Gridification and Workflow Scheduling for the German D-Grid

Dipl.-Inf. Dietmar Sommerfeld

Abstract

In recent years, scientists from many domains have discovered Grid technology which offers new possibilities for solving problems that are difficult to handle with traditional distributed computing. In the first part of this thesis we examine the process to enable life science applications for execution on the Grid. Such applications often require the analysis of very large data sets and consist of several successive program runs. As a case study we describe the adaptation of the gene-finding tool AUGUSTUS to Grid computing in the context of the MediGRID virtual organization which is part of the German D-Grid. At first we show how the application can be run manually on Grid resources by means of the Grid middleware. Afterwards we describe how the application is represented as a Petri net workflow comprising several tasks which are automatically distributed to appropriate Grid resources by an advanced workflow engine. Finally, we show how a convenient graphical user interface for end users is created within a portal framework.

In the second part of the thesis we concentrate on general Grid workflow scheduling in MediGRID which is concerned with the mapping of workflow tasks to Grid resources. We start our analysis with measurements of the availability of resources in D-Grid and identify four problems inherent to meta-scheduling that make Grid scheduling nearly always a delicate task. To address these issues we present three methods to estimate the queue waiting times of the prevalent cluster resources. Additionally, we propose three selection algorithms to automatically select the best current estimation method. Evaluation shows that prediction works very well on most resources and recognizes the peaks in waiting times that can last up to hours. For the scheduling of workflow tasks we propose and evaluate a two-tier approach that combines contemporary scheduling strategies for workflows and for Grids, and that is suitable for production environments. The first tier uses a list scheduling heuristic to create a full-ahead static schedule of tasks based on predictions of task execution times. The second tier performs a just-in-time mapping of the prioritized tasks to resources according to performance predictions that incorporate the estimated queue waiting times. In the end, the two-tier approach is integrated into the MediGRID workflow engine. We assess the improved scheduling by measurements and demonstrate a significant acceleration of up to 28% in workflow processing compared to the existing strategies.