Editorial

A.Y. Al-Dubai and M. Ould-Khaoua

Research on group communications has received considerable attention over the past years as they place a high demand on network bandwidth and have a consequent impact on applications execution time. The design of efficient group communication algorithms and their performance evaluation are of immense importance to ensure their successful introduction into practical large-scale parallel, distributed and networked systems. We have six papers selected out of 14 papers, submitted in response to our CFP for a special issue on the 'Design and Performance Evaluation of Group Communication in Parallel and Distributed Systems'. The purpose of this special issue is to report on the recent research and development of innovative group communication schemes and the techniques proposed to analyse their performance behaviour under various operating conditions. The contributions of the six papers are outlined below.

One popular approach to improving the implementation of MPI collective operations is to use intelligent or programmable network interfaces to offload the burden of communication activities from the host processor(s). Such implementations have shown significant improvement for micro benchmarks that isolate group communication performance, but these results have not been shown to translate to significant increases in performance for real applications. In the first paper, Brightwell et al. describe several characteristics of applications and application benchmarks that impact group communication performance. They then analyse network resource usage data in order to guide the design of collective offload engines and their associated programming interfaces. In particular, they provide an analysis of the potential benefit of non-blocking group communication operations for MPI.

IP multicast is becoming increasingly important for a wide range of applications such as IP telephony and video-conferencing. However, IP multicast is still far from being widely deployed in the internet, where is facing the explosive increasing of traffic. There are many issues that have delayed its deployment, including scalability, address allocation, billing and security, etc. In the second paper, Al-Begain and Cao propose a new source specific multicast protocol, called Scalable Recursive Explicit Multicast (SREM). SREM aims to address both scalability problem and high join/leave latency in existing multicast schemes as a step towards providing a solution for application in mobile IP networks.

Efficient algorithms for reduction operations across a group of processes are crucial for good performance in many large-scale, parallel scientific applications. While previous algorithms limit processing to the host CPU, Petrini et al. utilise the programmable processors and local memory available on modern cluster Network Interface Cards (NICs) to explore a new dimension in the design of reduction algorithms. They present the benefits and challenges, design issues and solutions, analytical models, and experimental evaluations of a family of NIC-based reduction algorithms.

Multicast is among the most important group communication. In the fourth paper, Barth, Cohen and Durbach investigate a theoretical problem related to communication resource allocation to answer multicast requirements for distributed interactive simulations. They present a new approach, which considers the issue as a multipoint communication problem for which they provide some complexity results with associated lower bounds, and present also two polynomial-time heuristics.

Broadcast communication refers to the delivery of the same message, originating from a given source, to all network nodes. Broadcasting has for long been among the most important topics in the field of parallel and distributed computing due to its increasing importance in many real-world parallel applications. Unfortunately, most existing algorithms have been studied within limited conditions, such as light traffic load and fixed network sizes. In the fifth paper, Al-Dubai, Ould-Khaoua and Mackenzie present an extensive comparative study among a number of broadcast algorithms, taking into account the scalability, parallelism, and low latency to some typical reflect QoS issues for broadcasting in wormhole-switched meshes.

OTIS-based multicomputers are parallel systems that can efficiently exploit optical connections between distant nodes while using electronic connections for physically close processors. There has not been a study that has attempted to determine an optimal multicast communication scheme for OTIS networks. In the last paper, Najaf-abadi and Sarbazi-azad present two approaches to implement efficient multicast communication in wormhole-routed OTIS networks, for two different possible unicast routing schemes on these networks in the absence of hardware support.

Finally, as guest co-editors of this special issue, we would like to express our deepest thanks to the Editor-in-Chiefs, Professor Y. Pan and Professor L.T. Yang, for hosting this special issue in the *International Journal of High-Performance Computing and Networking* (*IJHPCN*) and for their continued support and helpful guidance throughout all the stages of preparing this special issue. Our sincere thanks also go to the Editorial-office staff of the journal for their excellent job during the course of preparing this special issue. We also thank the authors for their contributions, including those whose papers were not included. We thank and greatly appreciate the thoughtful work of the many reviewers who provided invaluable evaluations and recommendations.