

# Accelerating Big Data with Hadoop and Memcached Using High Performance Interconnects: Opportunities and Challenges

Dhabaleswar K. (DK) Panda and Xiaoyi Lu  
The Ohio State University

## Abstract

Apache Hadoop is gaining prominence in handling Big Data and analytics. Similarly, Memcached in Web 2.0 environment is becoming important for large-scale query processing. These middleware are traditionally written with sockets and do not deliver best performance on modern clusters with high-performance interconnects. In this tutorial, we will provide an in-depth overview of the architecture of Hadoop components (HDFS, MapReduce, HBase, RPC, etc.) and Memcached. We will examine the challenges in re-designing the networking and I/O components of these middleware with modern interconnects and protocols (such as InfiniBand, iWARP, RoCE, and RSocket) with RDMA. Using the publicly available Hadoop-RDMA (<http://hadoop-rdma.cse.ohio-state.edu>) software package, we will provide case studies of the new designs for several Hadoop components and their associated benefits. Through these case studies, we will also examine the interplay between high performance interconnects, storage systems (HDD and SSD), and multi-core platforms to achieve the best solutions for these components.

## Outline

1. Introduction to Big Data Applications and Analytics
2. Overview of Apache Hadoop Architecture and its Components
  - HDFS
  - MapReduce
  - HBase
  - RPC
3. Overview of Web 2.0 Architecture and Memcached
4. Overview of Benchmarks and Applications using Hadoop and Memcached
5. Overview of High Performance Interconnects and Protocols
6. Challenges in Accelerating Hadoop and Memcached

7. Acceleration Case Studies and In-Depth Performance Evaluation
  - HDFS
  - MapReduce
  - HBase
  - RPC
  - Memcached
8. Optimizations and Tuning of Accelerated Designs on Modern Clusters
9. Opportunities for Additional Enhancements and Accelerations
10. Conclusions and Q&A

### **Brief Bio of the Speakers**

Dhabaleswar K. (DK) Panda is a Professor of Computer Science and Engineering at the Ohio State University. His research interests include parallel computer architecture, high performance networking, InfiniBand, Exascale computing, Big Data, programming models, GPUs and accelerators, high performance file systems and storage, virtualization, and cloud computing. He has published over 300 papers in major journals and international conferences related to these research areas. Dr. Panda and his research group members have been doing extensive research on modern networking technologies including InfiniBand, High-Speed Ethernet, and RDMA over Converged Enhanced Ethernet (RoCE). The MVAPICH2 (High Performance MPI over InfiniBand, iWARP and RoCE) and MVAPICH2-X (Hybrid MPI and PGAS (OpenSHMEM and UPC)) software packages, developed by his research group (<http://mvapich.cse.ohio-state.edu>), are currently being used by more than 2,055 organizations worldwide (in 70 countries). This software has enabled several InfiniBand clusters to get into the latest TOP500 ranking during the last decade. More than 177,000 downloads of this software have taken place from the project's website alone. This software package is also available with the software stacks of many network and server vendors and Linux distributors. Recently, Dr. Panda and his team have also developed a high performance Hadoop-RDMA software package (<http://hadoop-rdma.cse.ohio-state.edu>) to accelerate Hadoop with RDMA for Big Data. Dr. Panda's research has been supported by funding from the US National Science Foundation, the US Department of Energy, and several industry companies including Intel, Cisco, SUN, Mellanox, QLogic, NVIDIA, and NetApp. He is an IEEE Fellow and a member of the ACM. More details about Prof. Panda are available at <http://www.cse.ohio-state.edu/~panda>.

Dr. Xiaoyi Lu received the B.S. degree in Electronics and Communications Engineering from Huazhong University of Science and Technology, Wuhan, China, in 2006 and the Ph.D. degree in Computer Science from Institute of Computing

Technology, Chinese Academy of Sciences, Beijing, China, in 2012. He is a post-doctoral researcher in the Department of Computer Science and Engineering at the Ohio State University, USA, since July 2012. His current research interests include high performance interconnects and protocols, Big Data, Hadoop Ecosystem, and Parallel Computing (MPI/PGAS). Recently, Dr. Lu is working with other team members on designing and developing the high performance Hadoop-RDMA software package (<http://hadoop-rdma.cse.ohio-state.edu>). More details about Dr. Lu are available at <http://www.cse.ohio-state.edu/~luxl>.