A dynamic model of a GRID market

Uli Harder, Peter Harrison, Maya Paczuski & Tejas Shah
Department of Computing & Department of Mathematics
Imperial College London
Huxley Building, 180 Queen's Gate, London SW7 2AZ
England
{uh, pgh}@doc.ic.ac.uk, maya@imperial.ac.uk, tss101@doc.ic.ac.uk

Abstract

In this paper we discuss possible market mechanisms for the GRID. We propose a complete dynamic model of a GRID market with providers, middlemen and users. Using a multi-agent simulation of the proposed model, we present preliminary results of the price volatility of abstract units in the GRID market.

1. Introduction

From the hardware point of view the GRID is simply a collection of computing, storage and networking devices. At first glance it looks like standard cluster computing. However, the GRID is equipped with middleware that makes it easy for the user to find the right resources for their application. A good analogy is the Internet where search engines are the "information middleware" that bring clients and server together [8].

There are currently two main options for middleware: Globus [1] and Legion [2]. Ideally the GRID should be a heterogeneous collection of computing equipment, in reality it is at present run on LINUX machines.

Currently only a few GRIDs exists, in the UK the main GRIDs are used by computer scientists to develop middleware and architectures. The other main users are particle theorists who already use GridPP for scientific computations. Soon this will be enlarged to cope with the data created by the Large Hadron Collider at CERN. Other scientists are starting to use the technology for their purposes. Another pool of potential users is the banking and insurance market where many computational applications are for instance based on Monte Carlo techniques. Further afield but no less important there are games companies beginning to look at how they can use GRID technology.

2. The future of the GRID

There is consensus in the research community that the GRID, if it is successful will be used in a commercial context. One possible development might be that companies specialise in providing GRID services in a similar way telephone companies operate today. A second scenario is that large organisations will sell spare computing power to others. A third scenario sees private PC owner selling spare CPU cycles of their machines. Since the GRID is in a way meant to do away with the PC this seems a less likely option. In this paper we present a model for the first scenario.

3. Charging for computing time

The name "GRID" was chosen in analogy with the electrical power GRID. Like electrical appliances are simply plugged into a socket, computing applications are plugged into the Computer GRID. So, it is only a matter of time until GRID providers will want to sell computing power to end users. Therefore they will have to find out how to charge for their services.

The vision is that similar to electricity supply there are few negotiations between producers and end-customers. In the privatised electricity market there tend to be middlemen that buy electricity from producers and sell it on to end-customers using a range of service plans and tariffs. There have been successful and detailed multi-agent simulations of the UK and other electricity markets [3].

In the case of the GRID we have an additional problem because the consumption of the user can not be measured easily in kWh, but rather the user needs several components to perform a computation: CPU, disks, memory and networking. If one of the components is missing all others loose their value! In this paper we assume that in future a "GRID computing unit" (GCU) will be established. In fact this may simply be measured in real time access like telephone usage. In this case the middleman would enter contracts which specify the minimum number of CPUs, disk, etc that they have to supply for specific time period. Payments in the GRID-world could be made in a similar way utility/phone bills are settled nowadays.

Charging for computer access, or indeed any commodity, service, etc., can have two reasons:

- maximisation of utilization, fair access to machines that have been purchased collectively
- profit maximisation.

In this paper we shall assume the aim of the providers and middleman is to maximise their profit. The way for producers, middlemen and users to interact are auctions or commodity markets, or a combination of both.

4. Background

The presumably oldest record of a charging system for computer access is from Sutherland in 1968 [15]. There is quite a large body of literature from the 70's where the authors researched the access to mainframe computers [11, 9].

More recently authors have begun to investigate the possible commodity market created by the GRID. There is "G-commerce" [17, 18, 19] which investigates auctions and commodity markets. The "Compute Power market" approach which is more focussed on a completely heterogeneous market with supplier and consumers of all sizes [5]. There is also cpu charging (SPAWN) [16] and Popcorn [12] and a scheme based on CUMULUS [14]. Authors have also explored auctioning systems as replacements for scheduling algorithms [6].

For completeness we need to mention ideas from the mobile agent community where financial ideas are used to manage mobile agents [4].

5. Our model

In our model there are three agents: the GRID, middlemen and users. Though the GRID is likely to be several organisations in our simplified model there is only one. The GRID produces GCUs at a constant rate. The middlemen pass the GCUs to the users. Middlemen can not deal with other middlemen and users can not deal with the GRID directly. This setup is very similar to the Resource Allocation Game [13], an extension of the Minority Game [7]. However, in this case we have the additional structure of a market and a charging scheme. The model has got two timescales a fast one and a slow one.

