PROCEEDINGS OF SCIENCE

HL-LHC and Beyond Computing Challenges

Evangelos Kourlitis[a,b,*]

[a] *Technical University of Munich*
Boltzmannstr. 2, D-85748, Garching, Germany
[b] *The HEP Software Foundation*

E-mail: evangelos.kourlitis@cern.ch

In anticipation of the High-Luminosity LHC (HL-LHC) era, the High Energy Physics (HEP) community faces considerable challenges in the evolving landscape of both experimental and computational domains. This paper explores the coordinated efforts led by the HEP Software Foundation (HSF) to address common challenges in software and computing. The HSF facilitates international collaboration and coordinates research and development through working groups. As the HL-LHC poses unprecedented data volumes and computational demands, the community adapts with advancements in hardware utilization, novel simulation techniques, and the adoption of array programming models. Additionally, the paper highlights the crucial role of the HSF in establishing a unified software training framework to meet the evolving needs of the HEP community.

The Eleventh Annual Conference on Large Hadron Collider Physics (LHCP2023)
22-26 May 2023
Belgrade, Serbia

*Speaker

https://pos.sissa.it/

PoS(LHCP2023)112

1. HL-LHC and Beyond Computing Landscape

The High-Luminosity Large Hadron Collider (HL-LHC) has a confirmed operational timeline extending until the early 2040s, with the target to collect 3000 fb^{-1} of data [1]. Concurrently, investigations into potential future colliders, to be situated in either Europe or Asia, are underway [2]. Presently, substantial investments have been made by High Energy Physics (HEP) experiments in software and computing infrastructure. Notably, the ATLAS experiment at CERN currently has more than 750 PB of data stored, while recently broke the record of over 1 million CPU cores simultaneously contributing to data processing over a global network of 120 data centers, 5 High-Performance Computing (HPC) and 3 cloud facilities [3]. The expected computing needs at the HL-LHC era, and beyond, are only projected to increase and necessitate significant research and development efforts to maintain within budgetary constraints [4].

Simultaneously, computing technology has undergone a notable evolution. From the mid-2000s onwards, a transformative shift in microprocessor performance delivery has occurred, characterized by a stalling of clock frequency coupled with an increase in transistor density. This paradigm shift necessitates a concurrent software programming model. Notably, non-x86 CPU architectures, such as ARM, have gained substantial market share due to their enhanced energy efficiency, demonstrated in HEP workloads as well [5]. Furthermore, contemporary HPC systems predominantly employ hardware acceleration (e.g. GPUs) to parallelize computations effectively, resulting in a significant reduction in energy consumption for software written or adapted for such devices.

The parallel evolution of the experimental and computing landscape reveal imminent challenges. The unprecedented data volume, the increasingly complex workflows and data analysis model along with the increasingly heterogeneous computing resources mandate a comprehensive re-engineering of scientific software. These challenges span in various domains and are detailed below.

1.1 The HEP Software Foundation

The HEP Software Foundation (HSF) [6] plays a pivotal role in addressing the future computing challenges in HEP. Established in 2014, HSF endeavors to foster collaboration and synergy within the HEP community, encouraging efficient and effective solutions to shared computing challenges. Recognizing the universal nature of these challenges across different experiments, HSF serves as a catalyst for international coordination, uniting researchers and institutions in a collective effort. By establishing various working groups, HSF facilitates the collaboration of experts in the field, pooling their expertise to confront common issues. A significant outcome of these collective endeavors is the development of a comprehensive roadmap not only to reflect the collaborative spirit within HSF but also to serve as a guiding document for setting goals and priorities to navigate the future of HEP computing [7].

2. Energy Frontier Computing Challenges

2.1 Triggering

The HL-LHC will deliver a more than threefold increase in instantaneous luminosity, consequently increasing pileup by an order of magnitude. Presently, hardware triggers reduce the data

PoS(LHCP2023)112

rate to $O(0.1-1)$ MHz for the ATLAS and CMS experiments, while heterogeneous computing farms reconstruct and trigger events, further reducing the data rate to $O(1)$ kHz. Ongoing research aims to increase utilization of hardware accelerators, such as GPUs and FPGAs, to cope with the increasingly busy experimental environment with reduced power consumption. In the case of the CMS experiment, GPU integration has notably enhanced throughput by 80%, accompanied by a 30% reduction in power consumption [8]. Going a step further, the LHCb and ALICE experiments pursue a strategy that completely removes the hardware triggers and processes events at collision rate directly in computing farms based on hardware accelerators [9, 10].

