PROCEEDINGS OF SCIENCE

Benchmarking distributed-interactive HEP analysis workflows on the new Italian National Centre analysis infrastructure

Adelina D'Onofrio,[a,*] Muhammad Numan Anwar,[b,a] Matteo Bartolini,[c,a] Antimo Cagnotta,[d,a] Tommaso Diotalevi,[e,a] Francesco Giuseppe Gravili,[f,a] Salvatore Loffredo,[d,a] Paolo Mastrandrea,[g,a] Elvira Rossi,[d,a] Gianluca Sabella,[d,a] Federica Maria Simone,[b,a] Bernardino Spisso,[a] Alessandro Tarasio[h,a] and Tommaso Tedeschi[a]

[a] *Istituto Nazionale di Fisica Nucleare (INFN), Italy*
[b] *Politecnico di Bari, Italy*
[c] *Universitá di Firenze, Italy*
[d] *Universitá degli Studi di Napoli Federico II, Italy*
[e] *Universitá di Bologna, Italy*
[f] *Universitá del Salento, Italy*
[g] *Universitá di Pisa, Italy*
[h] *Universitá della Calabria, Italy*

E-mail: onofrioa@na.infn.it

Computationally efficient data processing and the challenges of future colliders are pushing to re-think High Energy Physics (HEP) computing models. The work behind this contribution aims at providing transparent resources for users and experiments, with suitable tools and environment, coupled with flexible and cloud-independent deployment in the framework of the ICSC project (High-Performance Computing, Big Data e Quantum Computing Research Centre). The resources will be experiment-agnostic and applicable across HEP experiments, exploited to benchmark the proposed workflows. Seamless interactive or quasi-interactive analysis are extremely promising: starting from container technology and Kubernetes, analysis tools are provided via Jupyter interface and Dask scheduling system, masking complexity for front end users and rendering cloud resources flexibly. An overview of the technologies involved and the results of two benchmark use cases will be provided, with suitable metrics to evaluate preliminary performance of the workflow.

42nd International Conference on High Energy Physics (ICHEP2024)
18-24 July 2024
Prague, Czech Republic

*Speaker

PoS(ICHEP2024)1043

1. Introduction

In the next future, the scale of the datasets involved in HEP data analysis will increase from Petabytes to Exabytes and next-generation techniques should play a key role to enhance the scientific reach of each result. Dealing with these challenges impacts both the software and the underlying computing infrastructure. It is fundamental to develop new industry-standard analysis paradigms, based on declarative programming solutions and on interactive workflows. According to the DataLake model, a new infrastructure is envisioned: using distributed and heterogeneous resources, distributed computing elements will be able to run workflows from the HEP community. Funded by the Italian National Recovery and Resilience Plan (PNRR), ICSC will provide such resources [1], allowing the implementation of a High Rate Analysis platform on a national scale. The working environment is highly customisable using container-based technologies, while the entire platform is deployed on a Kubernetes cluster [2], allowing a seamless, flexible, scalable and fault-tolerant deployment on the available resources with a limited impact on the admin's work time. In this contribution, a description of some use cases will be provided, showing the advantages of the proposed approach in physics data analyses.

2. HEP analyses performance evaluation

Use cases from two HEP experiments were considered to evaluate the platform's performance.

2.1 Analysis of simulated $Z \to e^+e^-$ decays at Future Colliders (FCCee)

In the context of the Future electron-positron Circular Collider (FCCee [3]), preliminary feasibility studies have being conducted using simulated events, leveraging the High Rate platform. $Z \to e^+e^-$ simulations were studied to test the scalability of the computing resources exploited, using the IDEA [4] detector concept. The core concept is to utilise a standard candle process in HEP, reconstruct the di-electron system invariant mass and mimic systematic variations modifying the decay channel electrons kinematic, applying a Gaussian smearing to their energies. Events selection and histogramming were performed interactively, with ROOT RDataFrame [5] and JupyterLab [6], and Dask [7] was used as back-end. Preliminary results, obtained running the same analysis workflow and considering the overall execution time as a metric (Fig. 1a), highlight a performance improvement in the distributed approach with respect to the standard/serial approach.

