



Contribution ID: 15

Type: **not specified**

Support Future Data Analysis with a Parallelised ROOT

Monday, 10 September 2018 11:00 (20 minutes)

Future accelerators and detectors pose to HEP scientific software a series of challenges, among which the efficient analysis of the data collected by the experiments. In the past few years ROOT became to a large extent parallel but our endeavour is not complete.

This presentation is dedicated to the characterisation of the parallelisation effort which took place up to now and to the lessons learned thanks to this effort. The units of the ROOT framework are discussed where task and data parallelism have been introduced, with runtime and scaling measurements. We will give an overview of concurrent operations in ROOT, for instance in the areas of I/O (reading and writing of data), fitting / minimization, and data analysis.

Extraordinary results have been achieved with the parallelisation of ROOT and more will come in the short term. On the other hand, assuming a flat funding for HEP computing and a reasonable evolution of hardware, in order to implement the research programme of HL-LHC, our software needs to be 10x faster than today. This includes data analysis (and therewith IO) and sophisticated procedures such as fits.

What does ROOT need to provide to support data analysis and processing in such a harsh environment?

In this talk we propose our vision about the ROOT's core components which will need most attention. Notable examples are the backend for IO, the interpreter and its ability to support heterogeneous platforms and ROOT's type system.

Primary author: PIPARO, Danilo (CERN)

Presenter: PIPARO, Danilo (CERN)

Session Classification: Parallelism, Heterogeneity and Distributed Data Processing

Track Classification: Presentations