

MPI-based decentralized learning framework for large-scale networks

Yongding Tian

Zaid Al-Ars

Maksim Kitsak

Quantum and Computer Engineering

Email: Y.Tian-3@tudelft.nl

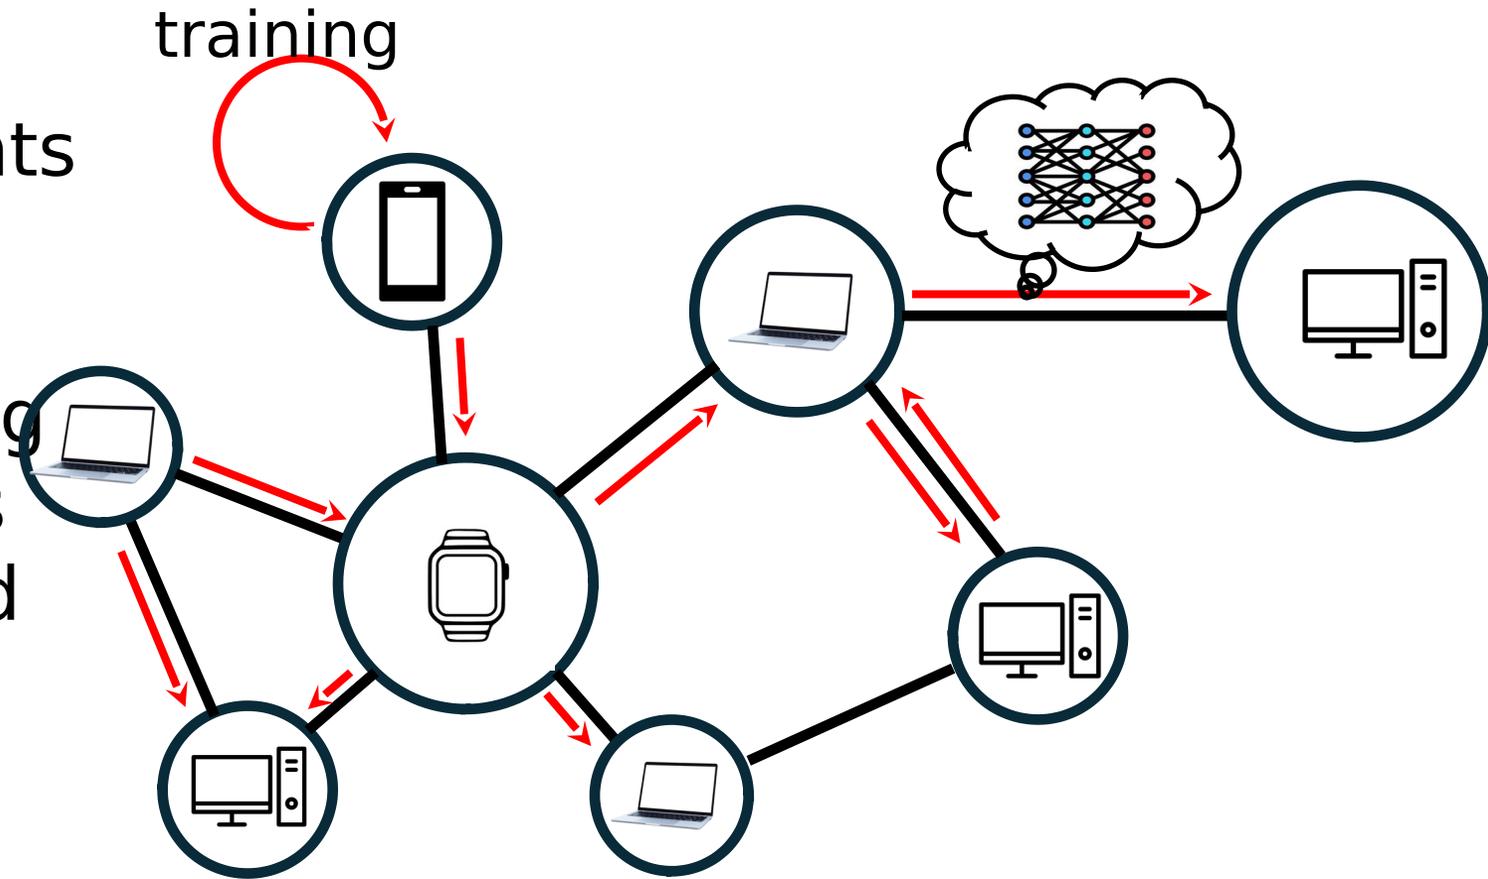
Delft University of Technology

What is decentralized learning?

