Population-based Metaheuristics for Tasks Scheduling in Heterogeneous Distributed Systems

F. Zamfirache, M. Frîncu, D. Zaharie

Tasks scheduling in heterogeneous distributed systems (e.g. grids) is known as a hard optimization problem because of the lack of reliable information about tasks and resources. In real distributed systems the tasks arrive continuosly and have to be assigned to resources either as they arrive or when a scheduling event is triggered. In this work we analyzed both the case when the scheduling event is triggered when a given number of tasks arrived and the case when the scheduling is activated at pre-specified moments of time.
Besides the traditional heuristics, a lot of population-based metaheuristics (evolutionary algorithms, memetic algorithms, ant systems) have been proposed lately. Most metaheuristics use local search operators involving either the move of a task from one machine to another one or the swap of two tasks belonging to different machines. Depending on the selection of the source and destination machines and of the task(s) to be relocated the local search can be more explorative or more exploitative. In order to make a compromise between greediness and randomness we propose a hybrid mutation operator involving both task relocation ensuring a decrease of the makespan and random relocation. Based on this mutation operator we designed an evolutionary algorithm which was tested both on a traditional benchmark based on the "expected time to compute" model and by using a simulator based on specific probability distributions. The evolutionary algorithm was compared with non-population heuristics (eg. MinMin, MaxMin, MCT, OLB, Sufferage etc.) and with other evolutionary algorithms and the results illustrate its competitiveness especially in the case highly heterogeneous and inconsistent distributed systems. Moreover population-based extensions of some heuristics used for online scheduling were designed and tested leading to improved schedules.