ACAT 2024



Contribution ID: 45

Type: Oral

A Mechanism for Asynchronous Offloading in the Multithreaded Gaudi Event Processing Framework

Wednesday 13 March 2024 15:10 (20 minutes)

High Performance Computing resources are increasingly prominent in the plans of funding agencies, and the tendency of these resources is now to rely primarily on accelerators such as GPUs for the majority of their FLOPS. As a result, High Energy Physics experiments must make maximum use of these accelerators in our pipelines to ensure efficient use of the resources available to us.

The ATLAS and LHCb experiments share a common data processing architecture called Gaudi. In Gaudi, data processing workloads are ultimately split into units called Algorithms, and Gaudi uses a smart scheduler (the Avalanche scheduler) to schedule these Algorithms on a fixed pool of CPU threads managed by Intel's TBB.

This is an architecture that efficiently fills the available CPU capacity provided the algorithms are primarily CPU-limited. However when the algorithms offload a large portion of their computational work to GPUs they can be left blocking a CPU thread, wasting precious core-time.

Here we present a prototype of an addition to this scheduler, which places such GPU-accelerated algorithms on a separate pool of dedicated threads. By making use of lightweight Boost Fibers, and the ability to suspend these fibers without suspending the underlying OS thread, we can run the GPU workload asynchronously, without blocking the thread. This allows more efficient use of the CPU resources, and where the work offloaded by a single Algorithm doesn't fill the GPU resources available can also improve GPU-efficiency by making use of separate CUDA streams.

Significance

This work presents an addition to the Gaudi Avalanche scheduler which enables it to deal with GPU-accelerated algorithms in a CPU efficient manner.

References

Experiment context, if any

ATLAS, LHCb

Primary authors: STANISLAUS, Beojan (Lawrence Berkeley National Lab. (US)); Dr LEGGETT, Charles (Lawrence Berkeley National Lab (US)); ESSEIVA, Julien (Lawrence Berkeley National Lab. (US)); CALAFIURA, Paolo (Lawrence Berkeley National Lab. (US)); TSULAIA, Vakho (Lawrence Berkeley National Lab. (US)); JU, Xiangyang (Lawrence Berkeley National Lab. (US))

Presenter: STANISLAUS, Beojan (Lawrence Berkeley National Lab. (US))

Session Classification: Track 1: Computing Technology for Physics Research

Track Classification: Track 1: Computing Technology for Physics Research