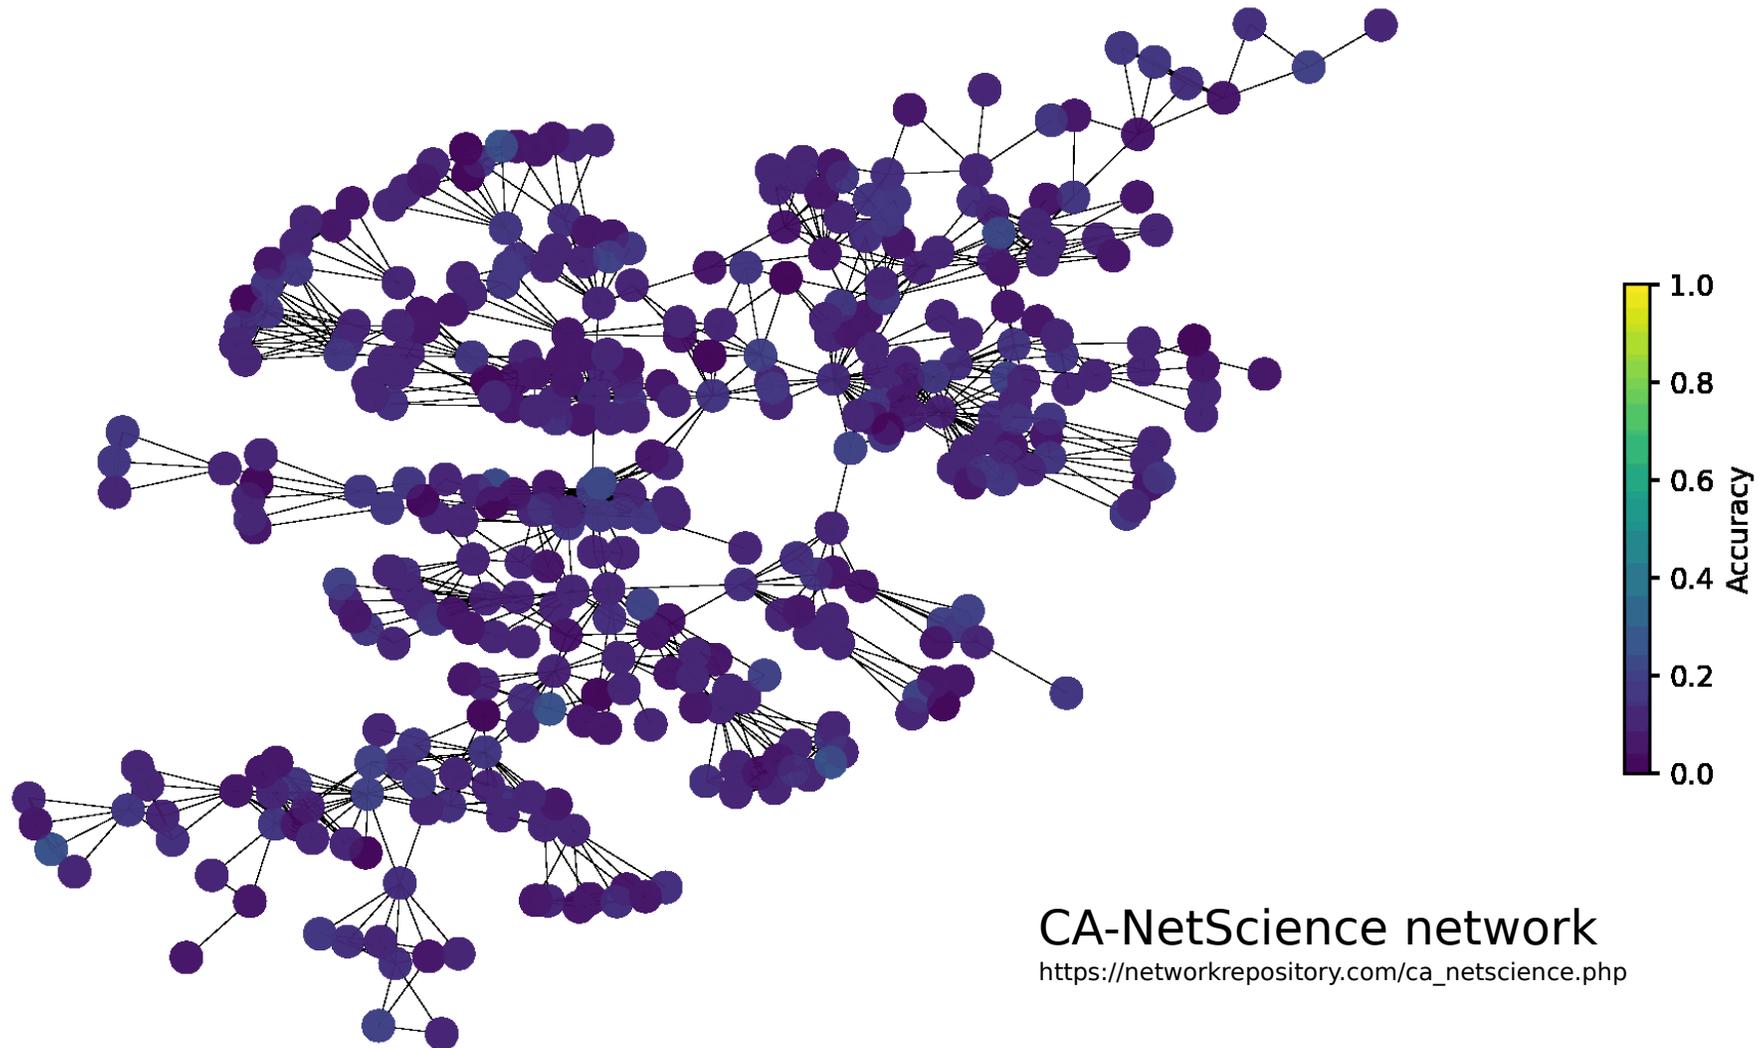
1. Share model weights

2. All nodes perform training and sharing models with others simultaneously and asynchronous

