

Topic 03

Scheduling and Load Balancing

Bettina Schnor

Local Chair

Scheduling and load balancing techniques are key issues for the performance of parallel applications. However, a lot of problems regarding, for example, dynamic load balancing are still not sufficiently solved.

Hence, many research groups are working on this field, and we are glad that this topic presents several excellent results. We want to mention only a few ones from the 13 papers (1 distinguished, 8 regular, and 4 short papers), selected from 28 submissions.

One session is dedicated to system level support for scheduling and load balancing. In their paper *The Impact of Migration on Parallel Job Scheduling for Distributed Systems*, Zhang, Franke, Moreira, and Sivasubramaniam show how back-filling gang scheduling may profit from an additional migration facility. Leinberger, Karypis, and Kumar present in *Memory Management Techniques for Gang Scheduling* a new gang scheduling algorithm which balances the workload not only due to processor load, but also due to memory utilization.

The wide range of research interests covered by the contributions to this topic is illustrated by two other interesting papers. *Load Scheduling Using Performance Counters* by Lindenmaier, McKinley, and Temam presents an approach for extracting fine-grain run-time information for hardware counters to improve instruction scheduling. Gursoy and Atun investigate in *Neighbourhood Preserving Load Balancing: A Self-Organizing Approach* how Kohonen's self-organizing maps can be used for static load balancing.

Further, there are several papers dealing with application-level scheduling in Topic 3. One of these, *Parallel Multilevel Algorithms for Multi-Constraint Graph Partitioning* by Schloegel, Karypis, and Kumar was selected as distinguished paper. It investigates the load balancing requirements of multi-phase simulations.

We would like to thank sincerely the more than 40 referees that assisted us in the reviewing process.