Usage Report for Fiscal Year 2020

Project Title:
Research and Developmemnt of system software for high performance big data applications

Name: Takaaki Fukai (1), ○ Kento Sato (1), Jens Domke (1)
Laboratory at RIKEN:
(1) RIKEN Center for Computational Science, High Performance Big Data Research Team

1. Background and purpose of the project, relationship of the project with other projects

Today, deep learning is an essential technology for our life. For example, deep learning contributes to high-quality machine translation, image classification and recognition, used in mobile devices and auto-motives. Deep learning is also indispensable for science applications. CosmoFlow is an example of large-scale science applications using deep learning.

To solve more complex problems with deep learning, both sizes of training datasets and neural networks are increasing. To train a model with large datasets and networks, a single computer is not enough. Therefore, distributed deep neural network (DDNN) training, using multiple computers for training a model, is necessary.

For large scale DDNN training, HPC clusters are a promising computation environment. Therefore, our next-generation supercomputer, Supercomputer Fugaku, is also expected to be used for the DDNN training. Actually, our project named DL4Fugaku is working on DDNN training on Fugaku, and achieved the 2nd highest score at the previous MLPerf HPC ranking.

In large scale DDNN, I/O performance is a challenging problem. In deep learning training, the order of the training samples to input the network should be randomly changed every epoch. Therefore, each sample is accessed once an epoch, and the order of accessing samples is random. So there is no locality of space and time, and general cache mechanisms do not work well. On the other hand, it is impractical that each node
keeps a huge dataset. Therefore, the large-scale DDNN training applications suffer performance overhead from a large number of remote access on loading data.

2. Specific usage status of the system and calculation method

To explore I/O optimization opportunities for DDNN training workload on Fugaku, we first measure shuffling performance on HOKUSAI instead of Fugaku.

3. Result and Conclusion

In our preliminary results on HOKUSAI, our synthetic benchmark for DDNN shuffling operations reveal, we found DDNN still have potential to improve shuffling operation.

4. Schedule and prospect for the future

Several researchers revealed the bottleneck of I/O performance and proposed methods to improve I/O performance on their clusters. However, Fugaku has the following different points from other systems such as HOKUSAI. (1) Fugaku consists of a large number of small compute nodes (48 cores and 32 GiB HBM2 memory without accelerator), unlike clusters that introduce GPUs; (2) Instead of node-local storage, Fugaku has the 1st level storage between the compute nodes and the global file system; (3) The compute nodes are connected through ToFuD six dimensions mesh-torus network.

We will further evaluation the I/O and shuffling performance on Fugaku while verifying the results with ones on HOKUSAI.

Fiscal Year 2020 List of Publications Resulting from the Use of the supercomputer

[Poster presentation]

- Takaaki Fukai, Kento Sato, “Measurement of I/O performance for distributed deep neural networks on Fugaku”, The 3rd R-CCS International Symposium, Feb, 2021