



Contribution ID: 7

Type: **Invited**

Monte Carlo Integration on GPUs – The gVEGAS Algorithm

The HL-LHC upgrade will require use of new computing resources such as HPCs in their production workflows. Many of the planned HPCs will include significant computing power in the form of accelerators, in particular of Graphic Processing Units (GPUs). To take advantage of these new computing resources ATLAS software and workflows must be adapted to use GPUs. One part of the ATLAS workflow that is well-suited for use on HPCs is Monte-Carlo (MC) generation. It requires little input but can be compute intensive, especially for complex physics processes involving many particles. This study is an effort to explore and revive a CUDA based GPU version of MadGraph, a MC integration and event generation package, that was last developed in 2013. A widely used MC integration program, VEGAS is parallelized for running on a GPU. The CUDA architecture is used for harnessing the power and the capability of the GPU to execute the algorithm in parallel. The performance of the CUDA version will be compared to the standard serial implementation which runs on a CPU. The study of the CUDA implementation of VEGAS will yield valuable insight that will facilitate the rewriting of VEGAS to more portable frameworks such as SYCL or Kokkos.

Primary authors: DARMORA, Smita (Argonne National Laboratory); HOPKINS, Walter (Argonne National Laboratory)

Presenter: DARMORA, Smita (Argonne National Laboratory)

Session Classification: Tuesday Morning 1