



Contribution ID: 237

Type: Oral

GPU Usage in ATLAS Reconstruction and Analysis

Monday, November 4, 2019 3:00 PM (15 minutes)

With GPUs and other kinds of accelerators becoming ever more accessible, High Performance Computing Centres all around the world using them ever more, ATLAS has to find the best way of making use of such accelerators in much of its computing.

Tests with GPUs – mainly with CUDA – have been performed in the past in the experiment. At that time the conclusion was that it was not advantageous for the ATLAS offline and trigger software to invest time and money into GPUs. However as the usage of accelerators has become cheaper and simpler in recent years, their re-evaluation in ATLAS's offline software is warranted.

We will show code designs and performance results of using OpenCL, OpenACC and CUDA to perform calculations using the ATLAS offline/analysis (xAOD) Event Data Model. We compare the performance and flexibility of these different offload methods, and show how different memory management setups affect our ability to offload different types of calculations to a GPU efficiently. So that an overall throughput increase could be achieved even without highly optimising our reconstruction code specifically for GPUs.

Consider for promotion

No

Primary authors: KRASZNAHORKAY, Attila (CERN); METE, Alaettin Serhan (University of California Irvine (US)); TSULAIA, Vakho (Lawrence Berkeley National Lab. (US)); SNYDER, Scott (Brookhaven National Laboratory (US)); LEGGETT, Charles (Lawrence Berkeley National Lab (US))

Presenter: KRASZNAHORKAY, Attila (CERN)

Session Classification: Track 5 – Software Development

Track Classification: Track 5 – Software Development