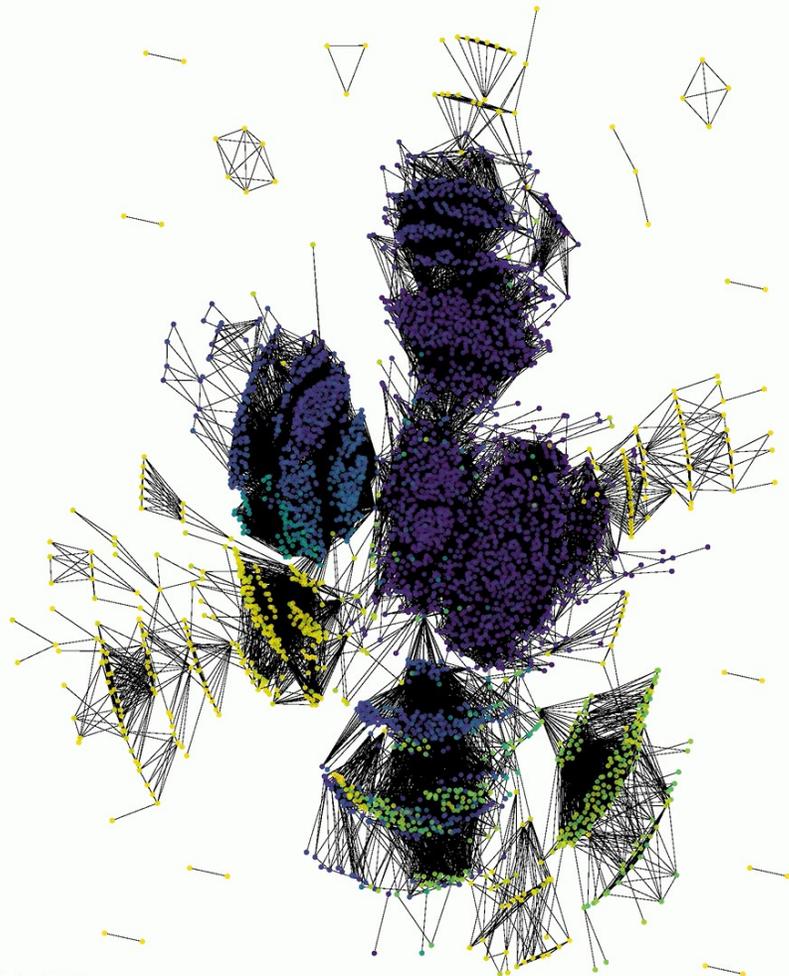
3. This system allows us to investigate how a good model is propagated in a network



Why do we investigate the model propagation?



