



Contribution ID: 299

Type: Talk

## Experience with the alpaka performance portability library in the CMS software

*Monday 21 October 2024 16:51 (18 minutes)*

To achieve better computational efficiency and exploit a wider range of computing resources, the CMS software framework (CMSSW) has been extended to offload part of the physics reconstruction to NVIDIA GPUs. To support additional back-ends, as well to avoid the need to write, validate and maintain a separate implementation of the reconstruction algorithms for each back-end, CMS has adopted the Alpaka performance portability library.

Alpaka (Abstraction Library for Parallel Kernel Acceleration) is a header-only C++ library that provides performance portability across different back-ends, abstracting the underlying levels of parallelism. It supports serial and parallel execution on CPUs, and extremely parallel execution on NVIDIA, AMD and Intel GPUs.

This contribution will show how Alpaka is used in the CMS software to develop and maintain a single code base; to use different toolchains to build the code for each supported back-end, and link them into a single application; to seamlessly select the best back-end at runtime, and implement portable reconstruction algorithms that run efficiently on CPUs and GPUs from different vendors. It will describe the validation and deployment of the Alpaka-based implementation in the CMS High Level Trigger, and highlight how it achieves near-native performance.

**Primary authors:** Dr BOCCI, Andrea (CERN); KORTELAJNEN, Matti (Fermi National Accelerator Lab. (US))

**Presenter:** Dr BOCCI, Andrea (CERN)

**Session Classification:** Parallel (Track 3)

**Track Classification:** Track 3 - Offline Computing