Performance profiling infrastructure comparing CUDA and OneAPI/SYCL in the MadGraph4GPU project

Poster abstract

Detailed view of the performance profiling infrastructure profiling different GPUs and CPUs with different abstraction layers in use in the MadGraph4GPU project, with an emphasis on looking at how we build and look at the performance between CUDA/C++ and OneAPI/SYCL. We will also go in detail what we did to fully automate this infrastructure to provide the team with up to date plots and measurements of the current performance of the project on different hardware and a holistic view of the performance of the project over time as the team develops it.

Author: TEIG, Jorgen (Norwegian University of Science and and Technology (NTNU) (NO))

Presenter: TEIG, Jorgen (Norwegian University of Science and and Technology (NTNU) (NO))