Chapter 8

# Conclusions and Future Directions

## 8.1 Summary

Grids are emerging as the infrastructure for next generation computing. In Grid environments, the resources are heterogeneous and geographically distributed with varying availability and a variety of usage and cost policies for diverse users at different times and, priorities as well as goals that vary with time. The management of resources and application scheduling in such a large and distributed environment is a complex task. It is envisioned that the use of a distributed computational economy is an effective metaphor for the management of resources and we have developed an architectural framework called the Grid Architecture for Computational Economy (GRACE).

To support the thesis that an economic-based Grid resource management and scheduling system can deliver significant value to users, resource providers and consumers, compared to traditional approaches, we have:

- identified the key requirements that an economic-based Grid system needs to support,
- developed a distributed computational economy framework called the GRACE, which is generic enough to accommodate different economic models and maps well onto the architecture of wide-area distributed systems,
- designed deadline and budget constrained scheduling algorithms with four different strategies: cost, time, conservative-time, and cost-time optimisations,
- developed a Grid resource broker called Nimrod-G that supports deadline and budget constrained algorithms for scheduling parameter sweep applications on the Grid,
- developed a Grid simulation toolkit, called GridSim, that supports discrete-event based simulation of Grid environments to allow repeatable performance evaluation under different scenarios,
- evaluated the performance of deadline and scheduling algorithms through a series of simulations by varying the number of users, deadlines, budgets, and optimisation strategies and simulating geographically distributed Grid resources, and
- demonstrated the effectiveness and application of Grid technologies for solving real-world problems such as molecular modelling for drug design on the WWG (World-Wide Grid) testbed.

## 8.2 Conclusions

The resource management and scheduling systems used in Grid environments need to be adaptive so that they can handle dynamic changes in availability of resources and user requirements. At the same time, they need to provide scalable, controllable, measurable, and easily enforceable policies for management of the aforementioned resources. To address these requirements, a resource broker called Nimrod-G has been developed. The Nimrod-G resource broker supports the deadline and budget constrained algorithms for scheduling task-farming applications on large-scale distributed systems. The use of a component-based architecture has enabled the Nimrod-G broker implementations for different middleware technologies such as Globus, Legion, and Condor with a minimal development effort.

The Nimrod tools for modeling parametric experiments are maturing and in production use for cluster computing. A prototype version Nimrod-G resource broker is publicly available [101] and efforts are currently underway to deploy it on production Grids. The Nimrod-G task farming engine (TFE) services have been used for developing customized clients and applications. In addition, the TFE job management protocols and services can be used for developing new scheduling policies. We have developed a number of market-driven deadline and budget constrained scheduling algorithms, namely, cost, time, conservative-time, and cost-time optimizations. The results of scheduling experiments with different QoS requirements show promising insights into the effectiveness of distributed computational economy for management of resources and its usefulness in application scheduling with optimizations.

A number of deadline and budget constrained scheduling experiments have been performed with a variety of requirements at different times by selecting different sets of resources available in the WWG testbed during each experiment. They can be categorised into the following areas:

1. Cost optimisation scheduling during Australian peak and off-peak times,
2. Cost and time optimisation scheduling using cheap local and expensive remote resources,
3. Large scale scheduling using cost and time optimisation algorithms, and
4. Molecular docking application scheduling using cost and time optimisation algorithms.

In a competitive commodity-market economy, the resources are priced differently at different times based on the supply and demand. For example, they are priced higher during peak hours and lower during off-peak hours. In the first experiment, for the same deadline, the results of application scheduling show that the broker is able to process jobs with low cost during Australian off-peak hours and more during peak hours. This means that pricing resources higher during peak hours and lower during off-peak hours motivates users to process their low priority jobs during off-peak. In the second experiment, for a given deadline and budget, we explored application scheduling with the cost and time optimisation strategies. The results show that the broker is able to process earlier with a time-optimisation strategy than the cost-optimisation, but at the expense of a higher financial cost. Similar results have been observed in the third experiment with a large number of geographically distributed heterogenous resources.

