Aerodynamics and Fluid Mechanics
Engineering and Design
Technical University of Munich

Accelerating Scientific ML through Multi-GPU Scaling

Semester Thesis/ Interdisciplinary Project (IDP) /Guided Research

Rapid advancement of machine learning (ML) has opened new frontiers in understanding and modeling complex fluid dynamics phenomena. Fluid dynamics, governed by the Navier-Stokes equations, underpins a wide range of applications, from weather forecasting and aerodynamics to energy systems and biomedicine. However, the inherent non-linearity and high-dimensional nature of these problems present challenges for traditional numerical solvers, particularly in scenarios that require large-scale simulations.

Multi-GPU training and scaling are crucial for modern machine learning because models and datasets are growing rapidly in size and complexity. Training a large neural network on a single GPU can take days, limiting both productivity and experimentation. By distributing computation across multiple GPUs, training becomes significantly faster, enabling faster iteration cycles and reducing time-to-result. Scaling also makes it possible to train models that exceed the memory capacity of a single GPU, unlocking the ability to handle larger batch sizes and higher resolution data. For both research and real-world applications, efficient multi-GPU training is essential to push the boundaries of model performance while keeping computational costs and development timelines manageable.
In this work, we try to scale up the current implementation of an existing training pipeline to multi-GPUs and potentially to multinode training.

Milestones

- Understand the existing pipeline, which is mainly based on the HuggingFace trainer class.
- Try out the built-in multi-GPU implementation from HuggingFace and discover potential bugs.
- Benchmark performance on multi-GPUs

Requirements

- Experience with Python (esp. Pytorch) is a must.
- Some knowledge of machine learning and parallel programming.
- Ability to work independently.

Optional

- Familiarity with HuggingFace transformers library.
- Github workflows.

Contact

Harish Ramachandran harish.ramachandran@tum.de with the subject "Interested to contribute to the development of the benchmarking repository". Please also attach your CV, current grade report, and link to Github (if worked on any open-source projects).