MLperf™ Performance Report

Silverton Consulting, Inc. StorInt™ Dispatch

This Storage Intelligence (StorInt™) dispatch covers the MLperf™ v1.0 series of AI-ML-DL model training and inferencing benchmarks[1]. This report focuses on **training** activity for the **High Performance Computing (HPC) environment.** There are 3 DNN models represented in the MLperf v1.0 HPC training benchmark: the **CosmoFlow, DeepCAM and OpenCatalyst** models.

- **CosmoFlow** is a 3D CNN cosmology parameter prediction model from LBNL which takes as input 3D segments of the universe (with 4 redshift buckets) and predicts OmegaM, Sigma8 and Ns cosmological parameters for that universe, at a mean average error of 0.124.
- **DeepCAM** is a convolutional encoder-decoder segmentation model trained on weather simulation data intended to identify unusual weather events such as hurricanes and atmospheric rivers. DeepCAM was the first model to scale to the entire OLCF Summit system. Models are trained until object classification intersection over union is > 0.82.
- **OpenCatalyst** is a DNN model from Meta (Facebook) and CMU and is used to identify potentially new catalysts/absorbents to assist in chemical reactions for energy storage applications. The training dataset is based on simulations of atoms physics and quantum mechanics and includes atom position, geometry, energy levels, etc. There were only 4 submissions for the OpenCatalyst workload and as such we will save this for a later time.

MLperf v1.0 HPC training DeepCAM benchmark results

We start our discussion with HPC DeepCAM training results in Figure 1.

[1] All MLperf inferencing and training results are available at https://mlcommons.org/en/ as of 2022 Feb 27.

