



Contribution ID: 4

Type: Oral

Towards a framework for GPU event generation

Wednesday 13 March 2024 14:30 (20 minutes)

We demonstrate some advantages of a top-bottom approach in the development of hardware-accelerated code by presenting the PDFFlow-VegasFlow-MadFlow software suite. We start with an autogenerated hardwareagnostic Monte Carlo generator, which is parallelized in the event axis. This allow us to take advantage of the parallelizable nature of Monte Carlo integrals even if we do not have control of the hardware in which the computation will run (i.e., an external cluster). The generic nature of such an implementation can introduce spurious bottlenecks or overheads. Fortunately, said bottlenecks are usually restricted to a subset of operations and not to the whole vectorized program. By identifying the more critical parts of the calculation one can get very efficient code and at the same time minimize the amount of hardware-specific code that needs to be written. We show benchmarks demonstrating how simply reducing the memory footprint of the calculation can increase the performance of a $2\rightarrow 4$ process. Finally, we present summary results about the performance achieved so far for PDF query on GPU and Monte Carlo integration.

Significance

In view of recent interested from the theory community in hardware accelerators, our goal is to present a development paradigm which could accelerate in the introduction of GPUs for MC simulation.

References

https://arxiv.org/abs/2211.14056 https://arxiv.org/abs/2106.10279 https://arxiv.org/abs/2012.08221 https://arxiv.org/abs/2010.09341

Experiment context, if any

Primary authors: Dr CRUZ MARTINEZ, Juan M. (CERN); CARRAZZA, Stefano (CERN)

Presenter: CARRAZZA, Stefano (CERN)

Session Classification: Track 3: Computations in Theoretical Physics: Techniques and Methods

Track Classification: Track 3: Computations in Theoretical Physics: Techniques and Methods