Report of the doctorate thesis presented by Lucas Mello Schnorr

Rapporteur: Siang Wun Song
Professor in Computer Science
Universidade de São Paulo - São Paulo, Brazil

Title of the document evaluated:

Some Visualization Models applied to the Analysis of Parallel Applications

Context of the Study

The document refers to the doctorate thesis of Lucas Mello Schnorr, who is jointly advised by Professor Philippe Olivier Navaux of the Universidade Federal de Rio Grande do Sul (UFRGS), Brazil, and Professor Denis Trystram of the Institut National Polytechnique de Grenoble (INPG), France.

Summary of the document presented

This thesis is composed of a total of 128 pages organized in seven chapters, in addition to two appendices, that contain extended abstracts in Portuguese and in French. The original content and importance of the topic researched, as well as the promising results obtained, make it a document of great impact to grid computing. In the following I give a summary of the document followed by an analysis.

Chapter 1 - Introduction

This chapter presents a succint introduction of the thesis. It starts with an introduction to the main characteristics of distributed computing and, more specifically, of grid computing, namely, dynamism, heterogeneity of hardware and software resources, and presence of multiple administrative domains, in addition to a complex network interconnection. A grid system is scalable, in the sense that its size can be increased. These characteristics can influence the performance of parallel applications during development and execution. Visualization techniques can help the designer in the analysis of parallel applications. Traditional visualization techniques, or their adaptation, however, are not adequate for this purpose. The reasons are basically the network topology complexity and the visualization scalability. The main goal of this thesis is to overcome the drawbacks of using traditional visualization methods. It proposed two visualization models: the three-dimensional model and the visual aggregation model. Both models constitute in a novel way to visualize the behavior of parallel applications and they are useful to the analysis and understanding of parallel applications executed in distributed systems such as grids.

Chapter 2 - Parallel Application Visualization

Chapter 2 presents existing related works. It starts with brief descriptions of existing parallel application visualization systems. Some of them are of historical importance, while others are still in use. Among the examples of performance visualization tools, the following are presented: ParaGraph, TraceView, Pablo, Paradyn, Vampir, Virtue, Jumpshot ParaProf, and Pajé. The chapter ends with a summary of the main visualization techniques, as well in which tools the techniques are used. The presented techniques are of three types, namely, Behavioral (Gannt-charts, variables in two and three dimensions, time-tunnel, phase portraits), Structural (matrix, graph with communications) and Statistical (bar and pie charts, Kiviat diagrams, statistical 3D representations). The main purpose of this chapter is to present related works and serves as a preliminary preparation for the following two chapters, where the author presents the proposed models.

Chapter 3 - The Three-Dimensional Model

This chapter contains the first of the two proposed visualization models: the Three-Dimensional Model. The visual conception of this model is to combine the visualization techniques that show the behavior of the application with techniques that show the structure or statistical data. The chapter starts by presenting the model overview and explains how the trace data are collected and processed. It explains the several components such as the trace reader, the extractor, and the entity matcher. The author then gives three cases that cover the structural representation and statistical representation. The three cases are: parallel application communication pattern, network topology with communication pattern, and logical organization and the communication pattern. The main section of this chapter is section 3.6, where one can see the benefits of the proposed model. The main objective is to create a 3D visual representation, based on the flow of time-ordered visual objects and the base configuration chosen by the user. Figures 3.11, 3.12 and 3.13 illustrate the usefulness of the graphical representation to aid the parallel application designer to analyze the performance of the parallel application.

Chapter 4 - Visual Aggregation Model

This chapter presents the second of the two proposed visualization models: the Visual Aggregation Model. While the first model deals with the network topology, this second model concerns mainly with the issue that most grid applications can be composed of a large number of processes. These processes usually follow a hierarchical organization, the processes belong to machines, which belong to cluster and these to grids. The proposed model takes into account the hierarchical organization of the monitored data. The time-slice algorithm summarizes the behavior of a parallel application in a time interval and serves as input to the aggregation model. The basic concepts include the treemap representation, which is applied to visualize the hierarchies represented by the time-slice technique. Several examples are given to illustrate the proposed model.

Chapter 5 - Triva Prototype Implementation

This chapter presents the prototype implementation, called TRIVA, of the two proposed models. In addition to code developed especially for this implementation, it is based on some existing tools and libraries. These include the generic visualization tool called Pajé and external libraries Ogre3D and GraphViz. Pajé is an existing visualization tool with such characteristics as extensibility, interactivity and scalability. The chapter starts with a discussion of the pros and contras of adopting an existing tool. Then Pajé is described with some detail. The libraries Ogre3D and GraphViz are very useful to implement the proposed models. This chapter discussed the implemented prototype in detail. It is quite technical and contains a description of the modules, sometimes with pseudo-code included.

Chapter 6 - Results and Evaluation

While the previous chapter presents the technical details of the implemented prototype, this chapter concerns the experiments carried out and the evaluation of the proposed tools. To this end, two types of trace data are used: those collected from real experiments run on grids and the synthetic data generated by programs. The reason to use synthetic data is explained, mainly because these data are not easy to obtain due to the large amount of processors involved, or to experiment with complex network topologies. This chapter is very important in the sense that it attempts to validate the proposal and the author presents several cases to show the usefulness of the tools.

Chapter 7 - Conclusion and Future Work

This chapter concludes the document with several final considerations, main contribution and future work. It also contains a list of conference papers produced from this thesis, including, among others, SBAC-PAD 2009, CCGrid 2009, and GRID 2009. The chapter mentions one submitted journal paper and several planned journal paper submissions.

Analysis

The document is well written and easy to read. The author makes a special attempt to summarize the content, whenever appropriate, and this helps the reader to grasp the main objectives of each chapter. The thesis work deals with a very important crucial problem in the development of techniques and tools to aid the parallel application designer to analyze the performance issues of the designed applications. This thesis motivates the proposed work by pointing out the deficiencies of existing visualization tools and technique for this purpose, due to two main factors, namely the network complexity and the scalability of the existing techniques in dealing with a large amount of resources, which is common in actual grid systems.

The thesis contains original work, in the form of a proposal of novel visualization tools for parallel applications. The prototype is implemented on top of some existing tools, plus some considerable amount of new code. The experimental results were carried out both with generated and synthetic trace data. The experiments were run on a large grid consisting of many machines located in several laboratories, in France and in Brazil and required considerable effort and work.

There are some conference publications produced from this thesis. Journal submissions are planned and thus more publications are expected.

Given these considerations, I consider this thesis an excellent piece of work and I recommend approval of this thesis for the obtention of the doctorate degree.

São Paulo, October 3, 2009

Siang Wun Song