



Contribution ID: 377

Type: **Parallel**

I/O Strategies for Multicore Processing in ATLAS

Monday, May 21, 2012 5:50 PM (25 minutes)

A critical component of any multicore/manycore application architecture is the handling of input and output. Even in the simplest of models, design decisions interact both in obvious and in subtle ways with persistence strategies.

When multiple workers handle I/O independently using distinct instances of a serial I/O framework, for example, it may happen that because of the way data from consecutive events are compressed together, there may be serious inefficiencies, with workers redundantly reading the same buffers, or multiple instances thereof. With shared reader strategies, caching and buffer management by the persistence infrastructure and by the control framework may have decisive performance implications for a variety of design choices. Providing the next event may seem straightforward when all event data are contiguously stored in a block, but there may be performance penalties to such strategies when only a subset of a given event's data are needed; conversely, when event data are partitioned by type in persistent storage, providing the next event becomes more complicated, requiring marshaling of data from many I/O buffers.

Output strategies pose similarly subtle problems, with complications that may lead to significant serialization and the possibility of serial bottlenecks, either during writing or in post-processing, e.g., during data stream merging.

In this paper we describe the I/O components of AthenaMP, the multicore implementation of the ATLAS control framework, and the considerations that have led to the current design, with attention to how these I/O components interact with ATLAS persistent data organization and infrastructure.

Primary author: ATLAS, Collaboration (Atlas)

Co-authors: Dr MALON, David (Argonne National Laboratory (US)); CALAFIURA, Paolo (Lawrence Berkeley National Lab. (US)); VAN GEMMEREN, Peter (Argonne National Laboratory (US)); Dr BINET, Sebastien (Universite de Paris-Sud 11 (FR)); Mr TSULAIA, Vakhtang (Lawrence Berkeley National Lab. (US)); LAVRIJSEN, Wim (Lawrence Berkeley National Lab. (US))

Presenter: VAN GEMMEREN, Peter (Argonne National Laboratory (US))

Session Classification: Event Processing

Track Classification: Event Processing (track 2)