



ATLAS Event Data Organization and I/O Framework Capabilities in Support of Heterogeneous Data Access and Processing Models

Saturday, August 6, 2016 6:00 PM (2 hours)

Choices in persistent data models and data organization have significant performance ramifications for data-intensive scientific computing.

In experimental high energy physics, organizing file-based event data for efficient per-attribute retrieval may improve the I/O performance of some physics analyses

but hamper the performance of processing that requires full-event access.

In-file data organization tuned for serial access by a single process may be less suitable for opportunistic sub-file-based processing on distributed computing resources.

Unique I/O characteristics of high-performance computing platforms pose additional challenges.

This paper describes work in the ATLAS experiment at the Large Hadron Collider to provide an I/O framework and tools for persistent data organization

to support an increasingly heterogeneous array of data access and processing models.

Primary author: ATLAS, Collaboration (ATLAS)

Presenter: CRANSHAW, Jack (Argonne National Laboratory (US))

Session Classification: Poster Session

Track Classification: Computing and Data Handling