

Topic 13

High-Performance Networks

Thilo Kielmann, Pascale Primet, Tomohiro Kudoh, and Bruce Lowekamp

Topic Chairs

Communication networks are likely among the most crucial resources for parallel and distributed computer systems. The last couple of years has seen many significant technical improvements, ranging from advanced on-chip interconnects, via system and storage-area networks, up to dedicated, optical wide-area networks. While this technology push is bringing new opportunities, it also comes with many new research challenges.

This topic is devoted to all kinds of communication issues in scalable compute and storage systems, such as parallel computers, networks of workstations, and clusters. In total, 13 papers were submitted to this topic out of which we have selected the four strongest ones.

Two of the accepted papers are dealing with aspects interconnection networks. In “Integrated QoS Provision and Congestion Management for Interconnection Networks”, the authors present a new switch architecture that uses the same resources for both purposes, which is shown to be as effective as previous proposals, but much more cost effective. In “Performance Analysis of an Optical Circuit Switched Network for Peta-Scale Systems”, the authors compare the performance of an OCS network, as expected in a peta-scale system, to more traditional networks. They show that an OCS network is comparable with the best networks known so far, however with better cost and flexibility.

The other two accepted papers are dealing with implementing message passing (MPI) libraries on top of modern networks. In “Fast and Efficient Total Exchange on Two Clusters”, the authors propose a new algorithm for the collective all-to-all exchange operation across multiple compute clusters, resulting in large performance improvements. Last but not least, “Network Fault Tolerance in Open MPI” describes methods for handling several kinds of network errors while maintaining high-performance communications.