2.2 The ATLAS use case workflow

Another use case focuses on the search of direct pair production of top s-quarks and dark matter particles, using $139\,\text{fb}^{-1}$ of integrated luminosity collected by the ATLAS detector at the CERN Large Hadron Collider (LHC) [8]. The analysis workflow starts from standard ATLAS raw data formats, for both data and Monte Carlo simulations, that undergo several reprocessing and event reconstruction steps up to an analysis object data format available to the users. Further custom-based selections, based on a simple cut&count approach, are applied to retain only the physically interesting events. This step was the core of the feasibility studies: RDataFrame and Dask were applied to a background sample, corresponding to O(GB) of Monte Carlo simulations. According to the analysis procedures, several systematic variations were applied to test the code's overall

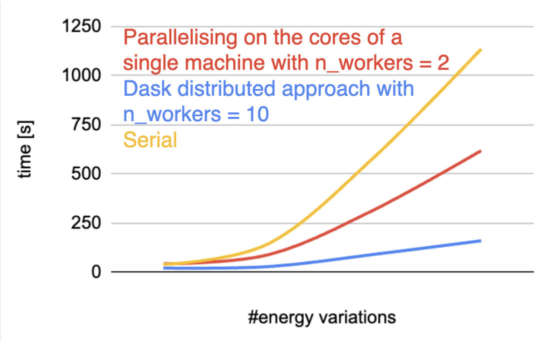

(a) $Z \to e^+e^-$ simple use case. The execution time as function of the number of Dask workers is represented.

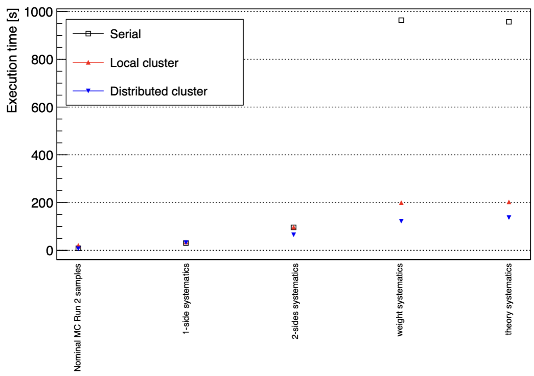

(b) ATLAS analysis. The execution time as function of the increasing number of systematic variations is represented.

Figure 1: Performance evaluation of two different use cases.

execution time over hundreds of iterations of the event selection step. Exploiting the distributed approach the execution time improved by a factor of 5 with respect to the standard/serial approach, iterating over a significant number of systematic variations (see Fig. 1b). At the same time, the amount of output files is reduced to O(MB).

3. Conclusions

An innovative approach has been presented to cope with the expectedly big amount of data to be processed in the upcoming LHC phase and at Future Colliders: an interactive, high throughput platform based on a parallel and distributed computing system, leveraging the ICSC DataLake model. Thanks to the new available tools and to a better handling and management of the computing resources, preliminary results on some benchmark HEP analyses show promising improvements in terms of the total execution time.

Acknowledgements

This work is supported by ICSC - Centro Nazionale di Ricerca in High Performance Computing, Big Data and Quantum Computing, funded by European Union - NextGenerationEU.

References

[1] ICSC Main page. Available at: https://www.supercomputing-icsc.it/.

[2] Kubernetes Main page. Available at: https://kubernetes.io/.

[3] FCC Main Page. Available at: https://fcc.web.cern.ch/Pages/default.aspx.

[4] M. Antonello, IL NUOVO CIMENTO, vol. 43C, 2020. doi:10.1393/ncc/i2020-20027-2.

[5] ROOT RDataFrame framework Main Page. Available at: https://root.cern/doc/master/classROOT_1_1RDataFrame.html.

[6] JupyterLab Main page. Available at: https://jupyter.org/lab.

[7] Dask Development Team, "Dask: Library for dynamic task scheduling", 2016. Available at: http://dask.pydata.org.

[8] The ATLAS Collaboration, *J. High Energ. Phys.* **2021**, 165 (2021). https://doi.org/10.1007/JHEP04(2021)165