Topic 13: High Performance Network and Communication

Chris Develder, Emmanouel Varvarigos,
Admela Jukan, and Dimitra Simeonidou

Topic Committee

This topic on *High-Performance Network and Communication* is devoted to communication issues in scalable compute and storage systems, such as parallel computers, networks of workstations, and clusters. All aspects of communication in modern systems were solicited, including advances in the design, implementation, and evaluation of interconnection networks, network interfaces, system and storage area networks, on-chip interconnects, communication protocols, routing and communication algorithms, and communication aspects of parallel and distributed algorithms.

This year we selected 4 papers and one of the papers was selected for the best paper session. The authors of the paper entitled "*Topology Configuration in Hybrid EPS/OCS Interconnects*" consider a hybrid interconnection network for future high performance computing (HPC) and data center systems, combining Electronic Packet Switching (EPS) and Optical Circuit Switching (OCS). They present heuristic algorithms to map the tasks of parallel HPC applications onto physical processors to allow efficient communication over the hybrid EPS/OPS network. The article entitled "*Towards an Efficient Fat-tree like Topology*" proposes extensions of the fat-tree topology and analyzes the way the routing algorithm affects the complexity of the switch. The authors not only look into the impact of topology and routing on performance (i.e., throughput, latency etc.), but also their influence on switch cost (e.g., in terms of number of switching elements required). In "*An adaptive, scalable, and portable technique for speeding up MPI-based applications*", the authors present a portable optimization of MPI called PRAcTICaL-MPI (Portable Adaptive Compression Library - MPI). PRAcTICaL-MPI enhances the performance and scalability of MPI applications by applying run-time lossless compression, in a transparent way for applications, to reduce the data volume exchanged among processes, selecting the most appropriate compression algorithm at run-time. Finally, the authors of the paper "*Cost-effective Contention Avoidance in a CMP with Shared Memory Controllers*" study, for large chip multiprocessors (CMP), the cause and effect of network congestion due to traffic local to the applications, and traffic caused by memory access. They present a mechanism to reduce head-of-line blocking in the switches, hence efficiently reducing network congestion, increasing network performance, and evening out performance differences between CMP applications.

C. Kaklamanis et al. (Eds.): Euro-Par 2012, LNCS 7484, p. 700, 2012.