IEEE distributed systems ONLINE
Expert-authored articles and resources

Nimrod-G Resource Broker for Service-Oriented Grid Computing

By Rajkumar Buyya, David Abramson, and Jonathan Giddy

Grid and Peer-to-Peer (P2P) computing platforms enable the creation of *virtual enterprises (VEs)*. These VEs facilitate sharing geographically distributed resources—such as computers and data sources—for solving large-scale problems in science, engineering, and commerce [1]. However, application composition, resource management and scheduling in these environments is a complex undertaking. The geographic distribution of resources owned by different organizations with different usage policies, cost models and varying loads and availability patterns is problematic. To address these resource management challenges, we have proposed and developed a computational economy framework for resource allocation and supply and demand regulation for resources. The new framework offers incentive to resource owners for being part of the Grid and motivates resource users to trade-off between time of results delivery and computational cost [2].

Assessing wants and needs

In P2P and Grid computing environments, resource management systems need to provide mechanisms and tools that facilitate the realization of goals for both service providers (resource owners) and consumers. Resource consumers need

- a *utility model*—how consumers demand resources and their preference parameters, and
- *brokers* that automatically generate strategies for choosing providers per user requirements and manage all issues associated with application execution.

The service providers need tools and mechanisms for *price generation schemes* to increase system utilization and *protocols* that help them offer competitive services. For the market to be competitive and healthy, coordination mechanisms are required to help reach equilibrium price—the market price at which the supply of a service equals the quantity demanded [3].

Numerous economic theories including micro and macroeconomic principles have been proposed. We can use some common economic models for trading goods and services and employ service price negotiation protocols in Grid computing [1]:

- commodity market models,
- posted price models,
- bargaining models,
- tendering or contract-net models,
- auction models,
- bid-based proportional resource sharing models,
- community—coalition—bartering models,
- monopoly and oligopoly models.

These economic models help in regulating the supply and demand for resources in Grid-based virtual enterprises.

Introducing the Nimrod-G resource broker

Our Grid resource broker— called Nimrod/G— supports soft deadline- and budget-constraining scheduling of task-framing applications on a P2P Grid distributed across the globe [1, 3, 4]. This approach provides economic incentive for resource owners to share their resources on the Grid and encourages the emergence of a new service-oriented computing industry. More importantly, it provides mechanisms to trade-off quality of service (QoS) parameters, deadlines, and computational costs and offers incentive for relaxing requirements. Depending on users' QoS requirements, our resource broker dynamically leases Grid services at runtime depending on their cost, quality, and availability. In addition to deadline and budget constraints in scheduling, our broker supports optimization of both or either.

Nimrod-G is a tool for automated modeling and execution of parameter sweep applications (parameter studies) over global computational Grids. It provides a simple, declarative parametric modeling language for expressing parametric experiments. Domain experts can easily create a plan for parametric computing and use the Nimrod runtime system to submit legacy jobs for execution. It uses an economics paradigm for resource management and scheduling on the Grid. It supports user-defined deadline and budget constraints for schedule optimizations and regulates supply and demand of resources in the Grid by leveraging the services of Grace (Grid Architecture for Computational Economy) resource trading [2, 5].

Nimrod-G provides a programmable and persistent *task farming engine* that can create and plug in user-defined scheduling policies and customized task farming or parameter sweep applications (for example, ActiveSheets [7], that execute Microsoft Excel computations and cells on the Grid). The TFE coordinates resource trading, scheduling, data staging, execution, and gathering results from remote Grid nodes to the user's home transparently.

Nimrod tools for modeling parametric experiments are quite mature and in production use for cluster computing. A prototype version of Nimrod-G has been developed and is undergoing continuous refinements. Developers have used Nimrod-G broker's persistent and programmable TFE services in developing customized clients and applications. An associated dispatcher is capable of deploying computations (jobs) on Grid resources enabled by Globus [8], Legion [9], and Condor [10]. TFE Jobs Management protocols and APIs are used for plugging in new scheduling policies. We have developed a number of deadline-based, market-driven scheduling algorithms: time and cost optimizations with deadline and budget constraints [11]. Once price is established using suitable economic models, depending on users requirements, our scheduling algorithms can select or lease resource services depending on their cost, quality, and availability.

Current experiments

We have conducted scheduling experiments involving over 200 job executions on our World Wide Grid (WWG) testbed with resources located on five continents: Australia, Asia, Europe, North America, and South America. Each job had a specified budget and timeframe. The system then prioritized and distributed the jobs according to their deadlines to the cheapest available machines depending on the time of day in Australia and other continents. If the cheapest machine was unavailable at a certain time, the scheduler selected another more expensive resource as long as it was within budget, to meet the deadline. The results of our scheduling experiments can be found in [1, 3, and 5].

