aiXcelerate 2021: Part II - Machine Learning (ML)

HPC.NRW Competence Network

CLAIX support for Distributed Machine Learning

Radita Liem (RWTH)

aiXcelerate 2021: Part II - Machine Learning (ML)

Background & Motivation

- Machine learning growing popularity in scientific approach attracts initiatives to scale up machine learning performance in cluster environment

- Machine learning tasks are often compute intensive, making them similar to traditional high-performance computing (HPC) applications.

- There are several initiatives out there to scale machine learning into the cluster environment such: **Horovod, Pytorch 2, LBANN,** and **Tarantella**

- With new technology, there‘s a need to model its performance accordingly for further improvement

Short Introduction for Distributed Machine Learning

HPC.NRW

There are three most common strategies in distributed machine learning:

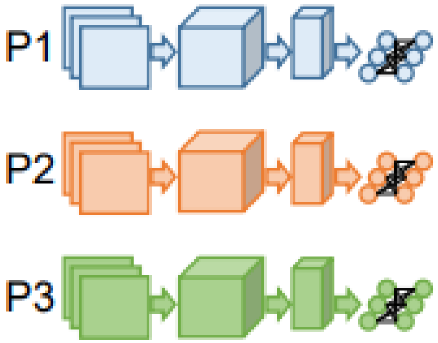

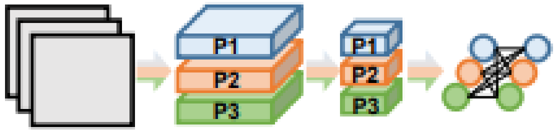

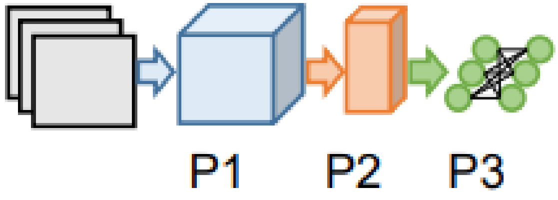

Data Parallelism

- Split data in micro-batches
- Easy to implement
- Best strong scalability
- Does not improve memory consumption

Model Parallelism

- Split layers in partitions
- Necessary for very large layers
- Can improve load balancing
- Hard and time consuming to implement

Layer Pipelining

- Split DNN in partitions
- Improves memory consumption
- Hard to achieve good load balance
- Difficult to implement

From Tarantella, a Framework for Distributed Deep Learning. Fraunhofer Center Machine Learning

