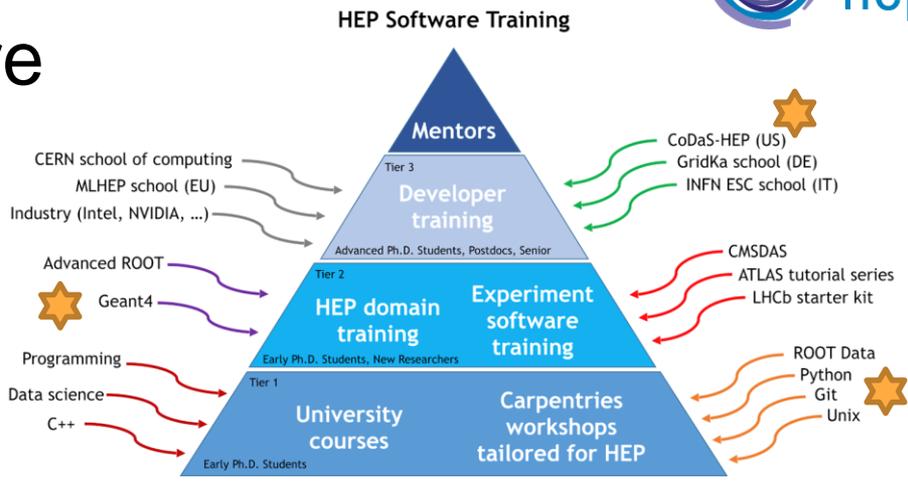
Training



Sample Topics: Git, OpenMP, SciPy, ML, Random Numbers, Columnar Data Analysis, Vectorization, etc.

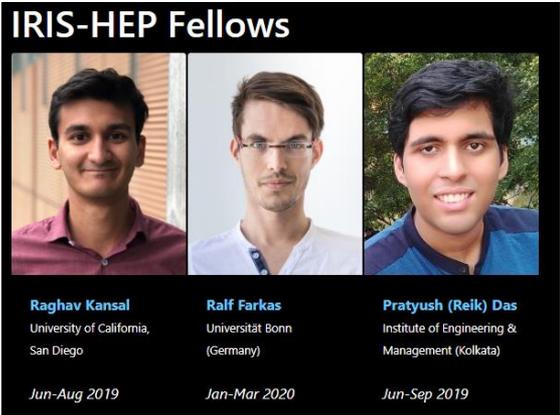
Direct Value to IRIS-HEP

We've had previous students become teachers, and previous students are now team-members in IRIS-HEP. Not just value to the community!



Fellows Program

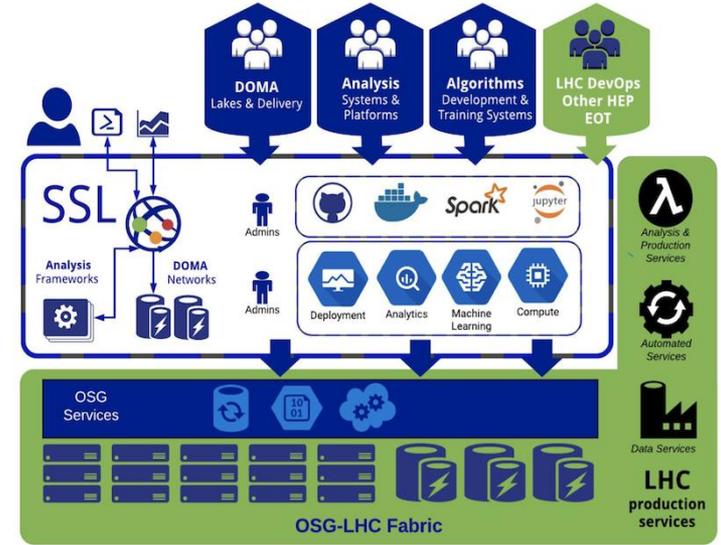
Provides opportunities for undergraduate and graduate students to connect with mentors within the larger HEP and Computational/Data Science community.



Scalable Systems Laboratory

Goal: Provide the Institute and the HL-LHC experiments with scalable platforms needed for development in context. Facilities R&D

River – a repurposed UChicago CS research cluster now being used to test/run IRIS-HEP projects.



CoDaS-HEP school environment, ServiceX test bed.

Kubernetes based cluster, can run the OSG-LHC environment, school environments, etc. Experimenting with “no-ops” management.

Collaborating with a CyberTraining project (OAC-1829707, 1829729) as well as a growing number of international collaborators.

Open Science Grid - LHC

The OSG is a consortium dedicated to the advancement of all of open science via the practice of Distributed High Throughput Computing, and the advancement of its state of the art.

- IRIS-HEP supports LHC operations and development of the consortium.



