

Advanced Data Set Composition with RNTuple

Wednesday, 22 May 2024 12:21 (1 minute)

RNTuple, the successor to ROOT's TTree I/O subsystem, is currently close to reaching production-level maturity and adoption in experiment core software as well as other analysis frameworks is well underway. As the (experimental) uses of RNTuple in production environments increases, the number of available data sets resulting from different production steps does as well, each with their own schemas. This presents the opportunity to start working towards more elaborate RNTuple access patterns. A common practice used across different stages in HEP workflows, is the in-memory vertical and horizontal composition of data sets. In the context of TTree, these compositions are referred to as "chains" and "friends", respectively. To successfully implement such compositions in RNTuple, several factors need to be taken into careful consideration. Importantly, (in)compatibility between different data sets needs to be handled transparently. Moreover, the rules that determine how the data sets can be composed have to be clearly defined. In this contribution, we will present the ongoing work to support composability of RNTuples. We will discuss the main design considerations through a selection of concrete use cases, and the steps necessary to make these designs fit naturally in the broader RNTuple implementation.

Primary author: DE GEUS, Florine (CERN/University of Twente (NL))

Co-authors: Dr PADULANO, Vincenzo Eduardo (CERN); BLOMER, Jakob (CERN); CANAL, Philippe (Fermi National Accelerator Lab. (US)); VARBANESCU, Ana-Lucia (University of Twente)

Presenter: DE GEUS, Florine (CERN/University of Twente (NL))

Session Classification: Poster Session & Lunch