2.2 Event Generation

Event generation is anticipated to utilize 10-20% of the HL-LHC computing resources [4]. In this area, the HSF has played a pivotal role in focusing research and development efforts through the Physics Event Generator Working Group. The group focused on understanding current CPU bottlenecks and finding optimal ways to transition to GPUs and vectorized implementations. Notably, matrix element computations have been successfully demonstrated on both GPUs and SIMD CPUs, showcasing significant speedup, with recent benchmarks indicating up to an 80-fold acceleration on an NVidia A100 GPU compared to single-threaded CPU execution for generating pairs of top quarks with additional gluon emission [11]. Furthermore, Machine Learning (ML) algorithms are improving the phase space sampling and integration [12] and novel theoretical approaches are effectively reducing the computational cost associated with negative weight events [13].

2.3 Detector Simulation

Detector simulation currently stands as the most significant CPU consumer, and the demand is expected to increase in the HL-LHC era [4]. To address these computing needs, LHC experiments are exploring fast simulation techniques, particularly for calorimeter simulation. Traditional fast simulators rely on parametrization methods, but there has been a recent surge in those leveraging ML algorithms [14]. These ML-based approaches broadly fall into two categories:

1. replacing or enhancing the detailed Geant4 simulation with ML-based surrogates,
2. using ML techniques to refine fast and lower-quality simulations to align with the comprehensive Geant4 models.

Current findings suggest that a hybrid approach combining Geant4 and fast simulators will likely be the path forward for the HL-LHC.

Despite these advancements, the use of highly accurate simulations based on Geant4 remains essential. Experiments are actively optimizing the Geant4 simulations for their specific needs, achieving notable throughput improvements. For instance, the ATLAS experiment achieved a 37% increase in throughput via both physics and computational optimizations [15]. Additionally, efforts are underway to leverage GPU hardware for calorimeter simulations, with projects like AdePT [16] and Celeritas [17] leading the way. Initial studies for the CMS detector indicate a twofold increase in throughput by offloading the simulation of electromagnetic showers onto GPUs [18].

PoS(LHCP2023)112

2.4 Event Reconstruction

Traditionally, reconstruction software has been tailored to specific experiments. However, recent advancements in hardware motivate universal solutions to avoid duplicating work and to have a broad common research and development approach. This imperative led to the inception of A Common Tracking Software Project (ACTS), designed to create a set of experiment-independent tools for track reconstruction [19]. The ACTS project seeks to offer high-level modules with versatile and efficient implementations applicable to tracking detectors across various experiments. Key features include tracking geometry description, a simple event data model, and standardized algorithms for seeding, track fitting, and vertexing. Moreover, ACTS has served as a testing ground for geometric deep learning, yielding promising initial results. For instance, in the ATLAS experiment, demonstrations exhibited fake track rates at the per mille level [20].

2.5 Data Analysis

To analyze the data anticipated from the HL-LHC and other future experiments is a challenge itself, given that the current workflows do not scale to the expected data volume. The cost of disk and tape storage is a concern, with analysis data projected to occupy a substantial portion of these resources. Consequently, adoption of new strategies is imperative. Ongoing efforts are directed towards advanced data encoding, as ROOT's RNTuple [21] and other more compact data formats, aiming for an approximate target of 10 kb/event. This approach facilitates improved handling of data in local batch compute facilities. Noteworthy examples include the ATLAS PHYSLITE [22] and CMS NanoAOD [23] data formats.

At the same time, a paradigscalablem shift is unfolding in data analysis workflows, with a focus on enhancing computational efficiency and reducing intermediate data files. The adoption of an array programming model, known as *columnar analysis*, is underway, enabling concurrent computations across multiple objects or events and aligning HEP data analysis with contemporary data analytics standards [24]. Python libraries, such as `awkward` [25] and `coffea` [26], have been developed to assist users in developing their frameworks and scaling workflows across various compute fabrics.

3. Conclusions and Community Training

As the LHC program advances toward its final run and preparations for the HL-LHC program and beyond are undergoing, the computing landscape goes through a significant evolution. This concurrent evolution of the experimental and computing facets presents substantial challenges for the HEP community, in various areas of the software. Given the commonality of these challenges, a coordinated effort is imperative for a comprehensive software and computing upgrade, and the HSF plays a pivotal role in fostering scientific computing collaborations to address these issues.

A compelling necessity exists for a unified, scalable, and sustainable software training framework driven by the entire community. The HSF Training Working Group, in collaboration with IRIS-HEP [27], is actively cultivating a community around these principles. The establishment of a unified Training Center for HEP currently encompasses 21 training modules within The Carpentries framework [28], featuring various workshops and training courses held on an annual basis.

