Performance analysis of Message Passing Interface collective communication on intel xeon quad-core gigabit ethernet and infiniband clusters

ABSTRACT

The performance of MPI implementation operations still presents critical issues for high performance computing systems, particularly for more advanced processor technology. Consequently, this study concentrates on benchmarking MPI implementation on multi-core architecture by measuring the performance of Open MPI collective communication on Intel Xeon dual quad-core Gigabit Ethernet and InfiniBand clusters using SKaMPI. It focuses on well known collective communication routines such as MPI-Bcast, MPI-AlltoAll, MPI-Scatter and MPI-Gather. From the collection of results, MPI collective communication on InfiniBand clusters had distinctly better performance in terms of latency and throughput. The analysis indicates that the algorithm used for collective communication performed very well for all message sizes except for MPI-Bcast and MPI-Alltoall operation of inter-node communication. However, InfiniBand provides the lowest latency for all operations since it provides applications with an easy to use messaging service, compared to Gigabit Ethernet, which still requests the operating system for access to one of the server communication resources with the complex dance between an application and a network.