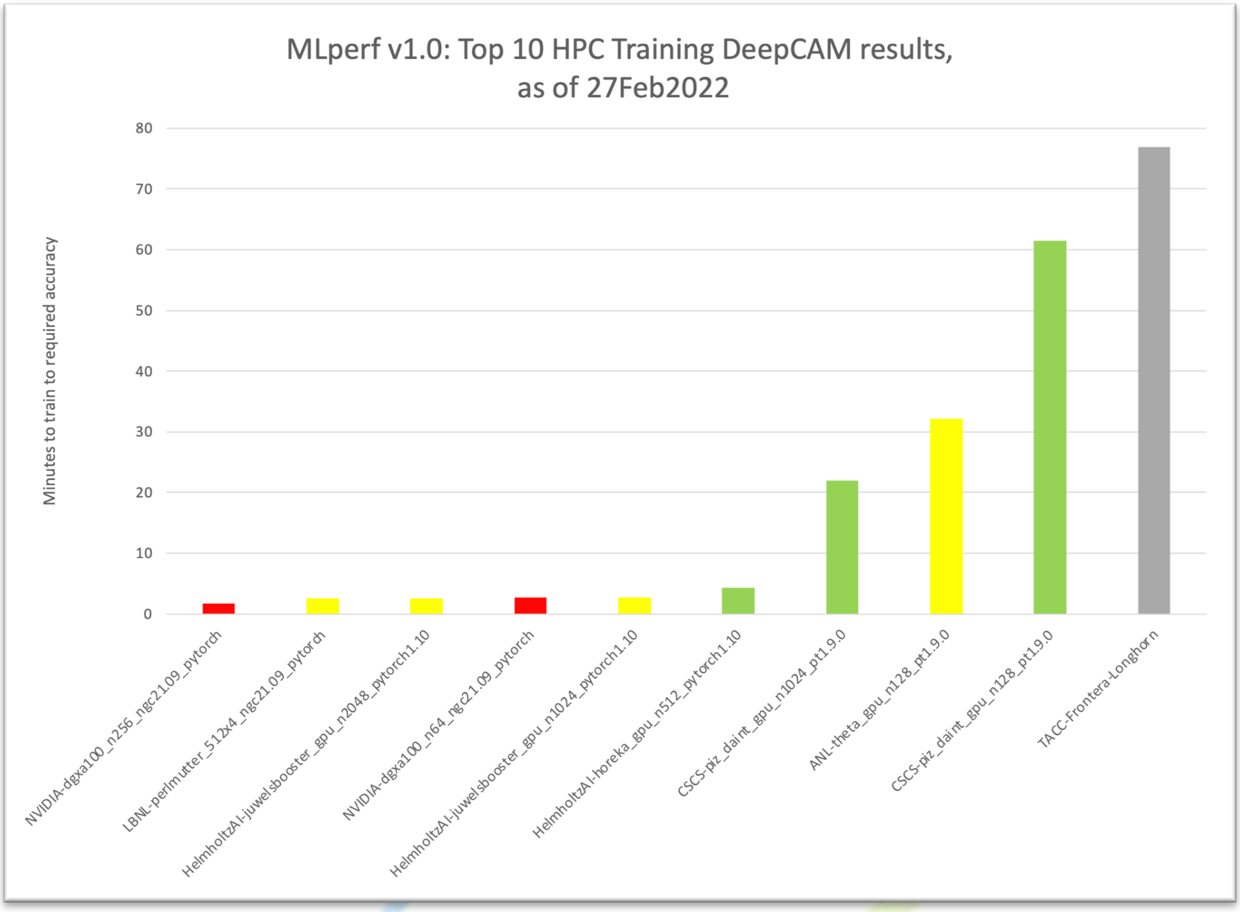

Figure 1 Top 10 MLperf v1.0 HPC DeepCAM training results

We have color coded the bars to indicate the type of NVIDIA GPUs in use by the submissions, red= A100 SXM-80GB GPUs, yellow=A100 SXM-40GB GPUs, green (shades)= A100 PCIe-40/16GB GPUs, and grey=V100-SXM2-16GB GPUs. In figure 1, the top 3 slots, came in at 1.7, 2.5- and 2.6-minutes training, respectively and all used 2048GPUs.

We find interesting the two NVIDIA DGXA100 submissions, #1 and #4. The #4 submission took 2.7 minutes to train approximately 1.6X longer. But #1 had 2048 (SXM-80GB) GPUs with 512 AMD EPYC 7742 CPUs while the #4 had 512 of the same GPUs and 128 of the same CPU or #1 had 4X the GPU and CPU hardware of #4 but only came in 1.6x faster. We conclude from this that at some level, DeepCAM training has some other bottleneck besides GPUs and CPUs. We would guess this has something to do with data transfer, storage memory to GPU memory perhaps.

In Figure 2, we show HPC CosmoFlow training results.

