

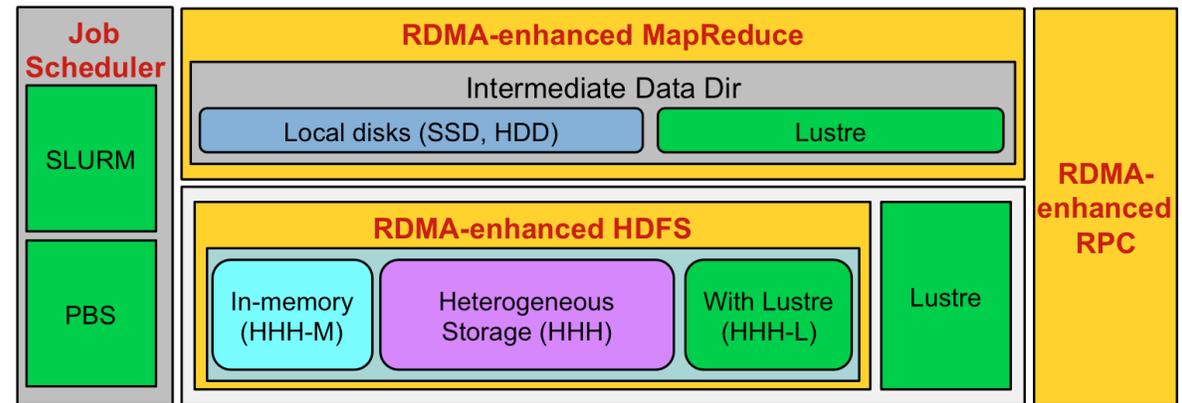


HiBD - High-Performance Big Data

Dhabaleswar K (DK) Panda, The Ohio State University

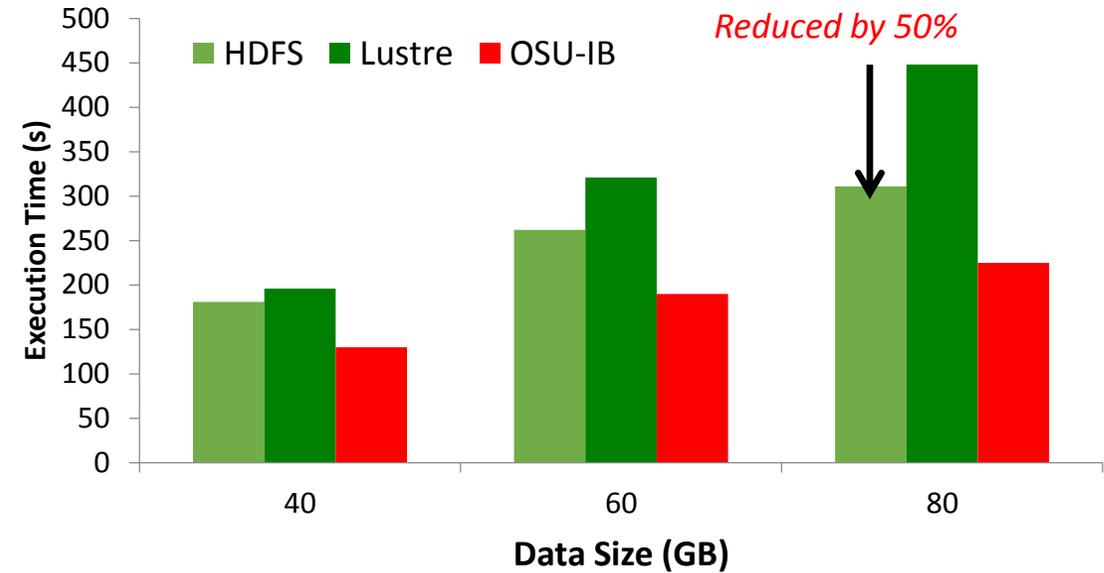
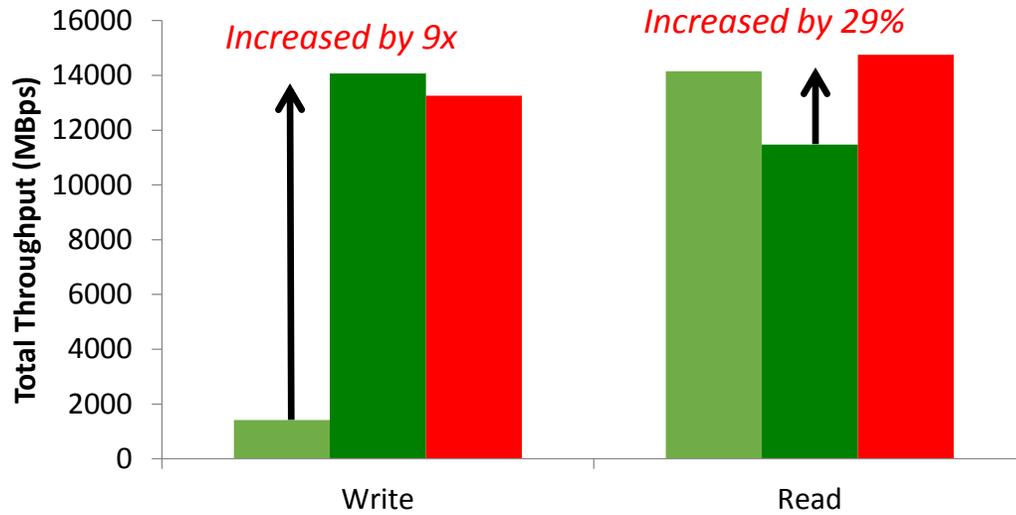
Overview of the HiBD Project for High-Performance Big Data Processing

