




Improving Large-Scale Deep Learning with Multi-Client DeepIO

Yue Zhu[†], Fahim Chowdhury[†], Huansong Fu[†], Adam Moody[‡], Kathryn Mohror[‡], Kento Sato[‡], Weikuan Yu[†]
Florida State University[†], Lawrence Livermore National Laboratory[‡]




Problem Statement




Our Goal: Propose an I/O solution to coordinate the use of memory, communication, and I/O resources for efficient DNN training with larger datasets and deeper neural networks

Existing Solutions




Multi-Client DeepIO




- DeepIO^[2]**: store dataset on memory and randomly fetch elements from storage buffers
 - Three optimization techniques and a TensorFlow API for DeepIO
 - Only support one client per node
- Multi-Client DeepIO**
 - Distinct read buffer for each training worker
 - DeepIO servers submit read requests in a round robin manner
 - Additional support for multi-readers on single node for TensorFlow

Experimental Results




- DeepIO delivers the best aggregate read performance**
 - Compare BeeGFS^[4] and Octopus^[3] on aggregate read bandwidth over 8 nodes



- Performance of alternative TensorFlow data API increases with clients counts**

- Compare raw DeepIO performance with proposed compatible TensorFlow data importing API



- Compare training time when using DeepIO and BeeGFS**

- Average Training Time with CIFAR10
- Mini-batch size: 128 per node

Conclusions & Future Work

- Analyze the existing solutions for storing datasets on HPC system
- Enable the multi-client support for DeepIO, an I/O solution for DNN training on HPC system
- Evaluate DeepIO with more DNNs and datasets
- Compare the results with the regular training over traditional parallel file system

[References]

- [1] Image Source. <https://blog.rescale.com/wp-content/uploads/2016/05/classification.png>.
- [2] Zhu, Yue, et al. "Entropy-Aware I/O Pipelining for Large-Scale Deep Learning on HPC Systems." *Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS), 2018 IEEE 26th International Symposium on*. IEEE, 2018.
- [3] Lu, Youyou, et al. "Octopus: an RDMA-enabled distributed persistent memory file system." *2017 USENIX Annual Technical Conference (USENIX ATC 17)*. 2017
- [4] BeeGFS: <https://www.beegfs.io/content/>