Efficiency Benchmarking of Distributed Machine Learning via Cloud

Xiaorui Liu

Computer Science and Engineering
Michigan State University
http://cse.msu.edu/~xiaorui/
xiaorui@msu.edu

Large-scale Machine Learning

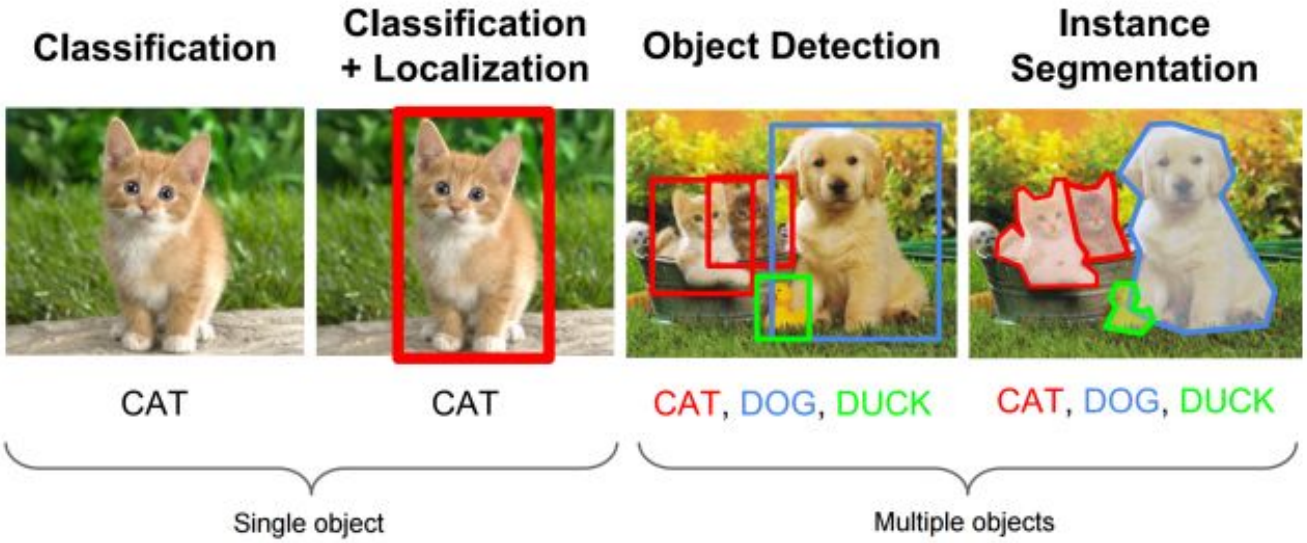

1.2 M images

Distributed Machine Learning

Limited computation power and memory space in a single machine

High Performance Computing

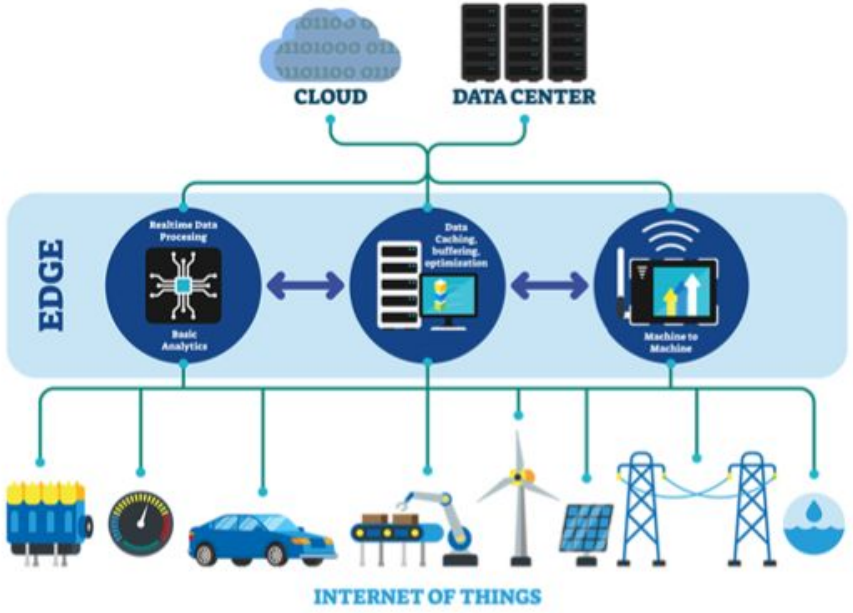

Edge Computing

Parallel SGD

$$\mathbf{x}^* := \arg\min_{\mathbf{x}\in\mathbb{R}^d} \left[f(\mathbf{x}) := \frac{1}{n} \sum_{i=1}^{n} f_i(\mathbf{x}) \right]$$

