Contribution ID: **688** Contribution code: **contribution ID 688**   Type: **Oral**

CMS High Level Trigger performance comparison on CPUs and GPUs

Tuesday, 30 November 2021 17:20 (20 minutes)

At the start of the upcoming LHC Run-3, CMS will deploy a heterogeneous High Level Trigger farm composed of x86 CPUs and NVIDIA GPUs. In order to guarantee that the HLT can run on machines without any GPU accelerators - for example as part of the large scale Monte Carlo production running on the grid, or when individual developers need to optimise specific triggers - the HLT reconstruction has been implemented both for NVIDIA GPUs and for traditional CPUs. This contribution will describe how the CMS software used online and offline (CMSSW) can transparently switch between the two implementations, and compare their performance on GPUs and CPUs from different architectures, vendors and generations.

Significance

References

Speaker time zone

Compatible with Europe

Primary author: COLLABORATION, CMS

Presenters: BOCCI, Andrea (CERN); COLLABORATION, CMS

Session Classification: Track 1: Computing Technology for Physics Research

Track Classification: Track 1: Computing Technology for Physics Research