

Introduction

Emmanuel Jeannot*, Ramin Yahyapour*, Daniel Grosu*, and Helen Karatza*

Scheduling and load balancing are key components of any resource management system. They are responsible for making the best use of available resources by planning, allocating and redistributing computational tasks onto these resources. Typical optimization criteria are the minimization of execution times and response times. However, additional criteria like power consumption, quality-of-service requirements or even multi-criteria problems have recently been investigated.

Scheduling and load balancing have been active areas of research throughout decades. These techniques are crucial for implementing efficient parallel and distributed applications. As parallel and distributed systems are evolving rapidly, new scheduling problems appear and novel solutions are sought which can be applied to these scenarios.

In short, the goal of this topic is to serve as a forum where efficient and robust scheduling, load-balancing, and/or resource management algorithms are proposed for modern parallel and distributed systems such as clusters, grids, and global computing platforms.

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. Finally, only 9 papers have been accepted which addressed new and exciting challenges and presented innovative solutions. The accepted papers range from theory to practice and address various problems such as energy-aware scheduling, dynamic load balancing for linear algebra, steady-state scheduling, dynamic scheduling for heterogeneous systems, the use of genetic algorithms, fault-tolerance, packet reordering, task or replica migration, and replica placement.

Finally, we would like to thank all the reviewers, for their time and effort, who helped us in the selection process.

* Topic Chairs.