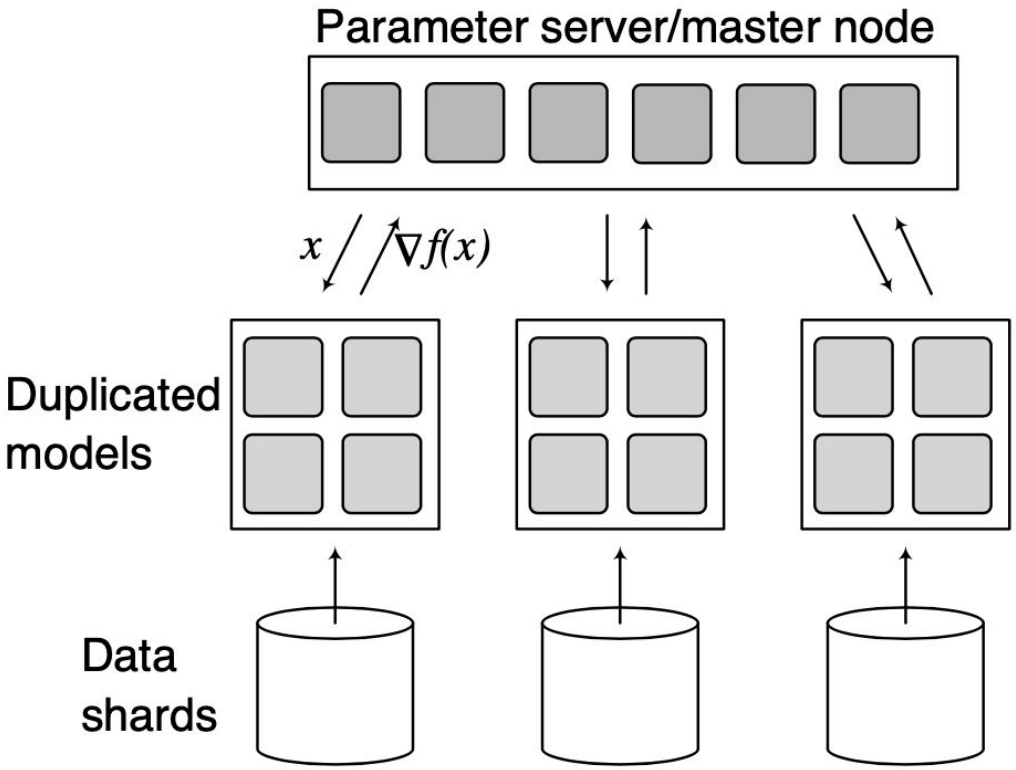

Parallel SGD

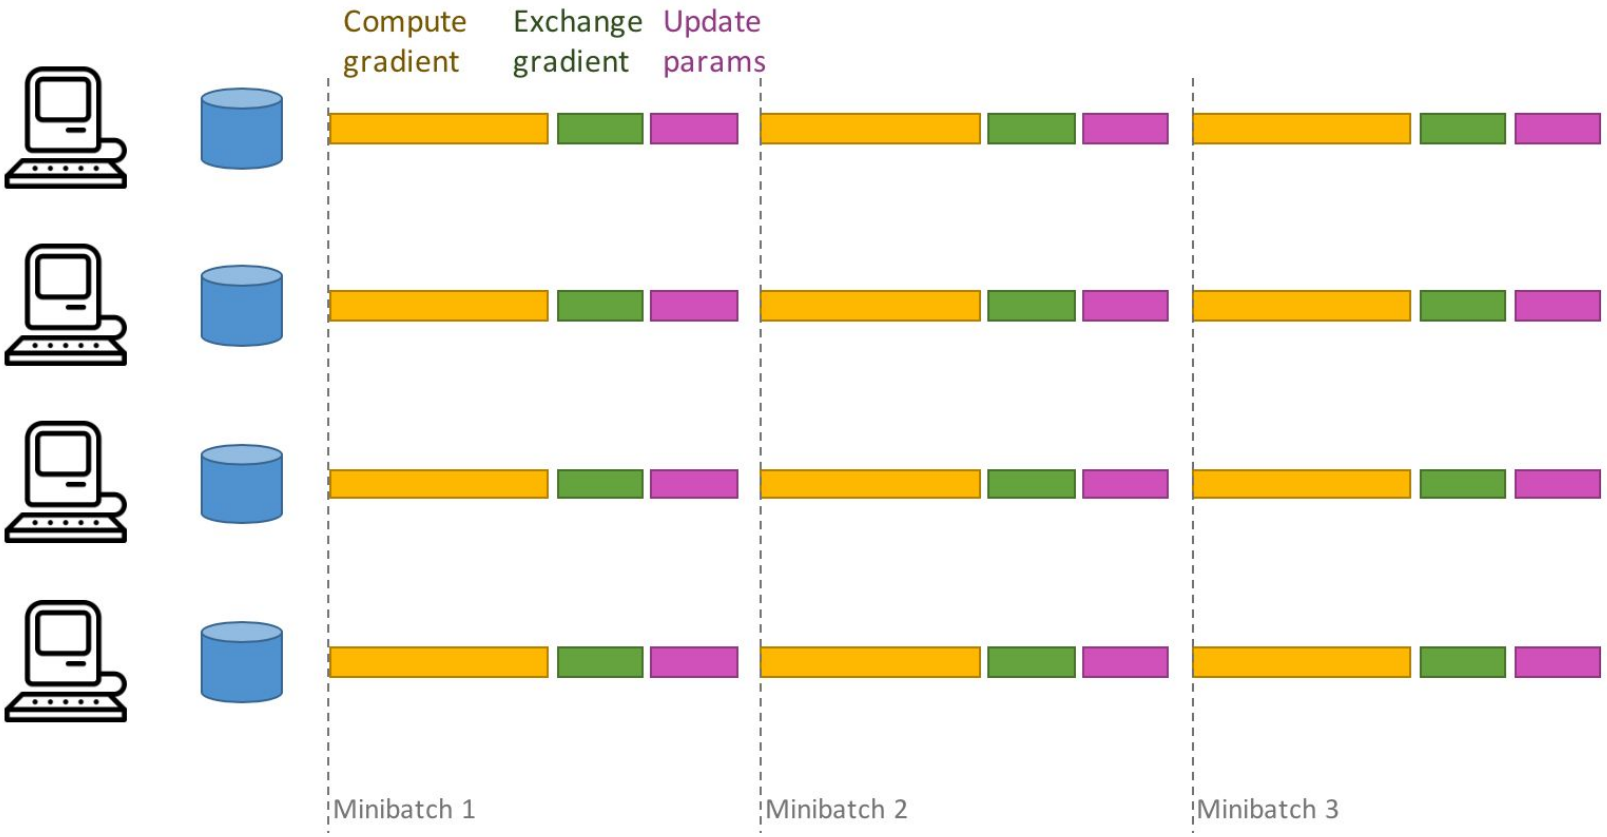

Time cost = local computation + communication + model update

Communication Bottleneck

$$\text{Communication time} = \text{latency} + \frac{\text{model size}}{\text{network bandwidth}}$$

