WP-D: GPU@LIV

GridPP52

Mark Wong

University of Liverpool, HEP computing

with Robert Fay, David Hutchcroft

Available GPUs

There are several users running jobs on machines with GPUs but not sure how they are using them.

- 6 A100 80GB
- 1 A40 48GB
- 2 RTX A4000 16GB
- 1 GTX 980Ti 6GB

Particle physics

- ATLAS - Hadronic jet reconstruction originating from b-quarks and tau-leptons and new physics searches: dark matter candidates such as dark photons and axion-like particles.
- LHCb - Online reconstruction (HLT1 & 2) and using machine learning for new physics searches. The LHCb lowest level trigger is implemented in RTX A5000, implemented for Run 3 and is working.
- Mu3e - Online track reconstruction and vertexing. I'm actually working on this.

Work in progress

- We are slowly ramping up work but constrained to 0.1FTE.
- Investigate, document, and produce recommendations on the different ways to present GPU resources touching on areas including GPU partitioning, optimization of hardware specifications, queue configuration, and API support, to allow for the optimal use of resources.
- Integrating the A100s to slurm and figuring out how to implement benchmarking and accounting with our GPUs (HEPScore23?).

Work in progress cont.d

A benchmark for other processors (ARM and GPUs)

- We have HEP workloads for ARM from a number of experiments.
- Workloads with GPUs are just emerging.
- Ongoing efforts to develop a benchmark for GPUs and power consumption.
- Assess progress on GPU benchmarking and test implementation if ready.