

High-throughput data analysis with modern ROOT interfaces

Tuesday 24 November 2020 16:22 (10 minutes)

With the upcoming start of LHC Run III and beyond, HEP data analysis is facing a large increase in average input dataset sizes. At the same time, balancing analysis software complexity with the need to extract as much performance as possible from the latest HPC hardware is still often difficult.

Recent developments in ROOT significantly lower the energy barrier for the development of high-throughput data analysis applications. This was achieved through a unique combination of ingredients: a high-level and high-performance analysis framework; just-in-time compilation of C++ code for efficient I/O and usability enhancements; automatic generation of Python bindings; transparent offloading of computations to distributed computation engines such as Spark.

The resulting simplified data analysis model has enabled a whole range of R&D activities that are expected to deliver further acceleration, such as context-aware caching.

This talk will provide an overview of recent developments in ROOT as an engine for high-throughput data analysis and how it is employed in several existing real-world usecases.

Primary authors: Dr GUIRAUD, Enrico (EP-SFT, CERN); Dr NAUMANN, Axel (CERN); Mr PADULANO, Vincenzo Eduardo (Valencia Polytechnic University (ES)); Dr TEJEDOR SAAVEDRA, Enric (CERN); Mr WUNSCH, Stefan (KIT - Karlsruhe Institute of Technology (DE))

Presenters: Dr GUIRAUD, Enrico (EP-SFT, CERN); Mr PADULANO, Vincenzo Eduardo (Valencia Polytechnic University (ES)); Mr WUNSCH, Stefan (KIT - Karlsruhe Institute of Technology (DE))

Session Classification: Software