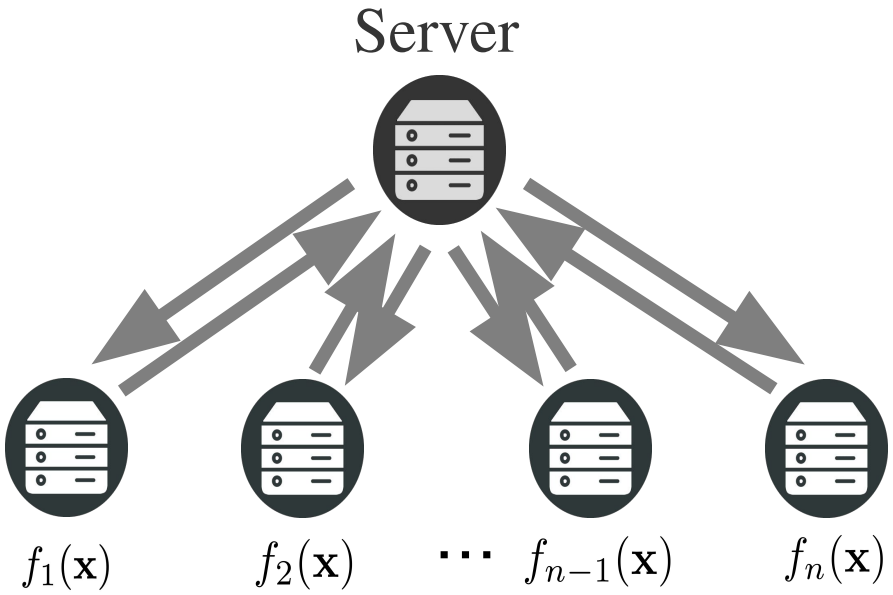

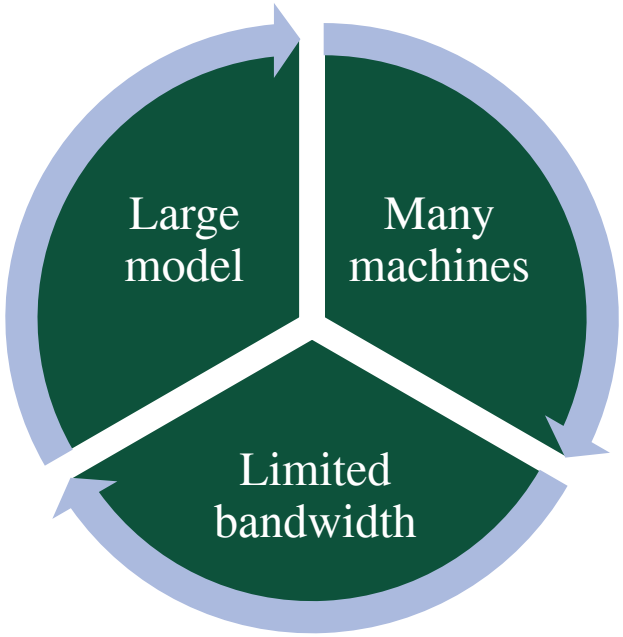

Communication bottleneck $\Rightarrow$ Poor efficiency & scalability

Project Goal: Efficiency Benchmarking

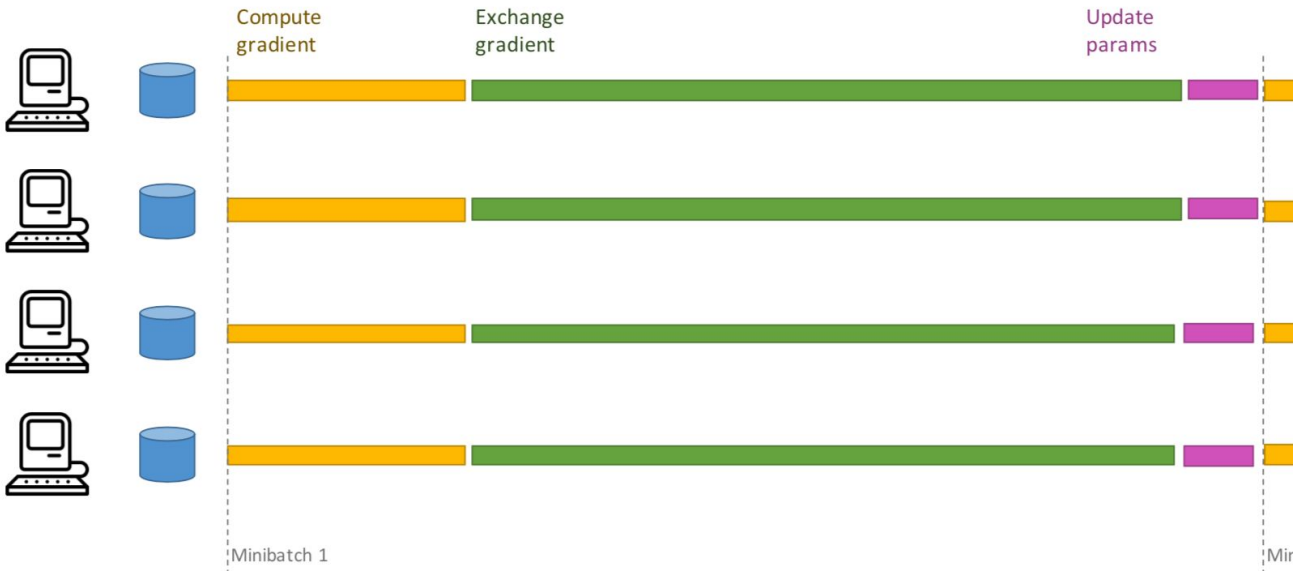

Measure the time cost decomposition and speedup under different **hardware and problem configurations** such as