At each *tick* of the fast timesacle the GRID produces a fixed number of GCUs. These are auctioned off to the middlemen who put in a bid for the number of GCUs they expect to need in the next tick to satisfy the demand of their customers. The GRID allocates the GCUs to the middlemen by trying maximise its profit and serving the highest bids first. If there is a greater demand than supply some middlemen will not get all or any of the GCUs they bid for. If there is oversupply the GRID will destroy surplus GCUs. The middlemen bid all their cash except a small safety margin. The users demand one or zero GCUs from their middleman at each tick randomly with a uniform distribution. They pay their middleman immediately at the agreed price. The middlemen use the money they collect this way for their next bid.

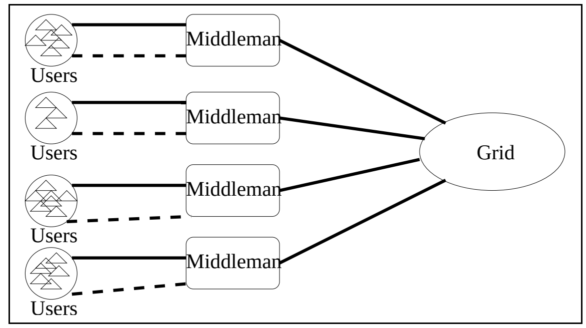

Figure 1. Our model, the broken line corresponds to transactions taking place at the slow timescale

At the slow timescale the middlemen decide whether they should change their current price p_t depending on the cash they currently hold C_t, the amount of cash they started out with initially C_0 and the average price p_a:

$$p_{t+1} = \begin{cases} p_t/2 & \text{if } C_t > 2C_0 \\ p_t & \text{if } C_0 < C_t \leq 2C_0 \\ p_a & \text{if } C_t < C_0. \end{cases}$$

The first rule stops inflation and might be thought of as intervention by a regulator or the Inland Revenue. The second rule preserves the status quo (*don't fix it, if it ain't broken*). The last rule is likely to make the middleman's price more attractive to users. At the same time users decide whether they want to change middleman. They do so by looking at two indicators. They will change middlemen if they are not satisfied with the service level. At each tick they make a note whether their demand was satisfied, middlemen may run out of GCUs if the have not secured enough in the auction. If S_{av} is the average service level achieved by all brokers and the user has received Rec GCUs and requested Req a change of middleman will take place if

$$sTol < \frac{S_{av} - \frac{Rec}{Req}}{S_{av}}$$

where $sTol$ is the tolerance level of a particular user. Similarly the users do not like price hikes. If P^{new} is the new price offered by its middleman, P^{old} the previous price and P^{old}_{av} the previous average price, the user will change middlemen if

$$cTol < \frac{P^{new} - P^{old}}{\sqrt{P^{old}_{av} P^{old}}}$$

where $cTol$ is their price tolerance level. Each user has its own price and service tolerance level which do not change. The new middleman is randomly chosen.

The choice of two timescales was motivated by work of Nagel et al. [10], where the authors have shown that a well-defined market only emerges if price changes happen at a different timescale to the consumption. In our simulation we used 10,000 ticks in the fast timescale for every tick in the slow timescale. This roughly approximates minutes and months.

6. Preliminary Results

We implemented the model described above as a multi-agent simulation. To get some indication how the model behaves we ran a configuration with 10 middlemen and 100 users demanding 50 GCUs on average whilst the GRID produces 40, 45 and 50 GCUs. The tolerance levels for price changes and service are set to 0.5 for all users. When we plot the changes in the average price middlemen demand from the users we see that the distribution is very broad indicating a high volatility of the GCU price, see figure 2. In fact a double-logarithmic plot reveals that the price change follows a power law with a gradient of -1.4, see figure 2.

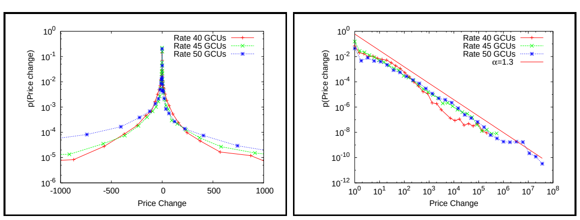

Figure 2. A half-logarithmic plot of the price changes and double logarithmic plot of the positive price changes

When there is an oversupply of GCUs, in this case 55, the model quickly settles down and all users are trapped with one middleman, as its service level is satisfactory and the price will not change at all. Still there is a bit dynamics left in the model when we look at the cash a middleman is left with at the end of tick of the slow timescale, see 3. This is due to the variance in the requests for GCUs by the users.

