



Contribution ID: 109

Type: Oral

Speeding up Madgraph5_aMC@NLO through CPU vectorization and GPU offloading: towards a first alpha release

Monday, 24 October 2022 14:30 (20 minutes)

The matrix element (ME) calculation in any Monte Carlo physics event generator is an ideal fit for implementing data parallelism with lockstep processing on GPUs and on CPU vector registers. For complex physics processes where the ME calculation is the computational bottleneck of event generation workflows, this can lead to very large overall speedups by efficiently exploiting these hardware architectures, which are now largely underutilized in HEP. In this contribution, we will present the latest status of our work on the reengineering of the Madgraph5_aMC@NLO event generator for these architectures. The new implementations of the ME calculation in vectorized C++, in CUDA and in the ALPAKA, KOKKOS and SYCL portability frameworks will be described in detail, as well as their integration into the existing MadEvent framework to keep the same overall look-and-feel of the user interface. Performance numbers will be reported both for the ME calculation alone and for the overall production workflow for unweighted event generation. First experience with an alpha release of the software supporting LHC LO processes, which is expected by the time of the ACAT2022 conference, will also be discussed.

Significance

- We plan to present the first functional release of the software usable by LHC experiments (or at least a clear timeline towards that).
- This contribution is relevant to both ACAT track1 and track3. It is relevant to track1 because it discusses approaches to exploiting heterogeneous resources which may be reused also by other HEP workloads, such as GPU/CPU data parallelism through compiler vector extensions, AOSOAs, portability frameworks and various threading implementations. It is relevant to track3 because we believe that similar large speedups on GPUs and vector CPUs are within reach also for any other MC matrix element event generator. For these reasons we kindly suggest to the organizers to also consider it for a plenary talk, which would cover the topic of speeding up Monte Carlo event generators also in more generic terms.

References

- HSFWS2020: <https://indico.cern.ch/event/941278/contributions/4101793/>
- vCHEP2021: <https://doi.org/10.1051/epjconf/202125103045>
- ICHEP2022: <https://agenda.infn.it/event/28874/abstracts/20368/>

Experiment context, if any

Madgraph5_aMC@NLO is routinely used, amongst others, by ATLAS and CMS

Primary authors: VALASSI, Andrea (CERN); Dr VUOSALO, Carl (University of Wisconsin Madison (US)); SMITH, David (CERN); FIELD, Laurence (CERN); NICHOLS, Nathan (Argonne National Laboratory (US)); MATTELAER, Olivier (UCLouvain); ROISER, Stefan (CERN); HAGEBOECK, Stephan (CERN); CHILDERS, Taylor (Argonne National Laboratory (US)); HOPKINS, Walter (Argonne National Laboratory (US))

Presenter: VALASSI, Andrea (CERN)

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

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