- *Number of machines or GPUs*
- *Machine learning model size*
- *Network communication bandwidth*

Cloud Computing

Advantages: resource elasticity and flexible configuration

- Multiple Virtual Machines with CPU/GPU/Memory
- Communication Network
- Data Storage

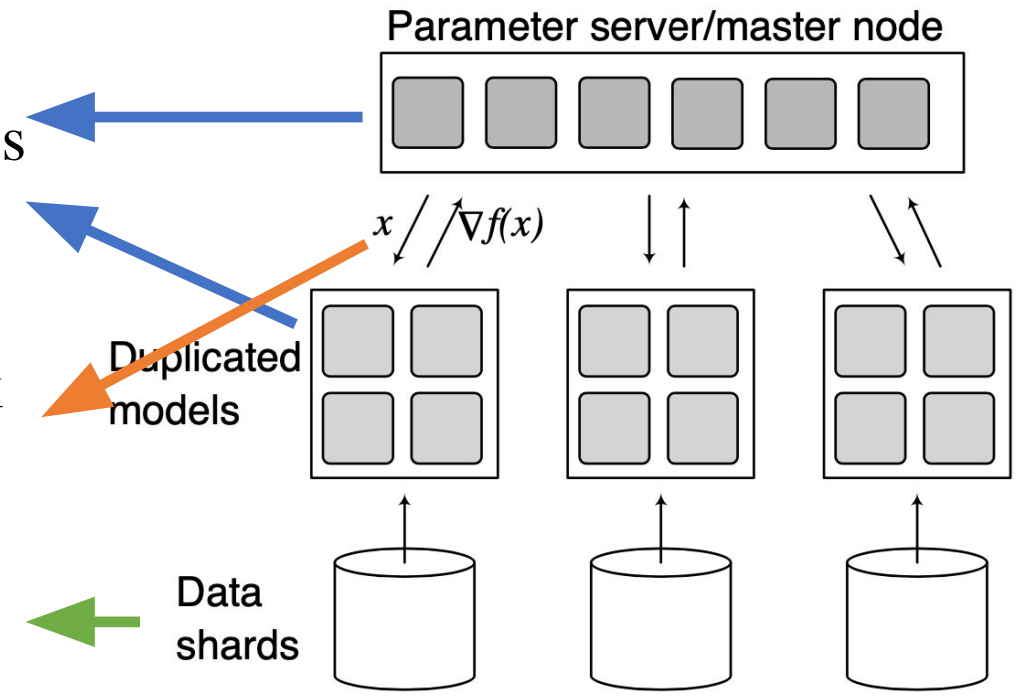

Cloud Service

- Azure Machine Learning Service: workspace, compute instance, compute cluster, and Blob data storage
- Azure CLI Interface, Azure Machine Learning SDK for Python
- Virtual Machines:
  - 16 x NC6 instances (GPU K80 + CPU + Memory + Ethernet network)
  - 8 x NC24r instances (4 x GPU K80 + CPU + Memory + Infiniband network)
- VM images with PyTorch and MPI installed