Investigate the knowledge propagation in a network



tick = 15600

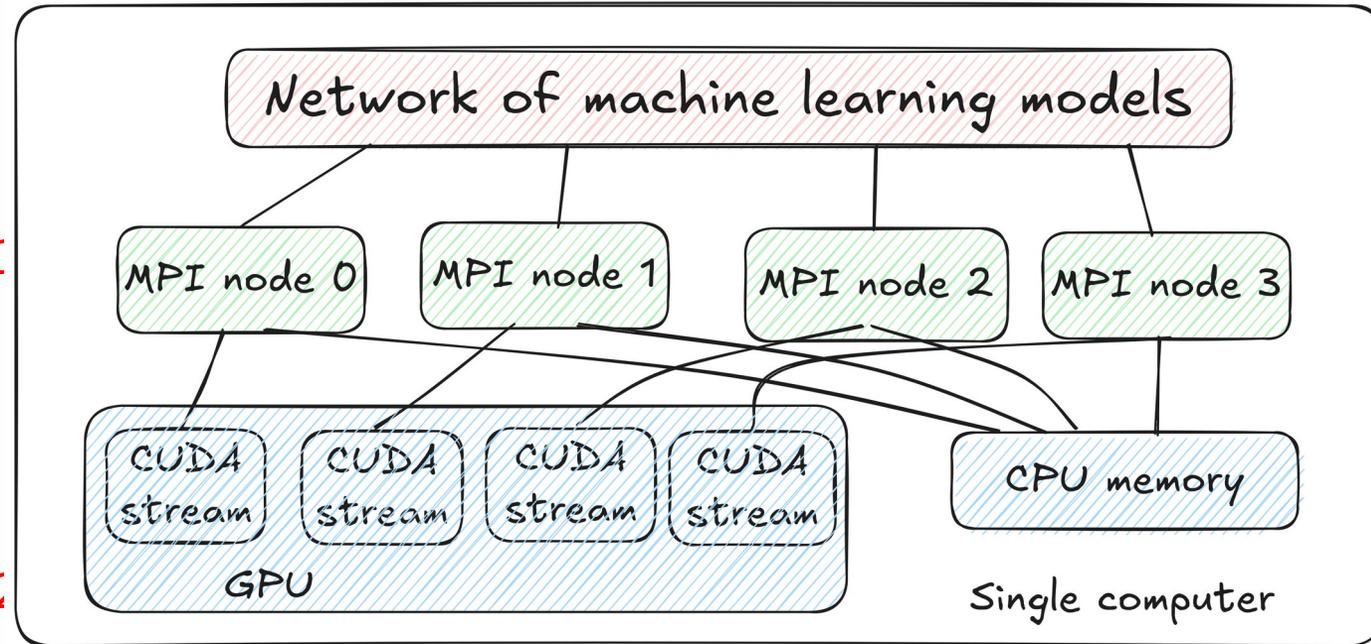
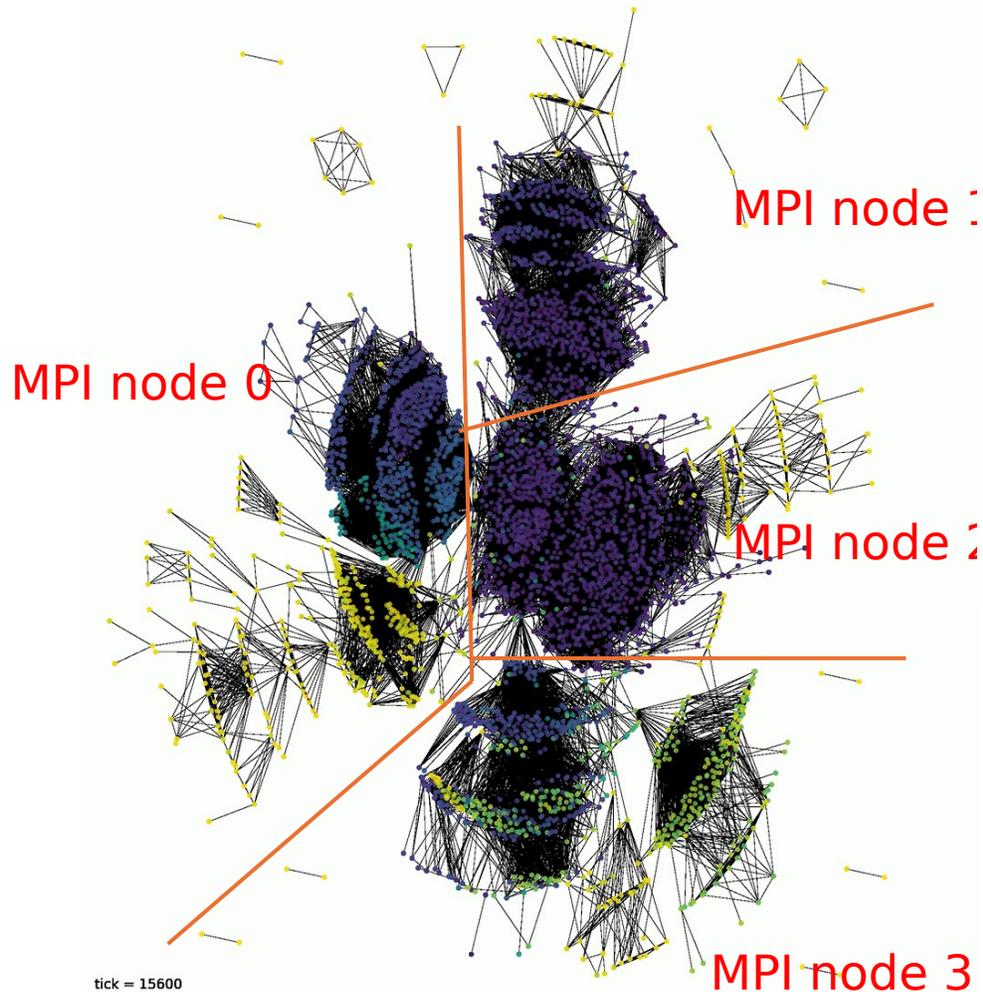
GPU memory size becomes the major bottleneck for large scale networks.