References

[1] "Longer term LHC schedule." https://lhc-commissioning.web.cern.ch/schedule/LHC-long-term.htm, 2022.

[2] F. Zimmermann, "Status of CERN's future collider projects." https://indico.cern.ch/event/1260648/, 2023.

[3] D. South, "R&D in ATLAS distributed computing towards HL-LHC." https://indico.jlab.org/event/459/contributions/11502/, 2023.

[4] The ATLAS Collaboration, "ATLAS Software and Computing HL-LHC Roadmap," Tech. Rep. CERN-LHCC-2022-005, LHCC-G-182, CERN, Geneva, 2022.

[5] D. Britton, "Data centres perspective on ARM adoption." https://indico.cern.ch/event/1230126/contributions/5315426/, 2023.

[6] "The HEP Software Foundation website." https://hepsoftwarefoundation.org/.

[7] J. Albrecht *et al.*, "A Roadmap for HEP Software and Computing R&D for the 2020s," *Comput. Softw. Big Sci.*, vol. 3, no. 1, p. 7, 2019.

[8] G. Parida, "Run-3 commissioning of CMS online HLT reconstruction using GPUs." https://indico.jlab.org/event/459/contributions/11822/, 2023.

[9] F. Lazzari, "Demonstration of track reconstruction with FPGAs on live data at LHCb." https://indico.jlab.org/event/459/contributions/11377/, 2023.

[10] G. Eulisse and D. Rohr, "The O2 software framework and GPU usage in ALICE online and offline reconstruction in Run 3." https://indico.jlab.org/event/459/contributions/12432/, 2023.

[11] S. Hageboeck, "Madgraph5_aMC@NLO on GPUs and vector CPUs: experience with the first alpha release." https://indico.jlab.org/event/459/contributions/11829/, 2023.

[12] L. Heinrich and M. Kagan, "Differentiable Matrix Elements with MadJax," *J. Phys. Conf. Ser.*, vol. 2438, no. 1, p. 012137, 2023.

[13] The ATLAS Collaboration, G. Aad, *et al.*, "Modelling and computational improvements to the simulation of single vector-boson plus jet processes for the ATLAS experiment," *JHEP*, vol. 08, p. 089, 2022.

[14] A. Adelmann *et al.*, "New directions for surrogate models and differentiable programming for High Energy Physics detector simulation," in *Snowmass 2021*, 3 2022.

[15] V. Kourlitis *et al.*, "Optimizing the ATLAS Geant4 detector simulation software," Tech. Rep. ATL-SOFT-PROC-2023-002, CERN, Geneva, 2023.

[16] G. Amadio *et al.*, "Offloading electromagnetic shower transport to GPUs," *J. Phys. Conf. Ser.*, vol. 2438, no. 1, p. 012055, 2023.

[17] S. R. Johnson *et al.*, "Celeritas is a new Monte Carlo transport code designed for high-performance simulation of high-energy physics detectors.." [Computer Software] https://doi.org/10.11578/dc.20221011.1, jul 2022.

[18] A. Gheata, "Accelerated demonstrator of electromagnetic Particle Transport (AdePT) status and plans." https://indico.jlab.org/event/459/contributions/11427/, 2023.

[19] X. Ai *et al.*, "A Common Tracking Software Project," *Comput. Softw. Big Sci.*, vol. 6, no. 1, p. 8, 2022.

[20] C. Rougier, "ATLAS ITk Track Reconstruction with a GNN-based Pipeline." https://indico.cern.ch/event/1103637/contributions/4821831/, 2022.

[21] J. Blomer *et al.*, "Evolution of the ROOT Tree I/O," *EPJ Web Conf.*, vol. 245, p. 02030, 2020.

[22] J. Schaarschmidt *et al.*, "PHYSLITE - A new reduced common data format for ATLAS," Tech. Rep. ATL-SOFT-PROC-2023-029, CERN, Geneva, 2023.

[23] M. Peruzzi *et al.*, "The NanoAOD event data format in CMS," *Journal of Physics: Conference Series*, vol. 1525, p. 012038, apr 2020.

[24] N. Hartmann *et al.*, "Columnar data analysis with ATLAS analysis formats," *EPJ Web Conf.*, vol. 251, p. 03001, 2021.

[25] J. Pivarski *et al.*, "Awkward Array." 10.5281/zenodo.4341376, 2018.

[26] N. Smith *et al.*, "Coffea columnar object framework for effective analysis," *EPJ Web Conf.*, vol. 245, p. 06012, 2020.

[27] "IRIS-HEP website." http://iris-hep.org.

[28] "The Carpentries website." https://carpentries.org/.