



Contribution ID: 131

Type: Poster

Exploring the use of accelerators for lossless data compression in CMS

Tuesday, 25 October 2022 11:00 (30 minutes)

The CMS collaboration has a growing interest in the use of heterogeneous computing and accelerators to reduce the costs and improve the efficiency of the online and offline data processing: online, the High Level Trigger is fully equipped with NVIDIA GPUs; offline, a growing fraction of the computing power is coming from GPU-equipped HPC centres. One of the topics where accelerators could be used for both online and offline processing is data compression.

In the past decade a number of research papers exploring the use of GPUs for lossless data compression have appeared in academic literature, but very few practical applications have emerged. In the industry, NVIDIA has recently published the `nvcomp` GPU-accelerated data compression library, based on closed-source implementations of standard and dedicated algorithms. Other platforms, like the IBM Power 9 processors, offer dedicated hardware for the acceleration of data compression tasks.

In this work we review the recent developments on the use of accelerators for data compression. After summarising the recent academic research, we will measure the performance of representative open- and closed-source algorithms over CMS data, and compare it with the CPU-only algorithms currently used by ROOT and CMS (`lz4`, `zlib`, `zstd`).

Significance

References

Experiment context, if any

CMS

Primary authors: BOCCI, Andrea (CERN); RUA, Stefan (Aalto University)

Presenter: RUA, Stefan (Aalto University)

Session Classification: Poster session with coffee break

Track Classification: Track 1: Computing Technology for Physics Research