Azure ML Studio

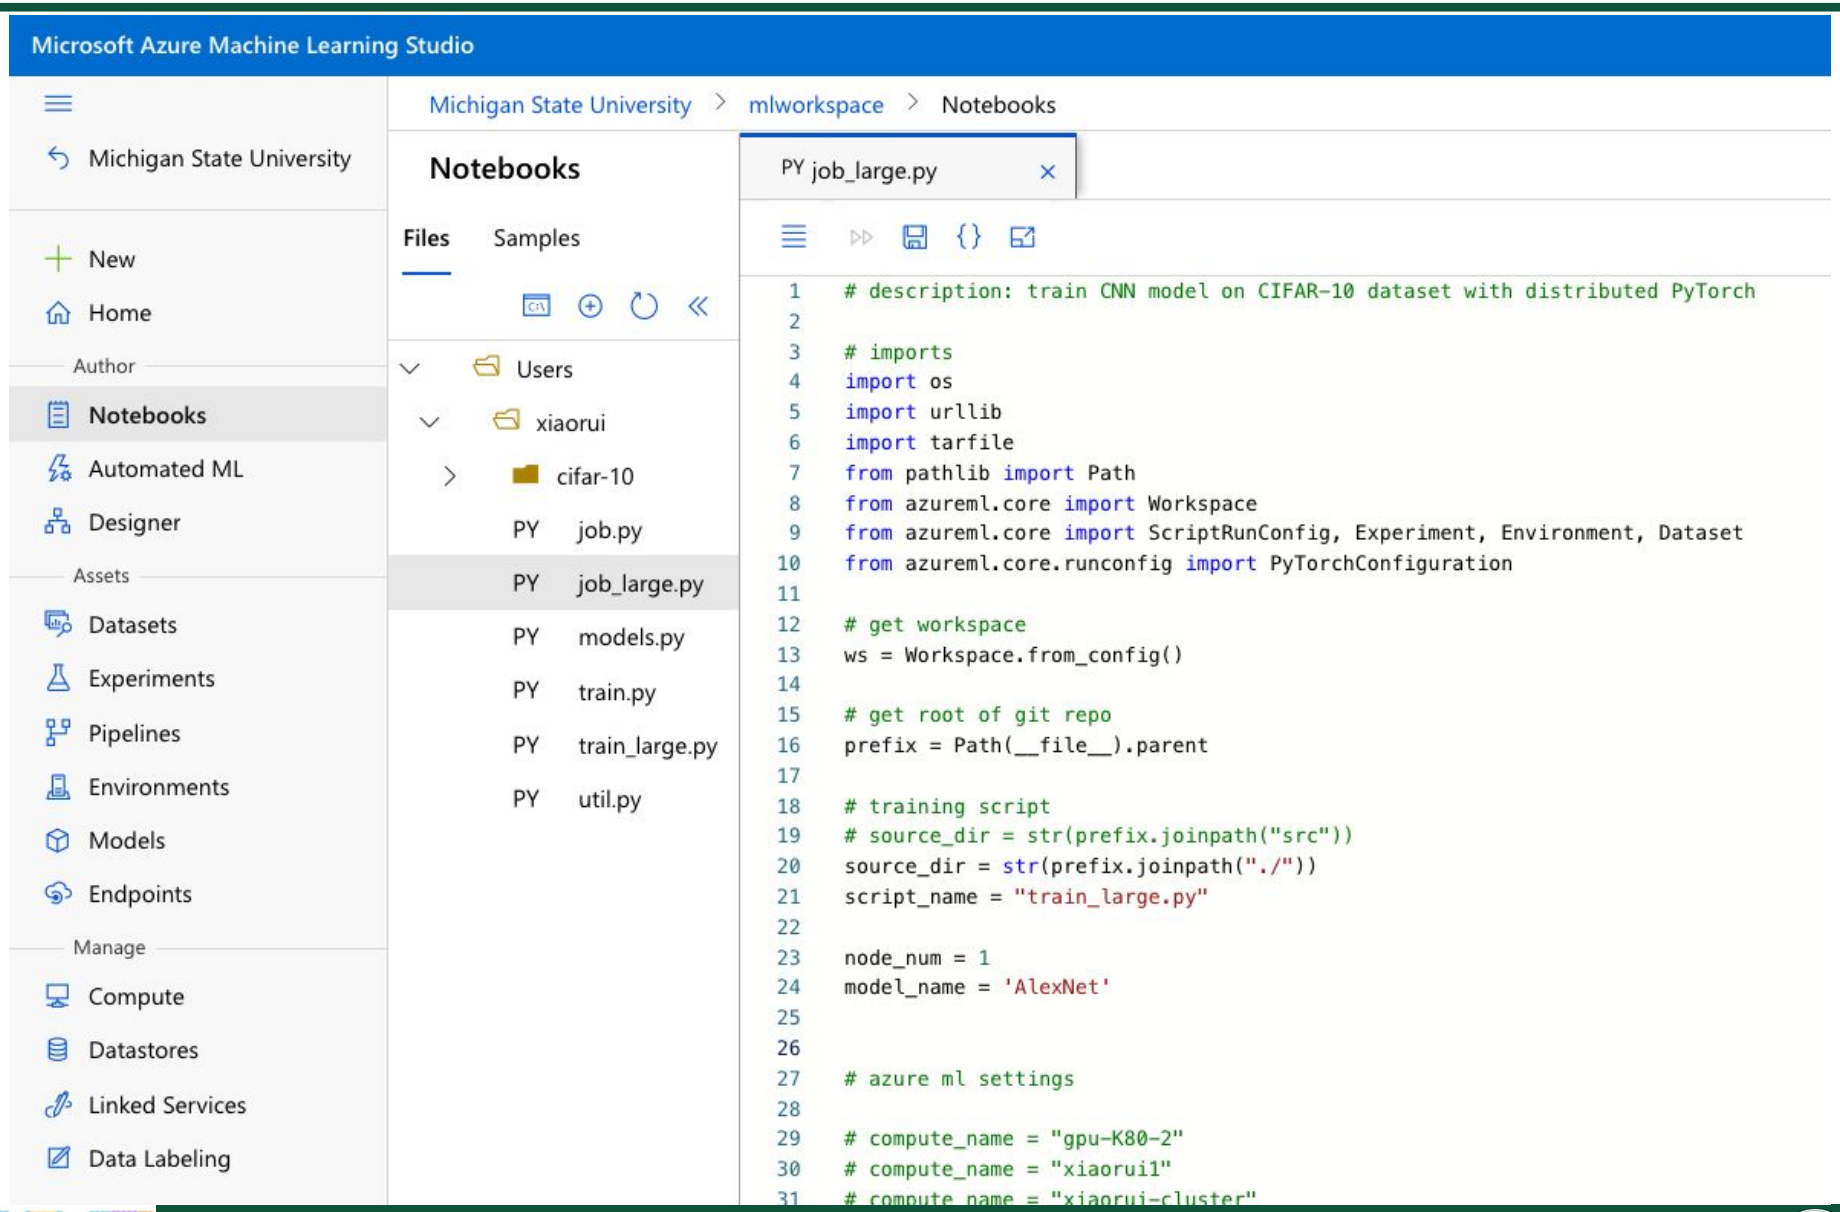

Azure ML Studio

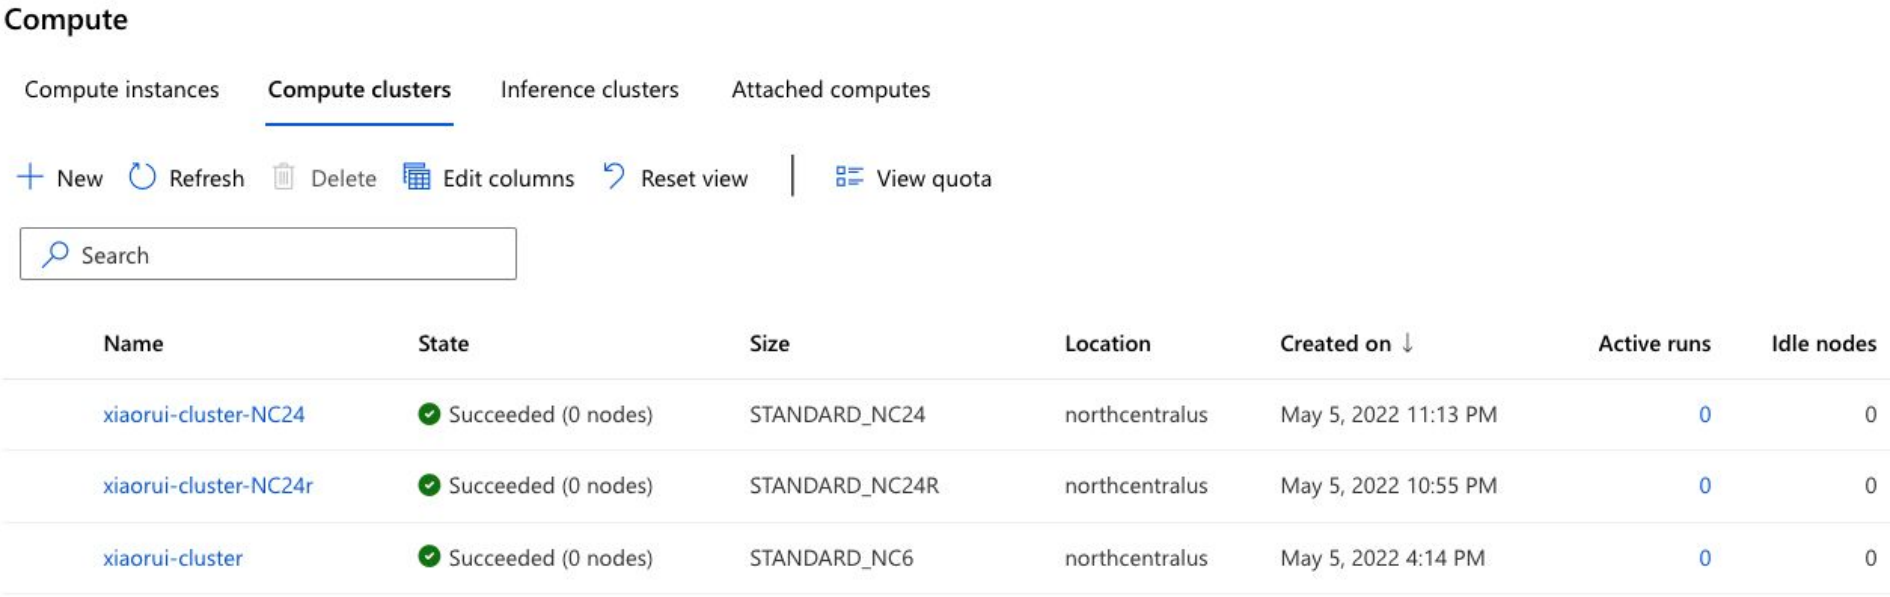

Azure ML Studio

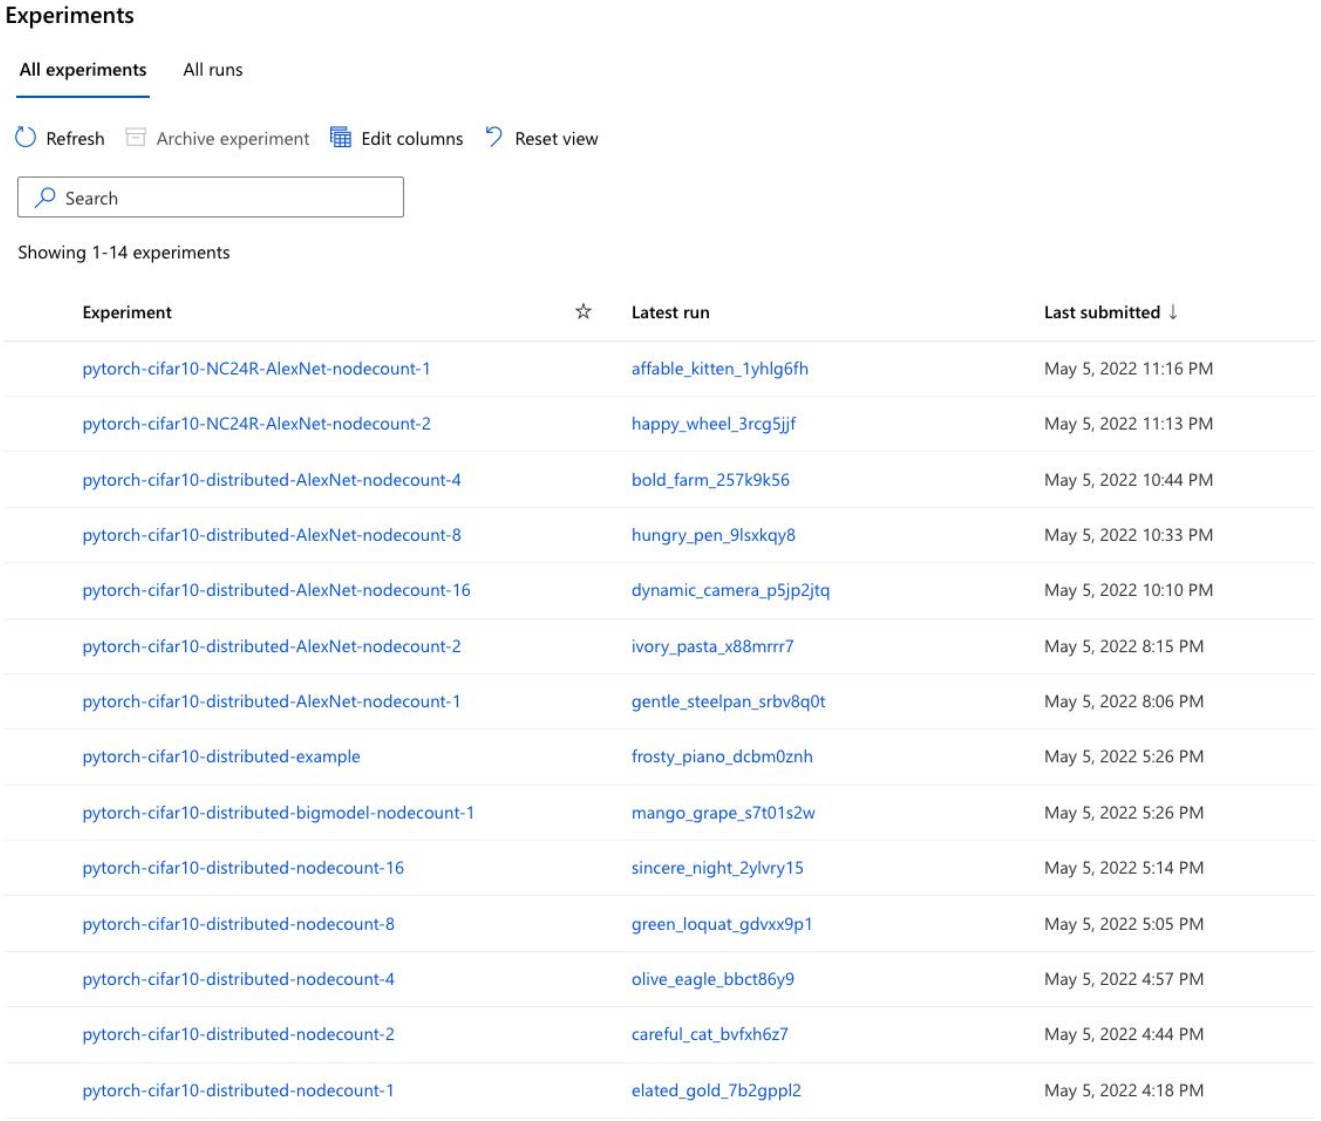

Experiment	Latest run	Last submitted
pytorch-cifar10-NC24R-AlexNet-nodecount-1	affable_kitten_1yhlg6fh	May 5, 2022 11:16 PM
pytorch-cifar10-NC24R-AlexNet-nodecount-2	happy_wheel_3rcg5jjf	May 5, 2022 11:13 PM
pytorch-cifar10-distributed-AlexNet-nodecount-4	bold_farm_257k9k56	May 5, 2022 10:44 PM
