

Writing ROOT Data in Parallel with TBufferMerger

Tuesday, July 10, 2018 12:00 PM (15 minutes)

The LHC experiments produce tens of petabytes of new data in ROOT format per year that need to be processed and analysed. In the next decade, following the planned upgrades of the LHC and the detectors, this rate is expected to increase at least ten-fold.

Therefore, optimizing the ROOT I/O subsystem is of critical importance to the success of the LHC physics programme. This contribution presents ROOT's approach of writing data from multiple threads to the same file in a highly efficient way. Technical aspects of its implementation, the TBufferMerger class, as well as examples of its programming model are described. Runtime performance and writing rates measurements and associated improvements with respect to the single writer case are discussed in the context of standalone examples and production parallel data processing workflows such as CMS reconstruction.

The interplay of parallel dataset writing with other parallelisation mechanisms in ROOT is discussed together with the future developments of the strategy aiming to maximise experiments' data processing throughput

Authors: AMADIO, Guilherme (CERN); PIPARO, Danilo (CERN); CANAL, Philippe (Fermi National Accelerator Lab. (US)); GUIRAUD, Enrico (CERN, University of Oldenburg (DE))

Presenter: AMADIO, Guilherme (CERN)

Session Classification: T5 - Software development

Track Classification: Track 5 –Software development