Our Nimrod-G resource broker has applications for drug design, human genome database analysis, aerospace modeling, automobile crash simulation, structural engineering, circuit simulation, network simulation, astrophysics, high-energy physics, oil and mining exploration, ray tracing, data mining, financial modeling, and so on.

We have already explored the capabilities of our Nimrod-G resource broker in a virtual drug design application for processing molecules in a protein databank on WWG testbed resources. A prototype version of Nimrod-G resource broker is available for download from our Web site.

Web Resources

Nimrod-G Website: http://www.csse.monash.edu.au/~davida/nimrod/

Economy Grid Website: http://www.buyya.com/ecogrid/

GridSim Website: http://www.buyya.com/gridsim/

Virtual Laboratory http://www.buyya.com/vlab/

World Wide Grid testbed Website: http://www.buyya.com/ecogrid/wwg/

References

1. R. Buyya, et al., "Economic Models for Management of Resources in Peer-to-Peer and Grid Computing", *Proc. Int'l Conf. Commercial Applications for High-Performance Computing.* (ITCom 2001), SPIE Press, Bellingham, WA, USA, 2001.
2. R. Buyya, et al, "An Economy Driven Resource Management Architecture for Global Computational Power Grids", *Proc. Parallel and Distributed Processing Techniques and Applications* (PDPTA 2000), CSREA Press, USA, 2000.
3. D. Abramson, et al., "High Performance Parametric Modeling with Nimrod/G: Killer Application for the Global Grid?", *14th Int'l Parallel and Distributed Processing Symposium* (IPDPS 2000), IEEE CS Press, Los Alamitos, Calif., USA, 2000.
4. R. Buyya, et al., "Nimrod/G: An Architecture for a Resource Management and Scheduling System in a Global Computational Grid", *Proc. 4th Int'l Conf. on High Performance Computing in Asia-Pacific Region* (HPC Asia 2000), IEEE CS Press, Los Alamitos, Calif., USA, 2000.
5. R. Buyya, et al., "A Case for Economy Grid Architecture for Service-Oriented Grid Computing", *Proc.10th IEEE Int'l Heterogeneous Computing Workshop* (HCW 2001), IEEE CS Press, Los Alamitos, Calif., USA, 2001.
6. D. Ferguson, et al., "Economic Models for Allocating Resources in Computer Systems", *Market-based Control: A Paradigm for Distributed Resource Allocation*, World Scientific Press, Singapore, 1996.
7. D. Abramson, et al., "ActiveSheets: Super-Computing with Spreadsheets", *Proc. High Performance Computing Symp.* (HPC'01), *Advanced Simulation Technologies Conf.*, Society for Modeling and Simulation (SCS) Press, San Diego, Calf., USA, 2001.
8. K. Czajkowski, et al., "A Resource Management Architecture for Metacomputing Systems", *Proc. 4th Workshop on Job Scheduling Strategies for Parallel Processing* (IPPS/SPDP '98), Springer Verlag, Germany, 1998.
9. S. Chapin, et al., "The Legion Resource Management System", *Proc. 5th Workshop on Job Scheduling Strategies for Parallel Processing*, Springer Verlag, Germany, 1999.
10. M. Litzkow, et al., "Condor-A Hunter of Idle Workstations", *Proc. of the 8th Int'l Conf. of Distributed Computing Systems*, IEEE Press, ***USA, 1988.***
11. R. Buyya, et al., "An Evaluation of Economy-based Resource Trading and Scheduling on Computational Power Grids for Parameter Sweep Applications", *2nd Int'l Workshop on Active Middleware Services* (AMS 2000), Kluwer Academic Press, Boston, USA, 2000.

Rajkumar Buyya is a PhD candidate in the School of Computer Science and Software Engineering, Monash University, located in Melbourne, Australia. Contact him at rajkumar@csse.monash.edu.au.

David Abramson is a Professor in the School of Computer Science and Software Engineering, Monash University, located in Melbourne, Australia. Contact him at davida@csse.monash.edu.au.

Jonathan Giddy is a Research Scientist in the Co-operative Research Centre for Enterprise Distributed Systems Technology, Monash University, located in Melbourne, Australia. Contact him at jon@csse.monash.edu.au.

Cite this article as:

Rajkumar Buyya, David Abramson, and Jonathan Giddy, "Nimrod-G Resource Broker for Service-Oriented Grid Computing", IEEE Distributed Systems Online, Volume 2, Number 7, November 2001.