pytorch-cifar10-distributed-AlexNet-nodecount-8	hungry_pen_9lsxkqy8	May 5, 2022 10:33 PM
pytorch-cifar10-distributed-AlexNet-nodecount-16	dynamic_camera_p5jp2jtq	May 5, 2022 10:10 PM
pytorch-cifar10-distributed-AlexNet-nodecount-2	ivory_pasta_x88mrrr7	May 5, 2022 8:15 PM
pytorch-cifar10-distributed-AlexNet-nodecount-1	gentle_steelpan_srbv8q0t	May 5, 2022 8:06 PM
pytorch-cifar10-distributed-example	frosty_piano_dcbm0znh	May 5, 2022 5:26 PM
pytorch-cifar10-distributed-bigmodel-nodecount-1	mango_grape_s7t01s2w	May 5, 2022 5:26 PM
pytorch-cifar10-distributed-nodecount-16	sincere_night_2ylvry15	May 5, 2022 5:14 PM
pytorch-cifar10-distributed-nodecount-8	green_loquat_gdvxx9p1	May 5, 2022 5:05 PM
pytorch-cifar10-distributed-nodecount-4	olive_eagle_bbct86y9	May 5, 2022 4:57 PM
pytorch-cifar10-distributed-nodecount-2	careful_cat_bvfxh6z7	May 5, 2022 4:44 PM
pytorch-cifar10-distributed-nodecount-1	elated_gold_7b2gppl2	May 5, 2022 4:18 PM

