



Contribution ID: 132

Type: **Parallel**

The operational performance of the ATLAS trigger and data acquisition system and its possible evolution

Monday, May 21, 2012 3:10 PM (25 minutes)

The ATLAS experiment at the Large Hadron Collider at CERN relies on a complex and highly distributed Trigger and Data Acquisition (TDAQ) system to gather and select particle collision data at unprecedented energy and rates. The TDAQ is composed of three levels which reduces the event rate from the design bunch-crossing rate of 40 MHz to an average event recording rate of about 200 Hz.

The first part of this presentation will give an overview of the operational performance of the DAQ system during 2011 and the first months of data taking in 2012. It will describe how the flexibility inherent in the design of the system has been exploited to meet the changing needs of ATLAS data taking and in some cases push performance beyond the original design performance specification.

The experience accumulated in the ATLAS DAQ/HLT system operation during these years stimulated also interest to explore possible evolutions, despite the success of the current design. One attractive direction is to merge three systems - the second trigger level (L2), the Event Builder (EB), and the Event Filter (EF) - within a single homogeneous one in which each HLT node executes all the steps required by the trigger and data acquisition process. Each L1 event is assigned to an available HLT node which executes the L2 algorithms using a subset of the event data and, upon positive selection, builds the event, which is further processed by the EF algorithms. Appealing aspects of this design are: a simplification of the software architecture and of its configuration, a better exploitation of the computing resources, the caching of fragments already collected for L2 processing, the automated load balancing between L2 and EF selection steps, the sharing of code and services on HLT nodes.

Furthermore, the full treatment of the HLT selection on a single node allows more flexible approaches, for example “incremental event building” in which trigger algorithms progressively enlarge the size of the analysed region of interest, before requiring the building of the complete event. To spot possible limitations of the new approach and to demonstrate the benefits out-lined above, a prototype has been implemented. The preliminary measurements are positive and further tests are scheduled for the next months.

Appealing aspects of this design are: a simplification of the software architecture and of its configuration, a better exploitation of the computing resources, the caching of fragments already collected for L2 processing, the automated load balancing between L2 and EF selection steps, the sharing of code and services on HLT nodes.

Furthermore, the full treatment of the HLT selection on a single node allows more flexible approaches, for example “incremental event building” in which trigger algorithms progressively enlarge the size of the analysed region of interest, before requiring the building of the complete event.

To spot possible limitations of the new approach and to demonstrate the benefits out-lined above, a prototype has been implemented. The preliminary measurements are positive and further tests are scheduled for the next months. Their results are the subject of this paper.

Primary author: NEGRI, Andrea (Universita e INFN (IT))

Co-authors: WHITESON, Daniel (University of California Irvine (US)); ZHOU, Ning (University of California Irvine (US)); WERNER, Per (CERN); HAUSER, Reiner (Michigan State University (US)); WIEDENMANN, Werner (University of Wisconsin (US))

Presenter: NEGRI, Andrea (Universita e INFN (IT))

Session Classification: Online Computing

Track Classification: Online Computing (track 1)