

Large Scale NN Training

Vladimir Loncar, Thong Nguyen, Sofia Vallecora,
Jean-Roch Vlimant

What happened

and what's next

Vladimir

Installed & compiled MPI with cuda supported.

Try with different worker optimizer algorithm for better multi-master architecture performance.

Thong

Implemented the single GPU version for pytorch.

Tested on MNIST dataset, obtained 98.4% accuracy after 9 seconds (was 90s yesterday on CPU).

Implement the pytorch multi-GPU version on single worker node.