Measurement

Small model: LeNet (<1MB) under Ethernet network

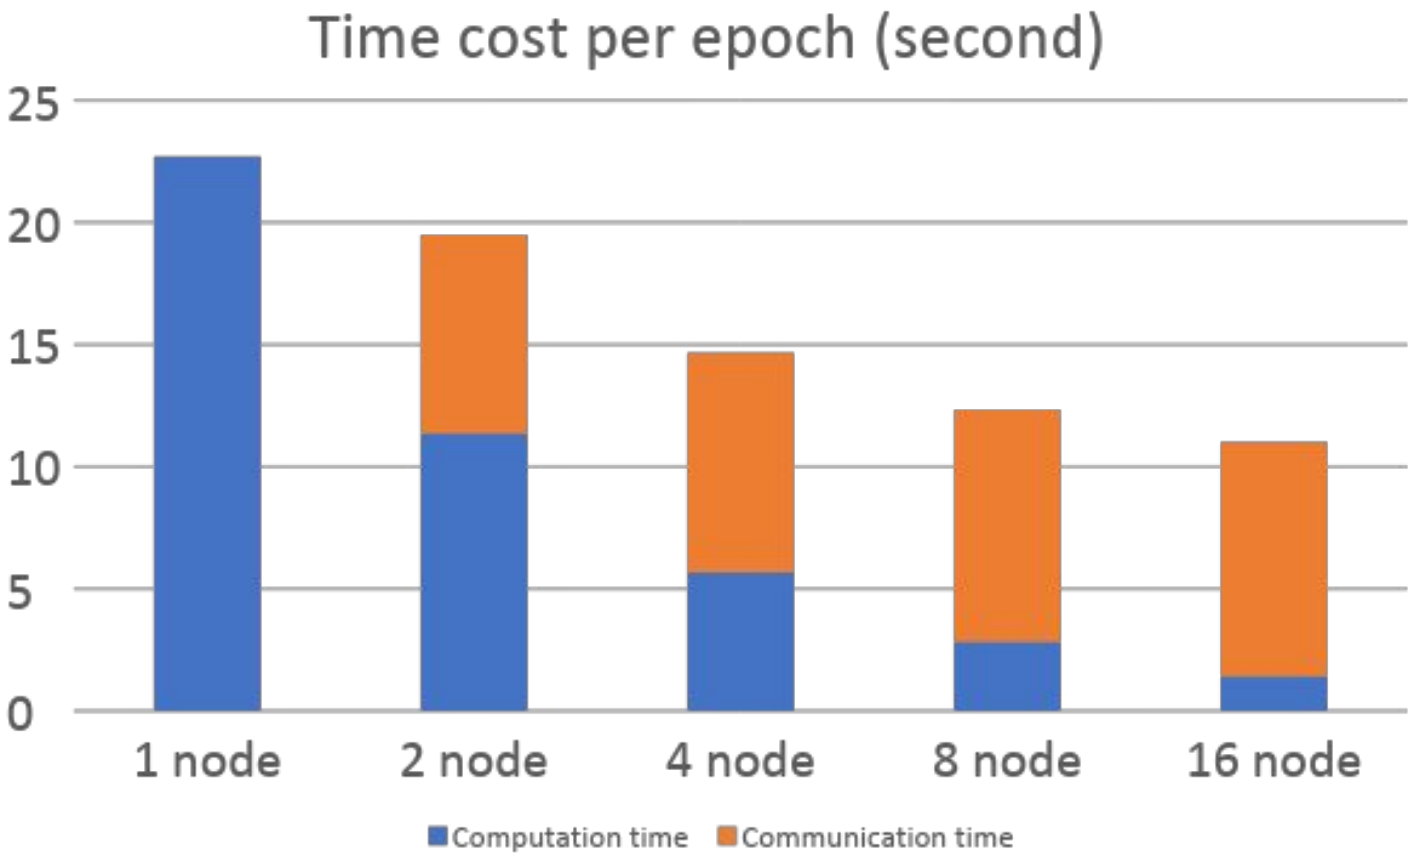

Measurement

Larger model: AlexNet (233 MB) under Ethernet network

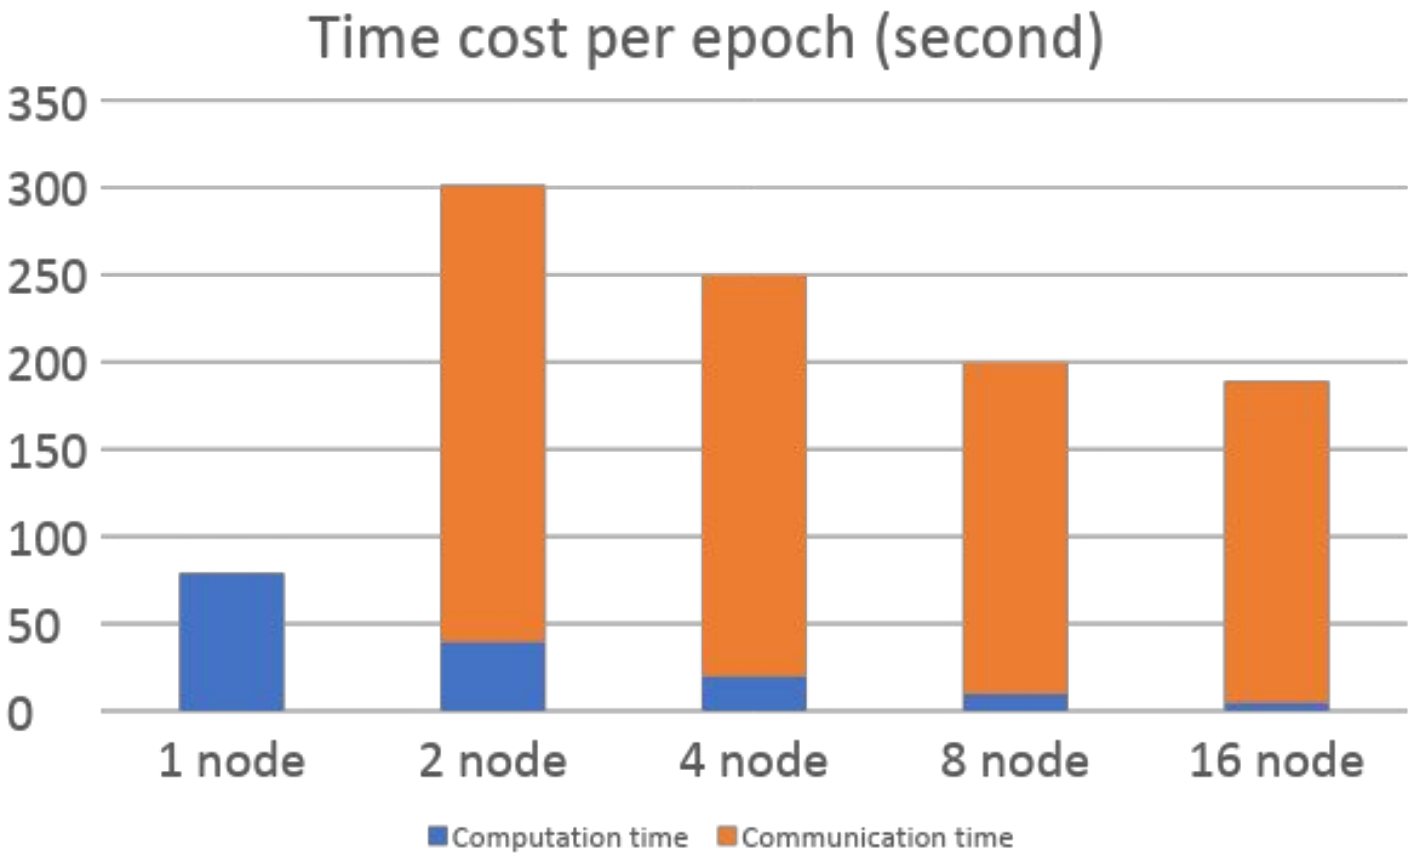

Conclusion

- ***Cloud computing*** provides a versatile platform to benchmark the efficiency of distributed ML (comprehensive documentation)
- ***Small models:*** distributed ML benefits from parallel computing but the speedup is way worse than linear speedup
- ***Large models:*** distributed ML might not benefit from parallel computing but just waste computation power
- Faster communication network (such as ***Infiniband***) is needed
- ***Communication-efficient algorithms*** are needed

Future Research: Communication-Efficient ML