Open Science Grid



- Work to separate local site hardware and software support by moving services into containers.
- Transitioning security service to use tokens

Particle physicists all over the world depend on these services and scheduling of processing hours (~10,000)

Some (biased) Impact Highlights



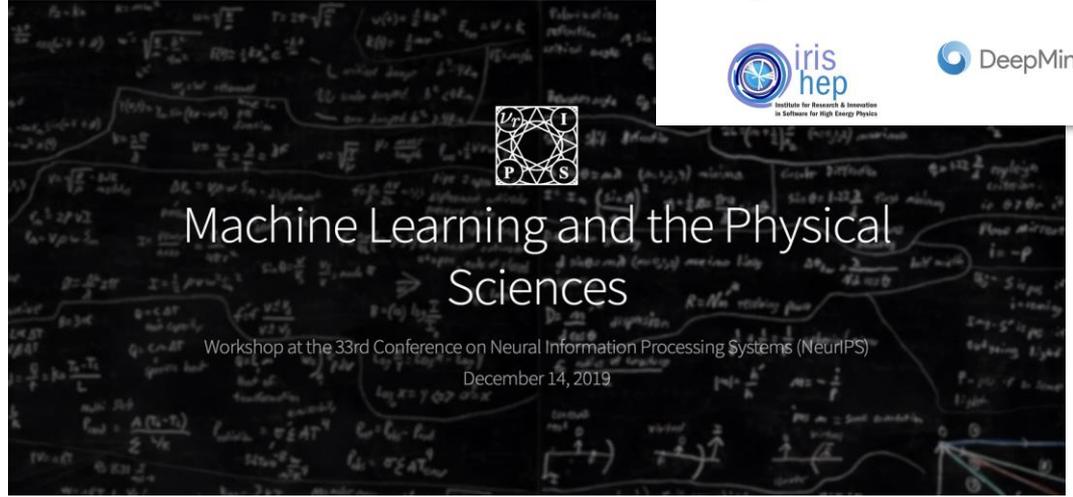
nature physics PERSPECTIVE
<https://doi.org/10.1038/s41567-018-0342-2>
 Corrected: Publisher Correction OPEN

Open is not enough

Xiaoli Chen^{1,2}, Sünje Dallmeier-Tiessen^{1*}, Robin Dasler^{1,3}, Sebastian Feger^{1,3}, Pamfilos Fokianos⁴, Jose Benito Gonzalez², Harri Hirvonsalo^{1,4,5}, Dinos Kousidis⁶, Artemis Lavasa⁷, Salvatore Mele⁸, Diego Rodriguez Rodriguez², Tibor Šimko^{1*}, Tim Smith¹, Ana Trisovic^{1,8}, Anna Trzcinska¹, Ioannis Tsanaktsidis¹, Markus Zimmermann¹, Kyle Cranmer⁶, Lukas Heinrich⁶, Gordon Watts¹, Michael Hildreth⁹, Lara Lloret Iglesias⁹, Kati Lassila-Perini¹ and Sebastian Neubert¹⁰

The solutions adopted by the high-energy physics community to foster reproducible research are examples of best practices that could be embraced more widely. This first experience suggests that reproducibility requires going beyond openness.

Sponsors

@NeurIPS

Co-Sponsored: interest in ML in physics and the sciences is very high in the global community.

For a Global Field



Global community is ~O(30K)

Community Building

IRIS-HEP came out of the S212-HEP: Conceptualization Process

This was a community building exercise:

- 17 workshops from 2016-2017
- More than 20 papers of ideas submitted to the physics archive
- Roadmap published in “Computing and Software for Big Science”

Part of IRIS-HEP’s mandate is to continue this process

- Blueprint meetings to build field-wide consensus on specific problems.
- The Fellows Program
- Topical Meetings: seminars on topics of interest.
- Sponsorship of conferences and workshops like PyHEP 2020, and LAWSCHEP 2019.

~900 have attended various small workshops we’ve run or sponsored



“The result: a Programme of Work for the field as a whole, a multifaceted approach to addressing growing computing needs on the basis of existing or emerging hardware.” – Eckhard Eisen (CERN Director of Research and Computing), editorial published with Roadmap



Summary



- IRIS-HEP was funded on September 1st, 2018
 - We are approaching the end of the design phase
 - Projects in all phases (design, prototype, and production) exist.
 - We are fully staffed, ~30 FTE's
 - Full description of projects available on our website, <http://iris-hep.org>
- Community Impact
 - Software is being adopted by others, in some cases dramatically.
 - Facilities work in SSL and OSG is leading the international field
- Community Outreach
 - We've reached almost 1000 people with our workshops, and another 300 with our training efforts
 - We continue to organize Blueprint workshops to build community consensus.
- Next
 - Start Execution Phase September 2020
 - Work on integrating projects in prototype stage into coherent and scalable software for the community
 - The "Snowmass Process- 2021" provides an opportunity for us to update the Community White Paper/Roadmap.