



Contribution ID: 185

Type: Poster

## Optimized GPU usage in High Energy Physics applications

*Wednesday, 26 October 2022 11:00 (30 minutes)*

Machine Learning (ML) applications, which have become quite common tools for many High Energy Physics (HEP) analyses, benefit significantly from GPU resources. GPU clusters are important to fulfill the rapidly increasing demand for GPU resources in HEP. Therefore, the Karlsruhe Institute of Technology (KIT) provides a GPU cluster for HEP accessible from the physics institute via its batch system and the Grid. As the exact hardware needs of such applications heavily depend on the ML hyperparameters, a flexible resource setup is necessary to utilize the available resources as efficient as possible. Therefore, the multi-instance GPU feature of the Nvidia A100 GPUs was studied. Several neural network training scenarios performed on the GPU cluster at KIT are discussed to illustrate possible performance gains and the setup that has been used.

### Significance

The basics we use are HTCondor and the MIG feature from NVIDIA and are described in other publications. However, we provide the resources, as one of a handful of Grid sites, to the Grid. Furthermore, the resources are shared with end-users with a more complex set of resource requirements than Grid jobs. Our experience and ideas on how to use GPUs efficiently in such an environment seem unique.

### References

### Experiment context, if any

CMS

**Primary author:** VOIGTLAENDER, Tim (KIT - Karlsruhe Institute of Technology (DE))

**Co-authors:** QUAST, Gunter (KIT - Karlsruhe Institute of Technology (DE)); GIFFELS, Manuel (KIT - Karlsruhe Institute of Technology (DE)); SCHNEPF, Matthias; WOLF, Roger (KIT - Karlsruhe Institute of Technology (DE))

**Presenter:** VOIGTLAENDER, Tim (KIT - Karlsruhe Institute of Technology (DE))

**Session Classification:** Poster session with coffee break

**Track Classification:** Track 1: Computing Technology for Physics Research