High-Performance Hadoop and Spark MapReduce on Modern HPC Systems

Md. Wasi-ur- Rahman, Advisor: Dhabaleswar K. Panda. Email: {rahmanmd, panda} @cse.ohio-state.edu

Done as a part of the High-Performance Big Data (HiBD) project (http://hibd.cse.ohio-state.edu)

MapReduce Programming Model

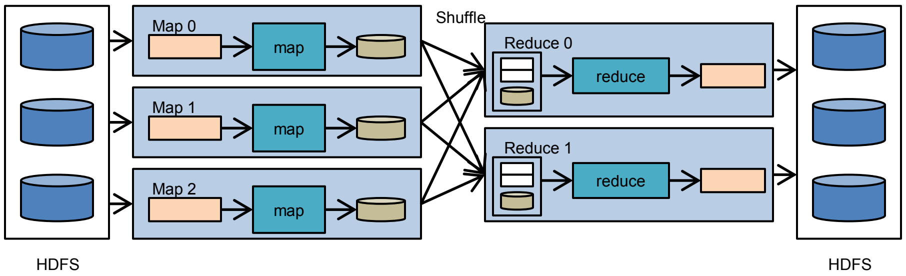

Limitations of MapReduce

- MapReduce implementations (e.g. Hadoop, Spark) use Ethernet infrastructure with Java sockets
- Large number of disk operations for Hadoop
- Concerns for performance and scalability

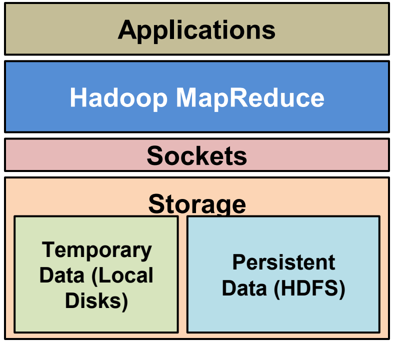

Number of InfiniBand Systems

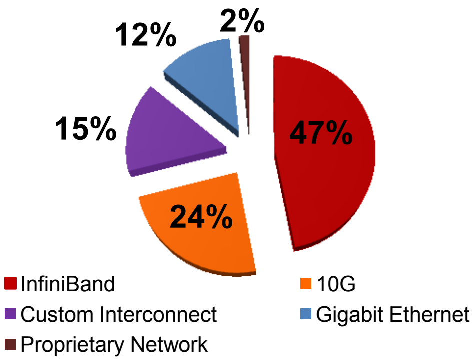

- HPC adopted advanced interconnects and protocols
- Low latency, high bandwidth
- HPC deploys vast parallel file systems, e.g. Lustre
- In-memory based computing is more popular

Broad Challenge

Can MapReduce execution frameworks (Hadoop, Spark) take advantage of HPC resources?

Selected Contributions

- RDMA-enhanced Hadoop MapReduce (ICS'14)
- RDMA-enhanced HDFS (SC'12)
- MapReduce over Luster (TPDS'16)
- Hybrid Hierarchical HDFS (CCGrid'15)
- RDMA-enhanced Spark (HotI'14)

High-Performance Hadoop and Spark MapReduce on Modern HPC Systems

Md. Wasi-ur- Rahman, Advisor: Dhabaleswar K. Panda. Email: {rahmanmd, panda} @cse.ohio-state.edu

Done as a part of the High-Performance Big Data (HiBD) project (http://hibd.cse.ohio-state.edu)

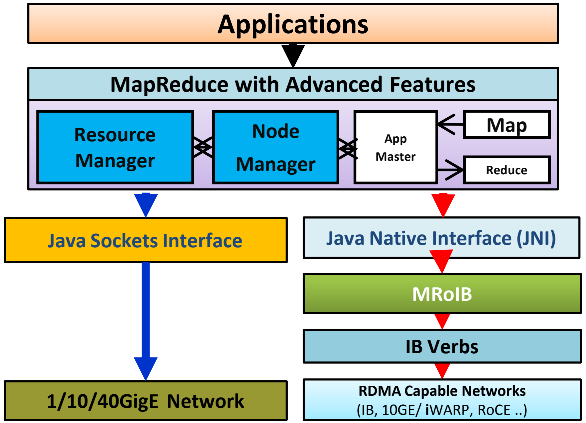

Enhanced phase overlapping

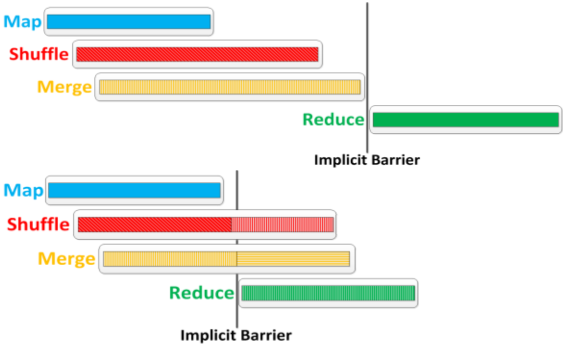

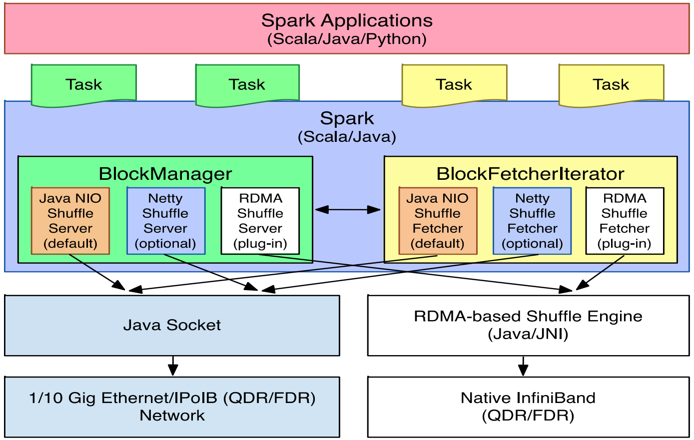

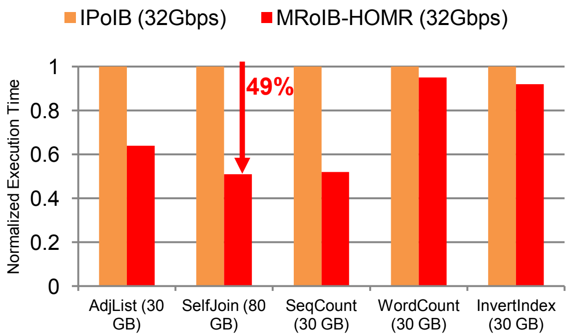

MapReduce over Lustre