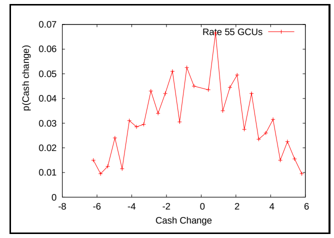

Figure 3. Changes in the cash of the only active middleman

7. Future work

In our future investigation into models for the GRID market we shall need to consider simpler models to see what really drives the model as well as look at different auction schemes and ways the users behave and make decisions. Additionally we need a more satisfactory pricing mechanism, perhaps like the ones used in the electricity markets.

Acknowledgments

The research was partially funded by EPSRC (research grant PASTRAMI, GR/S24961/01).

References

[1] Globus. www.globus.org.

[2] Legion. *www.cs.virginia.edu/ ˜legion/*.

[3] A. Bagnall. *Applications of Learning Classifier Systems, Studies in Fuzziness and Soft Computing vol. 150*, chapter A Multi-Agent Model of the UK Market in Electricity Generation. 2004.

[4] J. Bredin, R. T. Maheswaran, Ç. Imer, T. Başar, D. Kotz, and D. Rus. A game-theoretic formulation of multi-agent resource allocation. In *Proceedings of the Fourth International Conference on Autonomous Agents*, pages 349–356. ACM Press, June 2000.

[5] R. Buyya. Economic-based Distributed Resource Management and Scheduling for Grid Computing. *PhD Thesis, Monash University, Melbourne, Australia, April 2002*. cs.DC/0204048.

[6] R. A. Cagliano, M. D. Fraser, and M. E. Schaefer. Auction Allocation for Comuting Resources. *Communications of the ACM*, 38(6):88–96, 1995.

[7] D. Challet and Y.-C. Zhang. Emergence of Cooperation and Organization in an Evolutionary Game. *Physica A*, 246:407, 1997.

[8] P. Clarke. Introduction. *Getting ready for the Grid 11 February 2004 Institute of Physics HEPP group*, 2004. http://www.gridpp.ac.uk/meetings/iopfeb04/.

[9] I. W. Cotton. Microeconomics and the Market for Computer Services. *Computing Surveys*, 7(2):95–111, 1975.

[10] K. Nagel, M. Shubik, and M. Strauss. The importance of timescales: simple models for economic markets. *Physica A*, 340(4):668–677, 15 September 2004.

[11] N. R. Nielsen. The Allocation of Computing Resources – Is Pricing the Answer? *Communications of the ACM*, 13(8):467–474, 1970.

[12] O. Regev and N. Nisan. The POPCORN Market an Online Market for Computational Resources, 98. *Proceedings of the first international conference on Information and computation economies, pp. 148 - 157 Series-Proceeding-Article Year of Publication: 1998 ISBN:1-58113-076-7*.

[13] R. Savit, S. A. Brueckner, H. V. D. Parunak, and J. Sauter. Phase structure of resource allocation games. *Physics Letters A*, 311(4-5):359–364, 19 May 2003.

[14] B. Stiller, J. Gerke, P. Flury, P. Reichl, and Hasan. Charging Distributed Services of a Computational Grid Architecture. *Computer Engineering and Networks Laboratory TIK, Swiss Federal Institute of Technology, ETH Zurich, Switzerland*.

[15] I. E. Sutherland. A futures market in computer time. *Commun. ACM*, 11(6):449–451, 1968.

[16] C. A. Waldspurger, T. Hogg, B. A. Huberman, J. O. Kephart, and W. S. Stornetta. Spawn: A distributed computational economy. *Software Engineering*, 18(2):103–117, 1992.

[17] R. Wolski, J. S. Plank, and J. Brevik. G-commerce Building Computational Marketplaces for the Computational Grid. *University of Tennessee Technical Report UT-CS-00-439*. http://www.cs.utk.edu/ ˜rich/publications/CS-00-439.ps.gz.

[18] R. Wolski, J. S. Plank, J. Brevik, and T. Bryan. G-commerce: Market Formulations Controlling Resource Allocation on the Computational Grid. *University of Tennessee Technical Report UT-CS-00-450*. http://www.cs.utk.edu/ ˜rich/publications/CS-00-450.ps.gz.

[19] R. Wolski, J. S. Plank, J. Brevik, and T. Bryan. Analyzing Market-based Resource Allocation Strategies for the Computational Grid. *The International Journal of High Performance Computing Applications, Sage Science Press*, 15(3):258–281, 2001.