



Contribution ID: 364

Type: **Oral**

Evolution of the ROOT Tree I/O

Tuesday, November 5, 2019 5:00 PM (15 minutes)

The ROOT TTree data format encodes hundreds of petabytes of High Energy and Nuclear Physics events. Its columnar layout drives rapid analyses, as only those parts (branches) that are really used in a given analysis need to be read from storage. Its unique feature is the seamless C++ integration, which allows users to directly store their event classes without explicitly defining data schemas. In this contribution, we present the status and plans of the future ROOT 7 event I/O. Along with the ROOT 7 interface modernization, we aim for robust, where possible compile-time safe C++ interfaces to read and write event data. On the performance side, we show first benchmarks using ROOT's new experimental I/O subsystem that combines the best of TTrees with recent advances in columnar data formats. A core ingredient is a strong separation of the high-level logical data layout (C++ classes) from the low-level physical data layout (storage backed nested vectors of simple types). On the high level, we present new asynchronous and thread-friendly interfaces to support parallel reading and writing. On the low-level, we show how an optimized physical data layout speeds up serialization and deserialization and facilitates fast vectorized and bulk operations. This lets ROOT I/O run optimally on the upcoming ultra-fast NVRAM storage devices, as well as file-less storage systems such as object stores.

Consider for promotion

No

Authors: BLOMER, Jakob (CERN); CANAL, Philippe (Fermi National Accelerator Lab. (US)); NAUMANN, Axel (CERN); PIPARO, Danilo (CERN)

Presenter: BLOMER, Jakob (CERN)

Session Classification: Track 2 –Offline Computing

Track Classification: Track 2 –Offline Computing