The fourth scheduling experiment demonstrates the effectiveness and application of Grid technologies for solving real-world problems by creating a Virtual Laboratory environment. The molecular modelling for drug design application has been formulated as a parameter-sweep application using the Nimrod-G parameter specification language. The experiment then used the Nimrod-G broker to process molecular docking jobs on the Grid. The scheduling experiments with cost and time optimisations using Nimrod-G, demonstrates that the users can indeed express their valuations naturally by defining deadline, budget limits, and optimisation preference. It also demonstrates that Grids indeed enable the sharing and aggregation of geographically distributed resources for solving real-world data-intensive computing problems.

The GridSim toolkit has been used to evaluate the performance of deadline and scheduling algorithms through a series of simulations by varying the number of users, deadlines, budgets, and optimisation strategies and then simulating geographically distributed Grid resources. The scheduling simulations with varying deadlines and budgets for cost-optimisation strategy show that as the deadline is increased the cost of computation decreases until it reaches the optimal level—i.e., processing all jobs on the cheapest resources. The time-optimisation scheduling showed that as the budget is increased, the completion time decreases and the cost increases. Also, when the number of users competing for the same set of resources increases, there will be proportional impact on others depending on each user's strategies and constraints. Apart from complementing and strengthening the results of Nimrod-G scheduling studies, the simulations demonstrate the capability of GridSim and the ease with which it can be used to develop and evaluate the performance of new scheduling algorithms.

The results of scheduling applications with different QoS requirements demonstrates that market-based systems, such as the Nimrod-G broker, allow users to trade-off QoS parameters, deadlines and computational costs, and offer an incentive for relaxing their requirements—reduced computational cost for relaxed deadline when timeframe for earliest results delivery is not too critical. This approach of offering an economic incentive for resource owners to share their resources and resource users to trade-off between the deadlines and budgets, promotes the Grid as a platform for mainstream computing. This could in turn help lead to the emergence of a new service oriented computing industry.

The results of these experiments also demonstrate that the computational economy framework in the Grid environment helps in regulating the supply-and-demand for resources and offers an incentive to resource owners to share their resources and resource users to think about trade-off between the deadline and budget. It provides a decentralized resource management capability and is adaptable to changes in the Grid environment and user requirements. The economic-based Grid system is scalable, controllable, measurable, and uses easily understandable policy for management of resources.

The realization of the GRACE framework by utilising existing technologies such as Globus and providing new services that are essential for resource trading and the aggregation of resources, demonstrates that an economic-based Grid resource management systems can be developed and deployed. It also demonstrates that future network computing applications will have the capability to select resources/services dynamically at runtime depending on their availability, capability, cost, and the users QoS requirements.

## 8.3 Future Directions

This thesis formulated a comprehensive distributed computational economy architectural framework and strategies for service-oriented Grid computing. It demonstrated the benefits of developing economic-based Grid systems for distributed resource management and scheduling. A number of deadline and budget constrained scheduling algorithms for different optimisation strategies have been developed to meet user's QoS requirements.

This work has laid a foundation for the Grid economy and it opens up several avenues for future work in economic-based Grid resource management and scheduling.

### 8.3.1 Supporting Different Application Models

While a parameter-based parallel application model is a dominant model for many applications (e.g., molecular modelling, protein folding, high-energy physics, data mining, design explorations, and structural engineering) that are being used to explore the use of the Grid, there are applications (e.g., computational fluid dynamics) that need a different application model. These applications have tasks that need to communicate frequently with each other and may have interdependencies. To execute such tasks, resources need to be co-allocated to enable communications between tasks at runtime. This introduces various complexities into resource management and scheduling. To overcome these complexities an advance reservation capability is needed, but it is hard to get resources at multiple sites for co-allocation since each resource has a different allocation policy. In this case, computational economy approach could help in prioritizing allocations by supporting the cancellation of existing allocations, if the penalties are less than the benefits, and encourage resource providers to create alliances to support co-reservation and allocation.