We have to use larger but slower CPU memory.

Facebook network

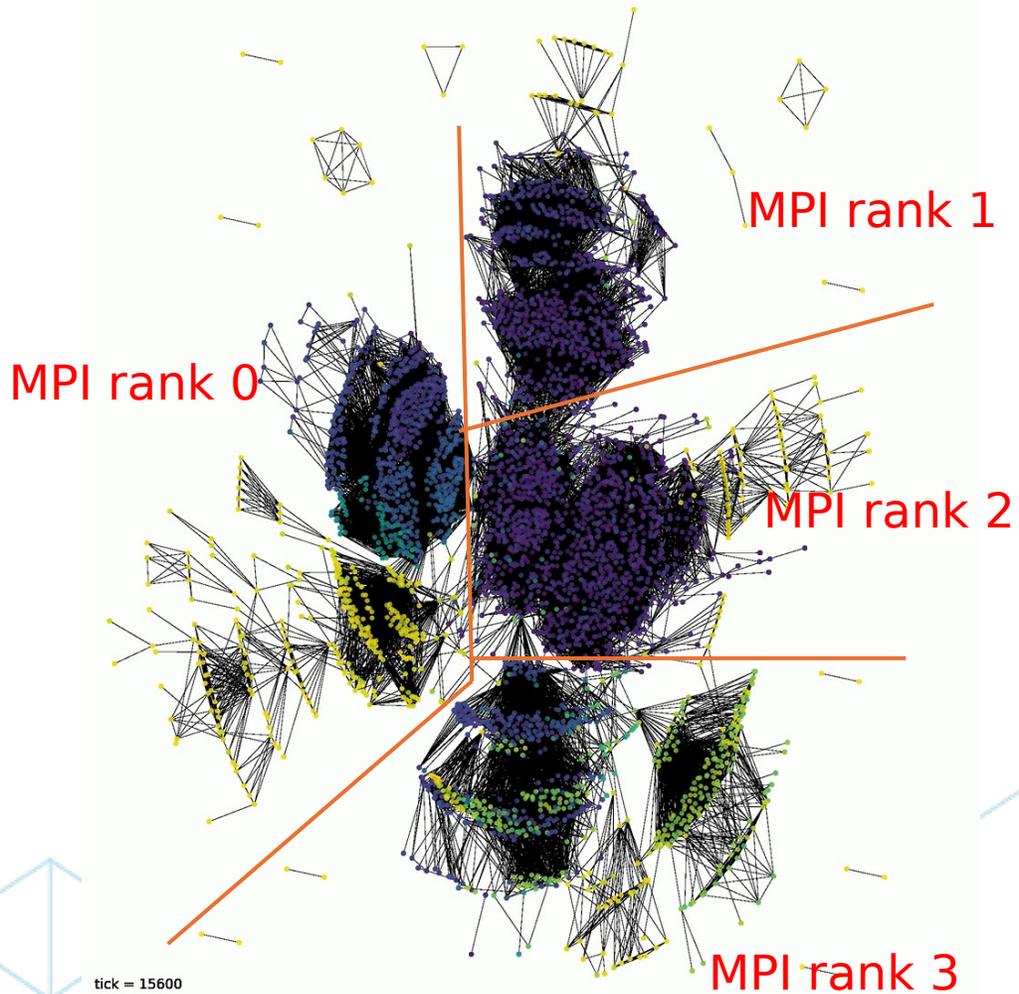
<https://snap.stanford.edu/data/ego-Facebook.html>

