

Introduction

Jesper Larsson Träff, Brice Goglin, Ulrich Bruening, and Fabrizio Petrini

Topic chairs

The Euro-Par topic on high-performance networks and communications is devoted to communication issues in scalable compute and storage systems, such as tightly coupled parallel computers, clusters, and networks of workstations, including hierarchical and hybrid designs featuring several levels of possibly different interconnects. All aspects of communication in modern compute and storage systems are of interest, for example advances in the design, implementation, and evaluation of interconnection networks, network interfaces, system and storage area networks, on-chip interconnects, communication protocols and interfaces, routing and communication algorithms, and communication aspects of parallel and distributed algorithms.

The papers submitted for the topic were reviewed by the four chairs and their selected subreviewers. Bar two (who had three reviews), the papers received 4, hopefully useful reviews (that in some cases admittedly could have been more extensive). The several submitted papers all fitted well to the call for papers as outlined above and the specific list of themes. Based on the reviews, the quality aspirations and the overall balance of Euro-Par, only two contributions were accepted for presentation at the conference, making for a selective topic. The topic chairs thank all submitting authors, the presenters, and the audience who will be listening and participating in the discussions. High-quality submissions to the topic also in the coming years are encouraged.

The first paper presented at the conference titled *Kernel-Based Offload of Collective Operations - Implementation, Evaluation and Lessons Learned* by Timo Schneider, Sven Eckelmann, Torsten Hoeffler and Wolfgang Rehm deals with issues in offloading collective communication algorithms to the communication network layer. A kernel-based architecture for implementing a framework for offloading such algorithms is described, implemented and experimentally evaluated with specific microbenchmarks on a standard cluster by comparing to traditional implementations of (non-blocking) collective operations with progress in user-space. Especially reduced CPU overhead and improvement in the capability to overlap communication with computation are shown.

The second paper by Alexandre Denis on *A High Performance Superpipeline Protocol for InfiniBand* discusses improved pipeline schemes for point-to-point communication, and in particular gives a more detailed analysis than usual in a *LogP* inspired performance model. Benchmarks show that in particular the costs of memory registration can be eliminated, making for significantly better performance even on “first touch” of a user-space communication buffer.