- Designs exploiting RDMA-based designs for HPC clusters with IB and RoCE
- Leverages parallel file systems like Lustre
- **Installed and available on SDSC Comet**
- **Public releases of various libraries**
 - **RDMA for Apache Hadoop 2.x (RDMA-Hadoop-2.x)**
 - Plugins for Apache and HDP Hadoop distributions
 - **RDMA for Apache Hadoop 1.x (RDMA-Hadoop)**
 - **RDMA for Memcached (RDMA-Memcached)**
 - **OSU HiBD-Benchmarks (OHB)**
 - HDFS and Memcached Micro-benchmarks
- <http://hibd.cse.ohio-state.edu>
- **Users Base: 135 organizations from 20 countries**
- **More than 13,200 downloads from the project site**
- **RDMA for Apache HBase, Spark and CDH (Upcoming)**



- **HHH:** Heterogeneous storage devices with hybrid replication schemes are supported in this mode of operation to have better fault-tolerance as well as performance. This mode is enabled by **default**.
- **HHH-M:** A high-performance in-memory based setup has been introduced in this package that can be utilized to perform all I/O operations in-memory and obtain as much performance benefit as possible.
- **HHH-L:** With parallel file systems integrated, HHH-L mode can take advantage of the Lustre available in the cluster.
- **MapReduce over Lustre, with/without local disks:** Besides, HDFS based solutions, this package also provides support to run MapReduce jobs on top of Lustre alone. Here, two different modes are introduced: with local disks and without local disks.
- **Running with Slurm and PBS:** Supports deploying RDMA for Apache Hadoop 2.x with Slurm and PBS in different running modes (HHH, HHH-M, HHH-L, and MapReduce over Lustre).

Performance Benefits with TestDFSIO and Sort on SDSC-Gordon



- *TestDFSIO for 80GB data size*

- Write: **9x** improvement over HDFS
- Read: **29%** improvement over Lustre

- *Sort*

- up to **28%** improvement over HDFS
- up to **50%** improvement over Lustre

N. Islam, X. Lu, M. W. Rahman, D. Shankar, and D. K. Panda, Triple-H: A Hybrid Approach to Accelerate HDFS on HPC Clusters with Heterogeneous Storage Architecture, CCGrid '15, May 2015

M. W. Rahman, X. Lu, N. S. Islam, R. Rajachandrasekar, and D. K. Panda, High Performance Design of YARN MapReduce on Modern HPC Clusters with Lustre and RDMA, IPDPS, May 2015

X. Lu, N. Islam, M. W. Rahman, J. Jose, H. Subramoni, H. Wang, and D. K. Panda, High-Performance Design of Hadoop RPC with RDMA over InfiniBand, Int'l Conference on Parallel Processing (ICPP '13), October 2013.