

Improved LS-DYNA Parallel Scaling From Fast Collective Communication Operations on High-Performance Compute Clusters

Lars Jonsson and Tim Prince
Intel Corporation

Abstract

Fast collective communications are a key to maintaining high parallel efficiency as the number of nodes increases on a cluster of high-performance servers. Profiling of LS-DYNA message traffic demonstrates that good parallel scaling requires fast communications of short messages - up to a few kilobytes - and in particular of collective operations involving short messages. Fast collective operations require both an efficient implementation of the message-passing operations in terms of message primitives and a high-bandwidth, low-latency interconnect. This paper demonstrates both these aspects by presenting parallel-scaling measurements on Intel Architecture based compute clusters with MPICH2 implemented over fast interconnects. The analysis evaluates both the benefits, at application level, of the emerging MPICH2 work from Argonne National Laboratories relative to MPICH1, and the benefits from the single-digit microsecond latencies offered by today's fastest interconnects. The paper also outlines how next-generation interconnect technologies and new, efficient, and flexible MPI implementations can even further improve both application performance and adaptability.

