

5th IEEE International Workshop on High-Performance Big Data and Cloud Computing

Welcome to HPBDC 2019

Managing and processing large volumes of data, or Big Data, and gaining meaningful insights is a significant challenge facing the parallel and distributed computing community. This has significant impact in a wide range of domains including health care, bio-medical research, Internet search, finance and business informatics, and scientific computing. As data-gathering technologies and data sources witness an explosion in the amount of input data, it is expected that in the future massive quantities of data in the order of hundreds or thousands of petabytes will need to be processed. Thus, it is critical that data-intensive computing middleware (such as Hadoop, Spark, Flink, etc.) to process such data are diligently designed, with high performance and scalability, in order to meet the growing demands of such Big Data applications.

The explosive growth of Big Data has caused many industrial firms to adopt High Performance Computing (HPC) technologies to meet the requirements of huge amount of data to be processed and stored. The convergence of HPC, Big Data, and Deep Learning is becoming the next game-changing business opportunity. Apache Hadoop, Spark, gRPC/TensorFlow, and Memcached are becoming standard building blocks in handling Big Data oriented processing and mining.

Modern HPC bare-metal systems and Cloud Computing platforms have been fueled with the advances in multi-/many-core architectures, RDMA-enabled networking, NVRAMs, and NVMe-SSDs during the last decade. However, Big Data and Deep Learning middleware (such as Hadoop, Spark, Flink, and gRPC) have not embraced such technologies fully. These disparities are taking HPC, Big Data, and Deep Learning into divergent trajectories.

International Workshop on High-Performance Big Data, Deep Learning, and Cloud Computing (HPBDC), aims to bring HPC, Big Data processing, Deep Learning, and Cloud Computing into a convergent trajectory. The workshop provides a forum for scientists and engineers in academia and industry to present their latest research findings in major and emerging topics for 'HPC + Big Data + Deep Learning over HPC Clusters and Clouds'.

Program Chairs

Xiaoyi Lu, The Ohio State University

Jianfeng Zhan, Institute of Computing Technology, Chinese Academy of Sciences, China

Dhabaleswar K. (DK) Panda, The Ohio State University

Program Committee

Luiz F. Bittencourt, University of Campinas, Brazil

Yong Chen, Texas Tech University

Ada Gavrilovska, Georgia Tech

Shadi Ibrahim, Inria, France

Zengxiang Li, Institute Of High Performance Computing, Singapore

Suzanne McIntosh, New York University

Sergey Maidanov, Intel

Raghunath Nambiar, AMD

Manoj Nambiar, Tata Consultancy Services Ltd., India

Judy Qiu, Indiana University

Tor Skeie, University of Oslo, Norway

Juan Touriño, University of A Coruña, Spain

Li Zha, Institute of Computing Technology, Chinese Academy of Sciences, China

Yunquan Zhang, Institute of Computing Technology, Chinese Academy of Sciences, China