



Contribution ID: 102

Type: **Parallel contribution**

Persistent Data Layout and Infrastructure for Efficient Selective Retrieval of Event Data in ATLAS

Thursday 11 August 2011 12:00 (30 minutes)

The ATLAS detector [1] at CERN has completed its first full year of recording collisions at 7 TeV, resulting in billions of events and petabytes of data. At these scales, physicists must have the capability to read only the data of interest to their analyses, with the importance of efficient selective access increasing as data taking continues.

ATLAS has developed a sophisticated event-level metadata infrastructure (TAG [2]) and supporting I/O framework [3] allowing event selections by explicit specification, by back navigation, and by selection queries to a TAG database via an integrated web interface (iELSSI). These systems and their performance have been reported on elsewhere.

The ultimate success of such a system, however, depends significantly upon the efficiency of selective event retrieval. Supporting such retrieval can be challenging, as ATLAS stores its event data in column-wise orientation using ROOT TTrees [4] for a number of reasons, including compression considerations, histogramming use cases, and more.

For 2011 data, ATLAS will utilize new capabilities in ROOT to tune the persistent storage layout of event data, and to significantly speed up selective event reading.

The new persistent layout strategy and its implications for I/O performance will be presented in this paper.

[1] ATLAS Collaboration, ATLAS Detector and Physics Performance Technical Design Report, CERN-LHCC-1999-14 and CERN-LHCC-1999-15.

[2] J. Cranshaw et al, "Event selection services in ATLAS", in J. Phys.: Conf. Ser., vol. 219, 042007, 2010

[3] P. van Gemmeren, D. Malon, "The event data store and I/O framework for the ATLAS experiment at the Large Hadron Collider", in IEEE International Conference on Cluster Computing and Workshops, 2009, pp.1-8.

[4] R. Brun and F. Rademakers, "ROOT –An Object Oriented Data Analysis Framework", Nucl. Inst. & Meth. in Phys. Rev. A 389 (1997) 81-86. See also <http://root.cern.ch>

Primary authors: MALON, David (Argonne National Laboratory); VAN GEMMEREN, Peter (Argonne National Laboratory)

Presenter: VAN GEMMEREN, Peter (Argonne National Laboratory)

Session Classification: Computing in HEP

Track Classification: Computing in HEP