

Minimising wasted CPU time with interruptible LHCb Monte Carlo

Tuesday 10 July 2018 14:00 (15 minutes)

During 2017 LHCb developed the ability to interrupt Monte Carlo simulation jobs and cause them to finish cleanly with the events simulated so far correctly uploaded to grid storage. We explain how this functionality is supported in the Gaudi framework and handled by the LHCb simulation framework Gauss. By extending DIRAC, we have been able to trigger these interruptions when running simulation on unoccupied capacity of the LHCb High Level Trigger farm, and are able to reclaim this capacity when needed for online data taking tasks. This has increased the opportunities for running Monte Carlo simulation during data taking runs as well as interfill periods and technical stops. We have also applied this mechanism to grid and cloud resources at external sites, providing the ability to reclaim capacity for operational reasons without long draining periods. In addition, the mechanism is used to increase the efficiency of the “job masonry” of packing single and multiprocessor jobs into the time slots on a single worker node, without the need for draining periods when multiple free processors must be assembled for a multiprocessor job. We explain how the Machine/Job Features mechanism is instrumental in communicating the desired finish time to LHCb jobs and virtual machines.

Authors: MCNAB, Andrew (University of Manchester); ROISER, Stefan (CERN); CLEMENCIC, Marco (CERN); JOST, Beat (CERN)

Presenter: MCNAB, Andrew (University of Manchester)

Session Classification: T3 - Distributed computing

Track Classification: Track 3 –Distributed computing