



UvA-DARE (Digital Academic Repository)

Hierarchical resource management in grid computing

Korkhov, V.V.

Publication date

2009

Document Version

Final published version

[Link to publication](#)

Citation for published version (APA):

Korkhov, V. V. (2009). *Hierarchical resource management in grid computing*.

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

Chapter 7. Summary and conclusions

The work presented in this thesis addresses the problems of resource management in distributed computing from the perspective of multi-layer complex distributed applications. The challenge is to study all the different layers of application parallelism bound into a single hierarchy and propose methods of workload balancing and resource management that are generic enough not to depend on a particular technology of distributed computing and can be reflected and used on the different layers of the application hierarchy.

While talking about technology invariant approaches to efficient resource management for complex parallel applications in distributed heterogeneous environment, we propose and study the adaptive workload balancing (AWLB) approach that can be used in different environments and with different distributed computing technologies. This method of adaptive workload balancing depends on the dynamic characteristics of both application and resources. Traditionally there are two approaches to workload balancing of parallel applications: (1) carefully calculate the distribution of the workload, taking into account all the properties of the environment and application – a time and resource consuming task requiring expert knowledge of the application structure and algorithms used; and (2) distribute the workload in a straightforward way, at best considering only the processing power of the worker nodes – a fast but not very efficient way in terms of parallel performance. We propose an intermediate approach that combines the fast calculation of initial approximation of the workload distribution and its further iterative refinement converging close to the optimum.

Discussing the complex application hierarchy and the ways to decompose complex applications to different layers and manage these layers, we consider a three-layer structure of distribution and parallelism: task (e.g. single parallel application), job (e.g. parameter sweep and task farming), and workflow (functional decomposition of the application). Each layer of this structure has its own requirements to maximize the performance and enable good scalability. Thus appropriate resource management and workload balancing has to be addressed while building the distributed multi-layer and multi-component application as a whole.

Chapter 3 addresses the lowest layer in the complex application hierarchy - a single parallel application running on heterogeneous resources. One of the most challenging problems in porting parallel distributed applications from homogeneous cluster environments to heterogeneous resources is to keep up a high level of parallel efficiency of the computational components. To tackle this problem, we developed a theoretical approach and a generic workload balancing technique that takes into account specific parameters of the resources dynamically assigned to a parallel job, as well as the

application requirements. Here we introduce the AWLB methodology and apply it to the parallel solvers of the Virtual Reactor - one of the driving application of this research.

The case of a single parallel application spanning over a set of parallel computers is described in Chapter 4. The possibility to use a large amount of resources raises the question if and when several parallel computers will indeed bring any performance benefit when used together for a single parallel application. The theoretical approach to estimate the possible speedup of a parallel application in a homogeneous computational Grid is introduced by Hoekstra and Sloot in [49]; here we examine the proposed approach using a real application as a test-case: a Lattice Boltzmann Equation (LBM) solver; A simple model of this application is used for prediction of possible performance gain from the multi-cluster distribution. Compared to the solution for parallel applications spanning over a set of heterogeneous resources proposed in Chapter 3, the approach presented in this chapter is valid for homogeneous Grids, but it gives rather high accuracy in predicting the application speedup using a simple application model and a set of basic environment characteristics.

Another application hierarchy layer that is called a multi-job application is presented in Chapter 5. In this type of the distributed application data processing is organized in a set of separate components that perform information exchange only at the start and finish of the execution. In this chapter the AWLB method developed in Chapter 3 is applied to multi-job applications with divisible workload corroborating the technology independence of this method. We present a hybrid resource management environment, operating on both application and system levels, developed for minimizing the execution time of parallel applications with divisible workload on heterogeneous Grid resources. This integrated environment consists of application-level AWLB applied to User-Level Scheduling (ULS) environment. The latter fills the gap between the application and Grid resource managers: this user-level middleware is a customizable, application-centric scheduler and application hosting environment. Dynamic benchmarking of resources and estimation of the application characteristics is used to optimize the usage of a dynamic user-level pool of Grid resources maintained by the ULS. To prove the concept we perform the experiments with a synthetic application with configurable requirements using EGEE Grid resources and the AWLB algorithm incorporated into the DIANE user-level scheduler. We present the experimental results and discuss the ways to manage the workload of divisible load parallel applications on the Grid, compare different workload distribution methods, illustrate the usage of dynamic resource pool and application performance dependencies with adaptive resource selection.

Chapter 6 discusses the highest layer of the complex distributed application hierarchy: the Grid workflow. The workflows controlled by a dataflow are in the focus of this study: a model of a data-driven workflow is discussed and different strategies of resource management for this type of workflows are evaluated. To be executed on the Grid a distributed workflow needs a framework that enables the workflow enactment and execution control: a workflow management system (WMS). This chapter presents the VLAM-G and its core component, the Run-Time System (RTS), as an implementation of a data-driven WMS. The RTS is a dataflow driven workflow engine

which utilizes Grid resources, hiding the complexity of the Grid from users. Special attention is paid to the concept of dataflow and direct data streaming between distributed workflow components. The architecture, components of the RTS, and the features of VLAM-G workflow execution are presented. As the top layer of the complex application hierarchy, a workflow can perform the management of components representing different lower layers, and the execution of each component is coordinated with the others. This capability of workflows and in particular the VLAM-G WMS is illustrated with a workflow developed for the Virtual Reactor application, which embraces the components of different layers.

The presented work covers the area of execution of multi-component and multi-layer applications in the distributed environment. From the bottom to the top layers of the application hierarchy we describe the methods of appropriate resource management, present implementation and experimental validation of the proposed approaches.

Now it's time to refer back and reconsider the questions posed in Chapter 1. The adaptive workload balancing algorithm was developed to be technology independent and to contain as little application specific properties as possible. On the other hand, it proved its efficiency being used in different environments and employed by different technologies: Message Passing Interface in Chapter 3 and User-Level Scheduling in Chapter 5. Moreover, we were able to use it on different layers of the application hierarchy. Mapping of the application layers and their management was studied throughout Chapters 3, 5, and 6 starting from parallel applications on heterogeneous resources and finishing with distributed workflows managed by a workflow management system. The question about the benefits of a distribution of a single parallel application across a set of Grid resources was addressed in Chapter 4 where the Grid speedup and efficiency metrics were introduced and used to evaluate possible benefits.

In the future research we plan to address the issues of resource and service management within Grids built with Service Oriented Architecture. This modern paradigm seems to require new approaches at a first glance, and the challenge is to apply the ideas and methods presented in this thesis to this type of computing environment. Efficient management and orchestration of Grid services is an essential part of modern workflow management systems, and the plans are to enhance the newest generation of VLAM-G WMS with the capabilities of advanced resource management for service-oriented workflows.