Score-P/Vampir Trace from Horovod MNIST Training

Tensorflow - **CPU**

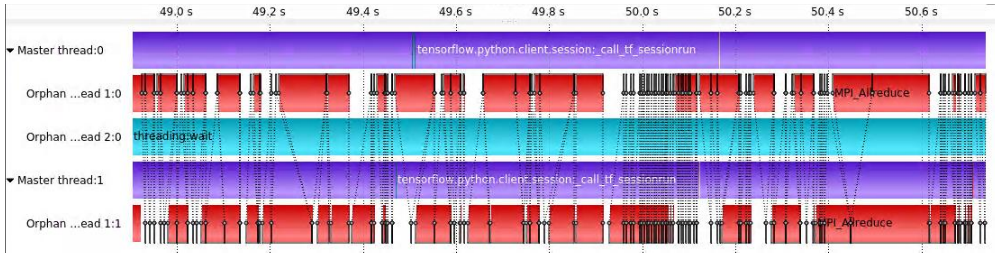

Pytorch - **CPU**

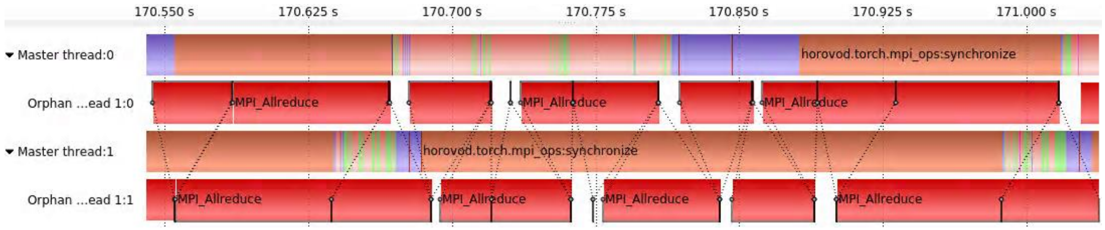

Tensorflow - **GPU**

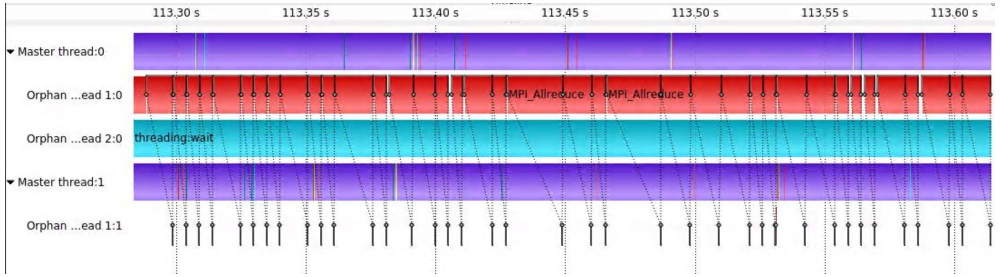

Pytorch - **GPU**

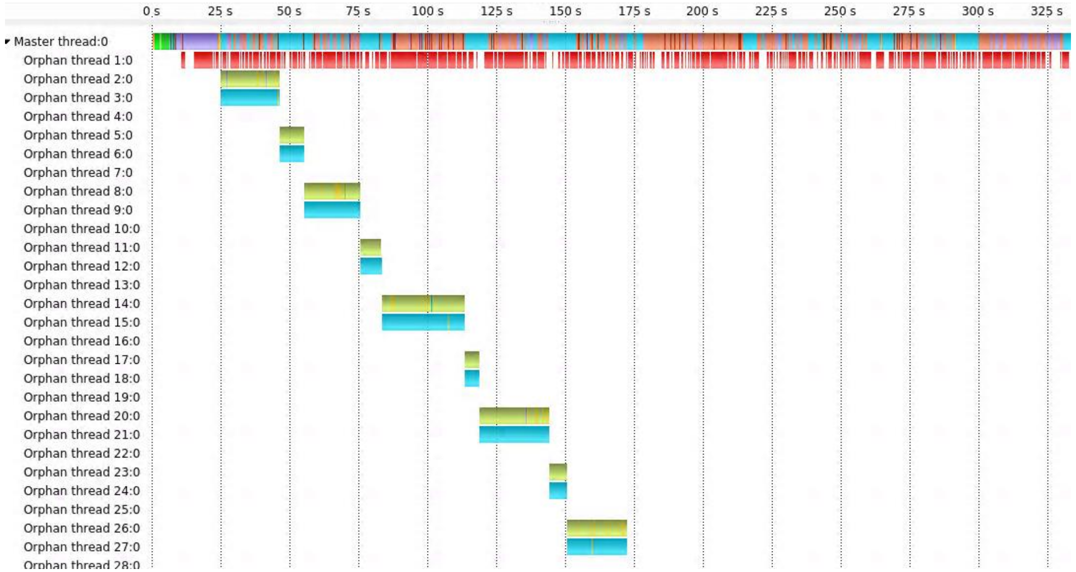

POP Assessment from Horovod MNIST Training

HPC.NRW

	Threads	Parallel Efficiency	Load Balance	Communication Efficiency
Tensorflow GPU	5	30%	73%	41%
Tensorflow CPU	5	13%	67%	19%
PyTorch GPU	124	3%	5%	49%
PyTorch CPU	4	20%	65%	31%

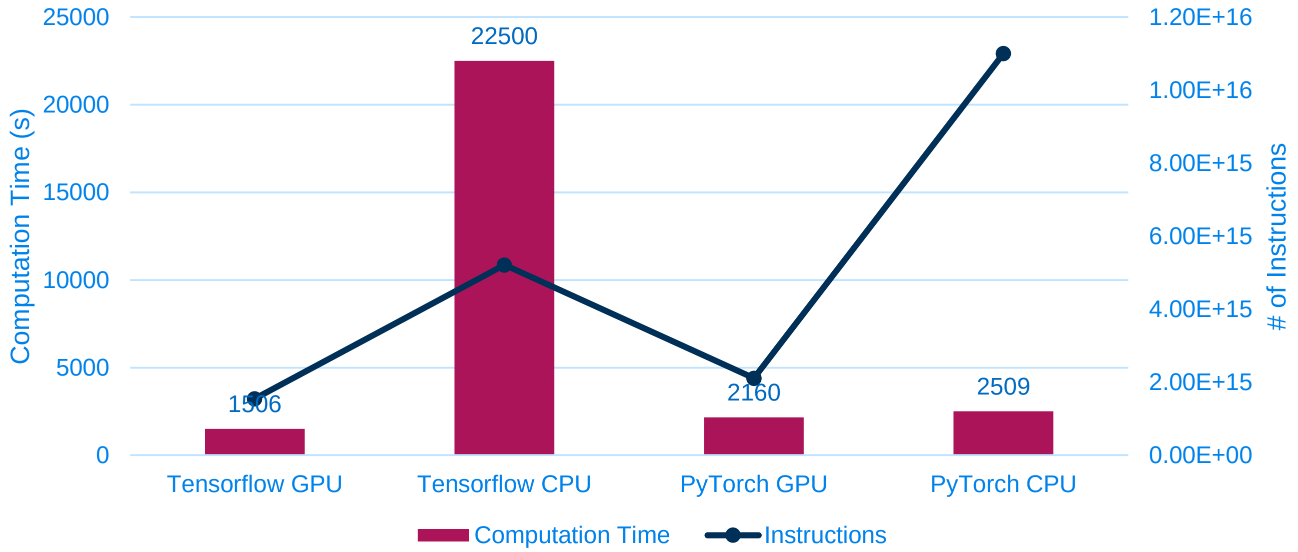

CPU and GPU performance depends on the model size and the problem (e.g. classification). The framework can do various work, this is only one sample

Support for CLAIX

HPC.NRW

- Running the application:
  - Container (Singularity)
  - Building the environment (conda, venv, additional libraries)
- Performance analysis
  - Profiling, tracing, and scaling tests
- Troubleshooting
  - Contact service desk