REPORT ON THE THESIS
ECONOMIC-BASED DISTRIBUTED RESOURCE MANAGEMENT AND SCHEDULING FOR GRID COMPUTING
by
RAJKUMAR BUYYA

Rajkumar Buyya has written a very good thesis indeed. It introduces new and interesting ideas, and makes a significant contribution to the dynamic field of grid computing.

It is introduced with a concise formulation of the problem addressed by the research, and continues with a comprehensive and valuable literature survey. Chapter 2 is, in fact, a thorough survey of existing grid systems, and thus valuable in its own right.

Chapter 3 makes a significant contribution. The author identifies the context thoroughly, and presents an interesting and useful model based on the computational economy metaphor for resource management and application scheduling.

Chapter 4 describes contributions to the design and development of the Nimrod-G system, both as a tool for setting up parametric experiments on computational grids, and for evaluation of the author's GRACE framework. This chapter also details the establishment of the WWG (World Wide Grid), a pioneering experimental grid comprising resources at various locations around the world; this testbed is in itself a very nice piece of work! The WWG is used to conduct a thorough experimental evaluation.

In order to overcome the inherent limitation of experiments involving disparate facilities around the world, and to more precisely answer detailed questions about allocation and scheduling behaviours, the author describes, in Chapters 5 and 6: a grid modelling and simulation toolkit, and its use to simulate a grid environment and economy-based scheduling within it. Again, we see a comprehensive experimental evaluation based on the simulation tools. Both toolkit and results are useful contributions to the grid community, and advances to the state of the art.

Chapter 7 explores further various challenges in making effective use of grid-based resources. It presents a case study in the application area of molecular docking, the formulation of problems in the area as a parametric case study, and the use of Nimrod-G in processing this study on grid-based resources. The final chapter identifies various directions in which this work can proceed.

Chapter 8 summarizes the thesis and suggests some future directions. One direction that might be interesting is an in-depth comparison with other non-grid methods of solving certain problems, to validate the argument that the grid concept itself is the best way to tackle the problem concerned (see comments below).

In summary, the thesis sets out to contribute in the following areas: the computational economy metaphor; means of expression of requirements, scheduling policies, and adaptive scheduling mechanisms; a prototype system for implementation and validation; real-world experiments with computational economy models; a simulation toolkit and testbed for more detailed, repeatable performance analysis; and to demonstrate the effectiveness of grid technologies in application. The author makes useful contributions in all these areas, and makes a strong case for the computational economy metaphor. In so doing, he also contributes useful tools to the grid community, in Nimrod-G and the grid simulation toolkit. He also takes some steps in what is perhaps the most important challenge currently in front of grid researchers and practitioners: to convince application developers that grid technolgies are beneficial in resource-intensive problems.

I recommend that the thesis be awarded, with some corrections as indicated below.