



Contribution ID: 112

Type: Poster

Scalable GNN Training for Track Finding

Wednesday, 13 March 2024 16:15 (30 minutes)

Graph Neural Networks (GNNs) have demonstrated significant performance in addressing the particle track-finding problem in High-Energy Physics (HEP). Traditional algorithms exhibit high computational complexity in this domain as the number of particles increases. This poster addresses the challenges of training GNN models on large, rapidly evolving datasets, a common scenario given the advancements in data generation, collection, and increase in storage capabilities. The computational and GPU memory requirements present significant roadblocks in efficiently training GNNs on large graph structures. One effective strategy to reduce training time is distributed data parallelism on multi-GPUs, which involves averaging gradients across the devices used for training.

This poster will report the speed-up of GNN training time when using distributed data parallelism with different numbers of GPUs and computing nodes. Running GNN training with distributed data parallelism leads to a decrease in accuracy. We are investigating the relationship between the number of devices and model accuracy degradation and strategies to mitigate it. Preliminary results on the TrackML dataset will be reported. GPU nodes from Perlmutter at NERSC will be used to run the experiments.

Significance

As the availability of HPC platforms with multi-GPUs increases, distributed deep learning training becomes an essential tool for exploring and experimenting with cutting-edge deep learning architectures and methodologies. By handling larger datasets and complex models, researchers and HEP scientists can push the boundaries of AI capabilities to improve the physics performance of track-finding experiments.

References

Ju, X., Murnane, D., Calafiura, P., Choma, N., Conlon, S., Farrell, S., ... & Lazar, A. (2021). Performance of a geometric deep learning pipeline for HL-LHC particle tracking. *The European Physical Journal C*, 81, 1-14.
Lazar, A., Ju, X., Murnane, D., Calafiura, P., Farrell, S., Xu, Y., ... & Lucas, A. (2023, February). Accelerating the Inference of the Exa. TrkX Pipeline. In *Journal of Physics: Conference Series* (Vol. 2438, No. 1, p. 012008). IOP Publishing.

Experiment context, if any

We report the results of training GNN models on the TrackML dataset. Even if this dataset is based on a simulation of a generic HL-LHC experiment tracker, the results could be extended to design and evaluate particle tracking algorithms for any of the experiments.

Primary authors: LADUSKA, Ivan (Youngstown State University); REEVES, Brenden (Youngstown State University); MANJEROVIC, Caroline (Youngstown State University); LAZAR, Alina (Youngstown State University); PHAM, Minh-Tuan (University of Wisconsin Madison (US)); CHAN, Jay (Lawrence Berkeley National Lab.

(US)); MURNANE, Daniel Thomas (Lawrence Berkeley National Lab. (US)); JU, Xiangyang (Lawrence Berkeley National Lab. (US)); CALAFIURA, Paolo (Lawrence Berkeley National Lab. (US))

Presenter: LAZAR, Alina (Youngstown State University)

Session Classification: Poster session with coffee break

Track Classification: Track 2: Data Analysis - Algorithms and Tools