



Contribution ID: 395

Type: Poster

RNTuple: Towards First-Class Support for HPC data centers

Thursday 27 October 2022 16:10 (30 minutes)

Compared to LHC Run 1 and Run 2, future HEP experiments, e.g. at the HL-LHC, will increase the volume of generated data by an order of magnitude. In order to sustain the expected analysis throughput, ROOT's RNTuple I/O subsystem has been engineered to overcome the bottlenecks of the TTree I/O subsystem, focusing also on a compact data format, asynchronous and parallel requests, and a layered architecture that allows supporting distributed filesystem-less storage systems, e.g. HPC-oriented object stores.

In a previous publication, we introduced and evaluated the RNTuple's native backend for Intel DAOS. Since its first prototype, we carried out a number of improvements both on RNTuple and its DAOS backend aiming to saturate the physical link, such as support for vector writes and an improved RNTuple-to-DAOS mapping, only to name a few. In parallel, the latest developments allow for better integration between RNTuple and ROOT's storage-agnostic, declarative interface to write HEP analyses, RDataFrame.

In this work, we contribute with the following: (i) a redesign and evaluation of the RNTuple DAOS backend, including a mechanism for efficient population of the object store based on existing data; and (ii) an experimental evaluation of single-node and distributed analyses using RDataFrame as a proxy between the user and RNTuple, showing a significant increase in the analysis throughput for typical HEP workflows.

Experiment context, if any

References

- [1] https://www.epj-conferences.org/articles/epjconf/abs/2021/05/epjconf_chep2021_02066/epjconf_chep2021_02066.html
- [2] <https://arxiv.org/abs/2204.09043>
- [3] https://www.researchgate.net/publication/346917416_Evolution_of_the_ROOT_Tree_IO

Significance

Our contribution lies at the intersection between High Energy Physics and High Performance Computing. In this contribution, we provide key updates to RNTuple, the designated successor of the ROOT TTree I/O subsystem. RNTuple comes with a user-friendly API and aims at higher throughput and smaller files. This work describes the latest developments on RNTuple and its integration with RDataFrame, focusing on their use on HPC data centers that leverage Intel DAOS as a distributed object store.

Primary authors: LAZZARI MIOTTO, Giovanna (Universidade Federal Do Rio Grande Do Sul (BR)); LOPEZ GOMEZ, Javier (CERN)

Presenter: LAZZARI MIOTTO, Giovanna (Universidade Federal Do Rio Grande Do Sul (BR))

Session Classification: Poster session with coffee break