

Hardware Portable Data analysis Building Blocks for High Luminosity LHC Era: The first building block

Friday 4 July 2025 09:00 (25 minutes)

The increasing data flow for LHC Run4 requires an architecture capable of handling massive parallelism. This shift demands a heterogeneous environment, creating the need for heterogeneous programming. Various ecosystems like KOKKOS and alpaka provide portability across different accelerators such as CPUs, GPUs, and FPGAs which are the most commonly used at CERN experiments. They achieve this by switching backends at runtime, making them well-suited for such diverse environments.

The CMS experiment has integrated alpaka (an abstraction library for parallel kernel acceleration) into its data analysis software, CMSSW, for several years. The current goal is to migrate as many components as suitable from traditional serial CPU processing to GPU-based processing (both CUDA and ROCm) using alpaka.

One key component is the Phase 2 Outer Tracker data unpacker, which translates RAW data from the DAQ output to the reconstruction step. This component presents an excellent opportunity for increased parallelization and performance gains. The project presented here marks the first step toward building hardware-portable data analysis components for the HL-LHC which is the porting of the Phase 2 Outer Tracker Unpacker from the CPU-based CMSSW model to an alpaka-based CMSSW model.

Primary author: MOMED, Mohammad (IT (Informationstechnik))

Presenter: MOMED, Mohammad (IT (Informationstechnik))

Session Classification: Block 3 - Applications