Scheduling and Load Balancing

Ramin Yahyapour[1], Raffaele Perego[1], Frédéric Desprez[2],
Leah Epstein[2], and Francesc Guim Bernat[2]

[1] Topic Chairs
[2] Members

Scheduling and load balancing techniques are crucial for implementing efficient parallel and distributed applications and for making best use of parallel and distributed systems. This includes planning and optimization of resource allocation as well as coping with the dynamics of the systems.

They have been subject for research for many decades but remain one of the core topics as the infrastructures evolved during the years as well as the approach to scheduling. These techniques can be provided either at the application level or at the system level, and both approaches are of special interest. At the application level, the mapping of distributed and parallel applications on to infrastructures an the development of dynamic load balancing algorithms that are able to exploit the particular characteristics and the actual utilization of the underlying system are of particular relevance. At the system level, areas of interest now include the support of modern many-core architectures as well as virtual systems like Cloud infrastructures. The optimization goals became more sophisticated by including now several criteria beyond the minimization of execution times and increasing utilization such as energy minimization and quality of service.

This year's iteration of the scheduling and load-balancing topic provides a very good coverage of these different perspectives. In the following, you will find contributions which focus on theoretical aspect or on practical questions. Similarly, we can see scheduling on application and on system level; locally, clustered, and distributed. To this end, several papers address hierarchical scheduling which is a common trend considering the different optimization aspects for multi-core systems, parallel clusters and distributed systems like Grids or Clouds. This also includes the coping with complexities due to the lack of central control and the support of autonomous management on individual compute sites.

In this year, the selection process for this topic area was highly competitive. A total of 29 papers have been submitted. All papers were reviewed by at least three independent reviewers. Eventually, only 8 papers were accepted for publication and presentation for Euro-Par 2010. This yields an acceptance rate of 27 %. The selection process was not easy as many papers provided very good research insights and interesting approaches.

We would like to thank all the reviewers, for their time and effort, who helped us in the selection process. At the same time, we would like to thank all authors who help to maintain Euro-Par as one of the premier scientific conferences at which innovative scheduling concepts for parallel and distributed systems are presented.

P. D'Ambra, M. Guarracino, and D. Talia (Eds.): Euro-Par 2010, Part I, LNCS 6271, p. 156, 2010.