

Event reconstruction in the LHCb Online cluster

Tuesday 24 March 2009 14:40 (20 minutes)

The LHCb experiment at the LHC accelerator at CERN will collide particle bunches at 40 MHz. After a first level of hardware trigger with output at 1 MHz, the physically interesting collisions will be selected by running dedicated trigger algorithms in the High Level Trigger (HLT) computing farm. It consists of up to roughly 16000 CPU cores and 44TB of storage space. Although limited by environmental constraints, the computing power is equivalent to that provided by all Tier-1's to LHCb. The HLT duty cycle follows the LHC collisions, thus it has several months of winter shutdown, as well as several hours a day of interfill gaps. This contribution describes the strategy for using these idle resources for data reconstruction.

Due to the specific features of the HLT farm, typical processing à la Tier-1 (1 core - 1 file) is not feasible. A radically different approach has been chosen, based on parallel processing the data in farm slices of $O(1000)$ cores. Single events are read from the input files, distributed to the cluster and merged back into files once they have been processed. A detailed description of this architectural solution and the obtained performance will be presented.

Primary authors: PUIG NAVARRO, Albert (Universidad de Barcelona); FRANK, Markus (CERN)

Presenters: PUIG NAVARRO, Albert (Universidad de Barcelona); FRANK, Markus (CERN)

Session Classification: Online Computing

Track Classification: Online Computing