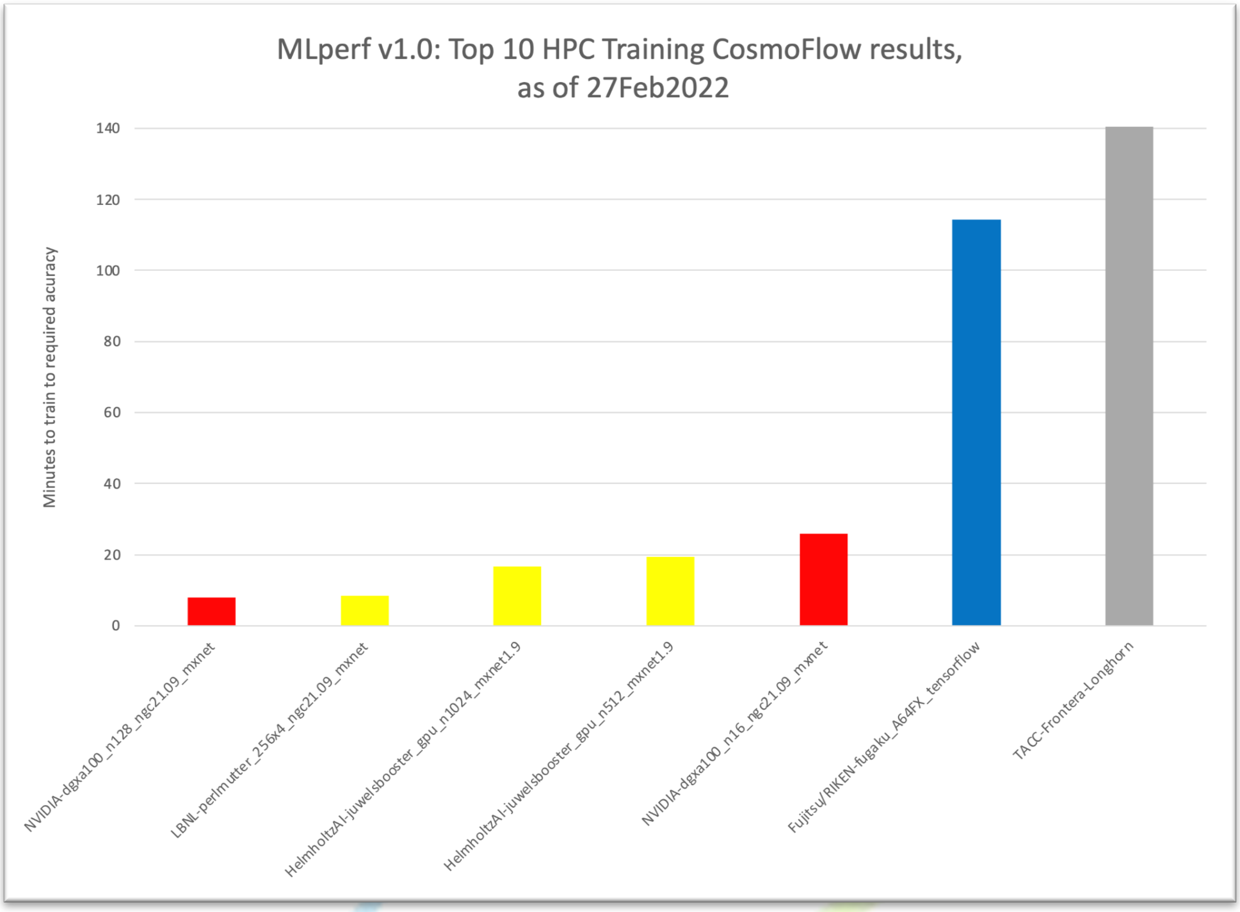

Figure 2 Top 10 MLperf v1.0 HPC CosmoFlow training results

For CosmoFlow, there were only 7 benchmark submissions. We used the same color coding here that was used above with the lone addition that blue=NO GPUs. In Figure 2, the top 3 solutions coming in at 8.0, 8.5 and 16.7 minutes of training time all had 1024 GPUs.

Again, we find the two NVIDIA DGX submissions especially interesting. The #5 ranked submission above trained CosmoFlow in 25.8 minutes. The #1 system had 1024 (SXM-80GB) GPUs plus 256 AMD EPYC 7742 CPUs while the #5 submission had 128 GPUs and 32 AMD EPYC 7742 CPUs, or #1 had 8X the GPU and CPU hardware as #5 but only achieved a 3.2X speed up. Once again in the CosmoFlow benchmarks there's some other bottleneck than GPU-CPUs and we would guess it's data transfer.

Unclear what to say about the #6 submission with 0 GPUs. It used TensorFlow with Mesh Tensorflow, and 512 FUJITSU Processor A64FX CPUs. We are familiar with some Intel activity to (DL) boost DNN training for CPUs but not aware of the A64FX processor architecture. However, we would assume that it has some serious Floating point matrix mathematical speedups in their instruction set.

Significance

One point is that it's nice to see that just throwing GPUs and CPUs at (DeepCAM and CosmoFlow) HPC DNN training models is not the only answer. There's very little information on what he storage systems are to support these submissions. The timing is taken from the start when the 1st byte of data hits the processor (GPU or CPU) complex until training reaches its required levels of accuracy.

Also, there's no indication on whether GDS is in use or not. As you should recall GDS is NVIDIA's proprietary protocol to speed up data transfers from storage memory to GPU memory. We assume that the DGX systems used GDS, but we don't have any idea about the other submissions. This would give DGX submissions some additional speedup if it were the only submission that used it.

If MLperf were listening, we would ask that they supply more information in their summary charts on the storage and interface configurations used for each of their submissions.

As always, suggestions on how to improve any of our performance analyses are welcomed.

This report was sent out to subscribers as part of our **free, monthly Storage Intelligence e-newsletter**. If you are interested in receiving future storage performance analyses along with recent product announcement summaries, please use the QR code (below right) to sign up for your own copy.

Silverton Consulting, Inc., is a U.S.-based Storage, Strategy & Systems consulting firm offering products and services to the data storage community

Newsletter signup QRcode