Simulator for large scale networks

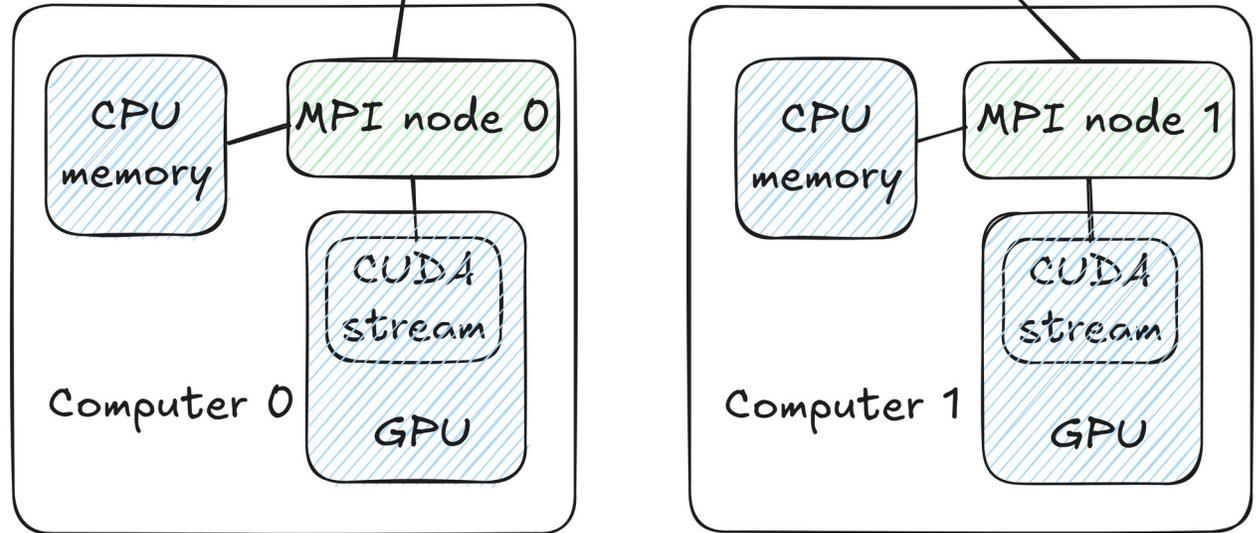


In a single computer environment
To reduce the memory copying overhead, we usually allocate >1 MPI process on a single computer to overlap the memory copying time and kernel execution time.

Simulator for large scale networks



Network of machine learning models



In a cluster environment



Memory consumption

Memory consumption	LeNet5 + MNIST	Resnet18 + CIFAR10
Optimizer	SGD with momentum	Adam optimizer
Batch size	64	128
50 nodes		30GB
500 nodes		280GB
10,000 nodes	94GB	
100,000 nodes	710GB	

All experiments are performed on a system with 1TB of memory and a single 2080ti 11GB GPU

The test network is a random regular network with an average degree of 8.

Computation performance

We test the scaling performance on a single computer with 2080Ti.

Train a single ResNet18 without MPI: **168ms / iteration**

1. Train a network composed of 50 ResNet18 nodes with MPI rank 4: **9.348s / iteration**

2.