



Contribution ID: 184

Type: **Poster presentation**

The online event selection architecture of the CBM experiment

Tuesday 7 June 2016 15:00 (1h 30m)

The Compressed Baryonic Matter (CBM) experiment is currently under construction at the upcoming FAIR accelerator facility in Darmstadt, Germany. Searching for rare probes, the experiment requires complex event selection criteria at an event rate of up to 10 MHz. To achieve this, all event selection is performed in a large online processing farm of several hundred nodes.

The “First-level Event Selector” (FLES) compute farm will consist primarily of standard PC components including GPGPUs and many-core architectures. The data rate at the input to this compute farm is expected to exceed 1 TByte/s of time-stamped signal messages from the detectors. The distributed input interface will be realized using custom FPGA-based PCIe add-on cards, which preprocess and index the incoming data streams.

At the intended high event rates, data from several events overlaps. Thus, there is no a priori assignment of data messages to events. Instead, event recognition is performed in combination with track reconstruction. Employing a new container data format to decontextualize the information from specific time intervals, data segments can be distributed on the farm and processed independently. This allows to optimize the event reconstruction and analysis code without additional networking overhead and aids parallel computation in the online analysis task chain. Time slice building, the continuous process of collecting the data of a time interval simultaneously from all detectors, places a high load on the network and requires careful scheduling and management. Using InfiniBand FDR hardware, this process has been demonstrated at rates of up to 6 GByte/s per node in a prototype system.

The system design is optimized for modern computer architectures by minimizing copy operations of data in memory, using DMA/RDMA wherever possible, reducing data interdependencies, and employing large memory buffers to limit the critical network transaction rate. A fault-tolerant control system will ensure high availability of the event selector.

In this presentation, we will give an overview of the online event selection architecture of the upcoming CBM experiment and discuss the premises and benefits of the design. The presented material includes results from studies on several prototype systems, including high-performance time slice building over an InfiniBand FDR network.

Author: DE CUVELAND, Jan (Johann-Wolfgang-Goethe Univ. (DE))

Co-authors: HUTTER, Dirk (Johann-Wolfgang-Goethe Univ. (DE)); HARTMANN, Helvi (Johann-Wolfgang-Goethe Univ. (DE)); LINDENSTRUTH, Volker (Johann-Wolfgang-Goethe Univ. (DE))

Presenter: DE CUVELAND, Jan (Johann-Wolfgang-Goethe Univ. (DE))

Session Classification: Poster session 1