Tutorial 1

Exploiting HPC for Distributed Deep Learning

Dhabaleswar K, Panda, ***Ohio State University***
Hari Subramoni, ***Ohio State University***
Arpan Jain, ***Ohio State University***
Ammar Ahmad Awan, ***Ohio State University***

Abstract:

The recent advances in Deep Learning (DL) has led to many exciting challenges and opportunities for CS and AI researchers alike. Modern DL frameworks like TensorFlow, PyTorch, and several others have emerged that offer ease of use and flexibility to train, and deploy various types of Deep Neural Networks (DNNs). In this tutorial, we will provide an overview of interesting trends in DNN design and how cutting-edge hardware architectures and high-performance interconnects are playing a key role in moving the field forward. We will also present an overview of different DNN architectures and DL frameworks. Most DL frameworks started with a single-node design. However, approaches to parallelize the process of DNN training are also being actively explored. The DL community has moved along different distributed training designs that exploit communication runtimes like gRPC, MPI, and NCCL. We highlight new challenges and opportunities for communication runtimes to exploit high-performance interconnects and efficiently support large-scale distributed DNN training. We also highlight some of our co-design efforts to utilize CUDA-Aware MPI for large-scale DNN training on GPU clusters. Finally, we include hands-on exercises to enable the attendees to gain first-hand experience of running distributed DNN training experiments on a modern GPU cluster.

Bios:

Dr. Dhabaleswar K. (DK) Panda is a Professor and Distinguished Scholar of Computer Science at the Ohio State University. He obtained his Ph.D. in computer engineering from the University of Southern California. His research interests include parallel computer architecture, high performance networking, InfiniBand, network-based computing, exascale computing, programming models, GPUs and accelerators, high performance file systems and storage, virtualization and cloud computing and BigData (Hadoop (HDFS, MapReduce and HBase) and Memcached). He has published over 400 papers in major journals and international conferences related to these research areas. Dr. Panda has served (or serving) as Program Chair/Co-Chair/Vice Chair of many international conferences and workshops.

Dr. Ammar Ahmad Awan is a researcher at Microsoft. He obtained his PhD from Ohio State University. He received his B.S. and M.S. degrees in Computer Science and Engineering from National University of Science and Technology (NUST), Pakistan and Kyung Hee University (KHU), South Korea, respectively. His research focus lies at the intersection of High Performance Computing (HPC) libraries and Deep Learning (DL) frameworks. He has published 20 papers in conferences and journals related to these research areas.

Dr. Hari Subramoni received a Ph.D. degree in Computer Science from The Ohio State University, Columbus, OH, in 2013. He is a research scientist in the Department of Computer Science and Engineering at the Ohio State University, USA, since September 2015. His current research interests include high performance interconnects and protocols, parallel computer architecture, network-based computing,

Exascale computing, network topology aware computing, QoS, power-aware LAN-WAN communication, fault tolerance, virtualization, big data, and cloud computing. He has published over 50 papers in international journals and conferences related to these research areas. Recently, Dr. Subramoni is doing research and working on the design and development of MVAPICH2, MVAPICH2-GDR, and MVAPICH2-X software packages. He is a member of IEEE. More details about Dr. Subramoni are available from http://www.cse.ohio-state.edu/~subramon.

Arpan Jain is a Ph.D. student in Computer Science and Engineering department at The Ohio State University, Columbus, Ohio. His research interests include deep learning, high-performance computing, GPU communication, and wireless networks. More details are available at http://nowlab.cse.ohio-state.edu/member/jain.575.

Tutorial 2

InfiniBand, High-speed Ethernet, Omni-Path, EFA and Slingshot

Dhabaleswar K, Panda, ***Ohio State University***
Hari Subramoni, ***Ohio State University***

Abstract:

InfiniBand (IB), Omni-Path, and High-Speed Ethernet (HSE) technologies are generating a lot of excitement toward building next-generation High-End Computing (HEC) systems including clusters, datacenters, file systems, storage, cloud computing, and Big Data (Hadoop, Spark, HBase, and Memcached) environments. RDMA over Converged Enhanced Ethernet (RoCE) technology is also being widely deployed.

This tutorial will provide an overview of these emerging technologies, their offered architectural features, their current market standing, and their suitability for designing HEC systems. It will start with a brief overview of IB, Omni-Path, and HSE. An in-depth overview of the architectural features of IB, Omni-Path, and HSE (including iWARP and RoCE), their similarities and differences, and the associated protocols will be presented. Next, an overview of the OpenFabrics stack which encapsulates IB, HSE, and RoCE (v1/v2) in a unified manner will be presented. An overview of libfabrics stack will also be provided. An overview of the emerging NVLink, NVLink2, and NVSwitch architectures will also be given. Hardware/software solutions and the market trends behind IB, Omni-Path, HSE, and RoCE will be highlighted. Finally, sample performance numbers of these technologies and protocols for different environments will be presented.

Bios:

Dr. Dhabaleswar K. (DK) Panda is a Professor and Distinguished Scholar of Computer Science at the Ohio State University. He obtained his Ph.D. in computer engineering from the University of Southern California. His research interests include parallel computer architecture, high performance networking, InfiniBand, network-based computing, exascale computing, programming models, GPUs and accelerators, high performance file systems and storage, virtualization and cloud computing and BigData (Hadoop (HDFS, MapReduce and HBase) and Memcached). He has published over 400 papers in major journals and international conferences related to these research areas. Dr. Panda has served (or serving) as Program Chair/Co-Chair/Vice Chair of many international conferences and workshops.

Dr. Hari Subramoni received a Ph.D. degree in Computer Science from The Ohio State University, Columbus, OH, in 2013. He is a research scientist in the Department of Computer Science and Engineering at the Ohio State University, USA, since September 2015. His current research interests include high performance interconnects and protocols, parallel computer architecture, network-based computing, Exascale computing, network topology aware computing, QoS, power-aware LAN-WAN communication, fault tolerance, virtualization, big data, and cloud computing. He has published over 50 papers in international journals and conferences related to these research areas. Recently, Dr. Subramoni is doing research and working on the design and development of MVAPICH2, MVAPICH2-GDR, and MVAPICH2-X software packages. He is a member of IEEE. More details about Dr. Subramoni are available from http://www.cse.ohio-state.edu/~subramon.