### 8.3.2 Supporting Different Economic Models

While a commodity economic model is expected to be the dominant model for pricing resources, similar to that of the Internet where access pricing is based on a flat-price model with some variance and large-slot size, a computational economy approach for Grid resource management requires extensive exploration. For example, currently the Nimrod-G scheduler does not support changes in the price of resources dynamically within a small period, less than the job execution time, once initial scheduling decisions are made. This is because, in scheduling the remaining jobs over the resources within the remaining budget, the scheduler assumes that the price of resources does not change. In addition, the scheduler uses the current price to calculate the cost of jobs that have completed in the past. Hence, using the current scheduler in a system where prices vary and cannot be guaranteed, cost estimations become meaningless and the budget cannot be guaranteed. In order to overcome this limitation, new scheduling strategies and algorithms that learn from historical and market dynamics are needed. These strategies and algorithms should not only be able to adapt dynamically to the changes in resource conditions at runtime, but also to changes in access prices, even during the execution of jobs.

There exist a few other economic models such as auctions, contract net, and bid-based proportional resource allocation for resource trading. Previous research in market-based systems have explored such frameworks, but they expect the user to build applications explicitly using a market-oriented programming framework, where the programmer has to develop a budget allocation strategy for each task and create bids. This makes application development harder and time consuming. Tools similar to Nimrod-G need to be developed that automatically take budget allocations into account and work with different economic models.

It is expected that Nimrod-G will be enhanced to support scheduling with advance resource reservation. Plans are also underway to explore other economic models such as tenders/contract-net and auctions for resource brokering. These new models require new scheduling algorithms.

### 8.3.3 Accounting

Within the Grid community, there is a great interest in building an accounting model and infrastructure. In the GRACE framework, we have proposed the concept of Grid Market Directory, Grid Bank (GB), and automated payment mechanisms. This is similar to a debit/credit card company mediating payments, i.e., the buyer presents the card to purchase items, the seller requests the credit/debit card company to make payment, and the credit/debit card company then claims aggregated amount periodically from the customer. We propose a similar model for implementing the Grid Bank and automating payments.

In a Grid environment, both resource owners and consumers need to have an account in the GB, which can record activities. Under such a framework, the Nimrod-G user, having a unique identity in the Grid, submits their application to the broker along with the deadline and budget. When the broker schedules a job on the resource, it can inform the resource owner about its GB account details to which expenses should be charged. The GB admission control model having the record of resource usage details can then charge the user account. The Nimrod-G broker agent also maintains the record of resource usage details for each job, to which it can refer for verifying charges if there is a fraud.

To support anonymous online payments, we need digital currency. Although we have not developed digital currency, we note that electronic currency technology is rapidly progressing with emerging e-commerce infrastructures [98] such as NetCash, NetCheque, and Paypal. When electronic currency is available, Nimrod-G can be enhanced to automate the payments with minimal effort.

### 8.3.4 Enhancing GridSim to Support QoS based Resource Entities

The GridSim toolkit is rapidly evolving. The network model needs to be enhanced by supporting various types of networks with different static and dynamic configurations and cost-based QoS services. To enhance resource model entity with file I/O operation, off-the-shelf storage I/O simulators need to be incorporated. GridSim currently supports a framework for a resource model with advance reservation. To enable the simulation of Grid resource management with economic models such as tenders and auctions, the FIPA (Foundation for Intelligent Physical Agents) standards [30] based interaction protocol infrastructure can be integrated, along with necessary enhancements to the resource model to support admission control.

### 8.3.5 Wide-Area Data-Intensive Programming and Scheduling Framework

An experience in developing a prototype Virtual Laboratory environment for distributed drug design demonstrates the ease of use and applicability of the Nimrod-G tools for data intensive computing on the Grid. The current system can be extended to support adaptive mechanisms for the selection of the best CDB service depending on the access speed and cost. A new project, called the HEPGrid (High Energy Physics and the Grid Network), has been initiated to develop a virtual laboratory environment for enabling high-energy physics events processing on distributed resources on a larger scale [96].

We expect that the economy driven approach to resource management and scheduling will make a great impact on the eventual success and widespread adoption of the Grid in day-to-day computational activities.