A peer-to-peer Market for Grid Computing

Uli Harder *

Abstract

In one of the latest developments in Grid Computing Peer-to-Peer (P2P) networks have become fashionable for the Grid system architecture. The current implementations are either Data Grids or CPU sharing Grids. In this paper we begin to investigate how a P2P network can behave as CPU sharing Grid. In our model the nodes of the network are very simple and either require or offer service for the exchange of money. In our we would like to investigate for instance how the nature of P2P overlay network affects this Grid economy.

1 Introduction

In recent years the concept of Grid computing has become fairly well established as a paradigm for distributed computing [1]. One relatively new idea is to employ peer to peer technologies and economic incentives for service discovery and scheduling in Grid systems.

In this paper we present a simplified model of Grid peer to peer system where nodes either sell or buy a single unit of a resource. This is motivated by the idea that users will buy or barter for access to computing power. We do not go into the detail whether this is the CPU or disk space or a combination of both. This is a similar approach taken in [2].

This model was inspired by the work of the Internet Centre on the MaGog Global Open Grid (GOG) software [3].

2 Related work

The idea to use economic ideas to manage access computer systems is fairly old and the first reference is usually credited to Sutherland [4]. He describes an auctioning scheme used to manage access to a PDP-1. Later more involved approaches were developed for mainframes [5, 6]

However, most approaches taken in the mainframe days can not be applied to the Grid as any central service will normally not scale and prevent the system from becoming bigger. There is a one new approach in Computer Science based on Hayek's idea of Catallaxy [7, 8].

In earlier a three-tier model of a future Grid market was explored [2, 9]. We now think that a different approach might be more interesting where ideas from Peer-to-Peer (p2p) networks are used to provide the Grid services. For a review

*Department of Computing, Imperial College London Huxley Building, 180 Queens Gate, London SW7 2AZ, UK

of p2p networks see [10]. A data grid that is openly using p2p technologies is P-grid [11]. However, it has not progressed to a production application. A centralized P2P network that uses a central server to locate resources is Tycoon [12, 13, 14, 15]

In the Physics community relevant related research has for instance been conducted by Rosvall and Sneppen [16]. They investigate the propagation of information on networks with agents that have only local knowledge. Amongst others there is also a recent paper by Goshal and Newman [17] which describes a method to self-organize a network to reduce search time. Whilst some P2P networks show signs of a scale-free small world network [18] this is not necessarily a desirable feature when nodes search for information in the network.

One of the main motivation for this work is the desire to assist to establish a market economy for Grid Computing similar to that of privatized electricity markets for instance [19]. In the electricity markets one of the main difficulties for a regulator is to detect collusion of providers in order to force up prices [20]. With our model it might be possible to determine whether any such behaviour is actually based on agents decisions or in fact determined by the market rules. In the latter case the regulator could try to change rules and therefore make detection of cartels easier.

3 The model

The peer to peer overlay network of our "Grid" is defined by a connected graph. The nodes of the graph are the computing resources connected to the Grid and are either

- Buyer. This nodes tries to acquire a resource and bids a maximum price it is willing to pay.
- Seller. A node that sells exclusive access to its resource for a minimum price.

Either nodes type is un-satisfied when they are looking for or are offering service. Otherwise they are in a satisfied state. In the model a buyer node sends messages to its nearest neighbours informing it that it wants to purchase a resource unit at a price p_b. It also forwards messages from its neighbours until a preset time to live (TTL) has expired. Seller nodes look at incoming messages and decide whether their own price p_s can satisfy an incoming message. They also forward messages they can not satisfy to their neighbours. If a match can be made both nodes go into the satisfied state and the message is not forwarded any further by the Seller node. Later incoming possible matches are ignored. Satisfied nodes continue to pass on notes.

In the simulation updates of the nodes happen asynchronously. At each time step a node is picked randomly and the following actions take place. In [21] the authors show how parallel updates can induce system behaviour which is presumably artificial. It is also unrealistic to assume that all nodes have exactly the time globally.

- For a buyer node, the node checks whether its messages have timed out. If that is the case the node sends out new messages with a slightly in-

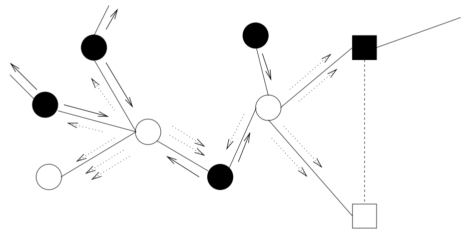

Figure 1: A partial sketch of a p2p system. Solid nodes are sellers and the others buyers. Circles indicate nodes looking for matches. Squares are satisfied nodes. The solid lines are peer relationships, the dashed line is satisfaction relationship. The solid arrows are first generation messages and the dotted ones second generations.

creased price. The node also forwards all messages that have arrived from neighbouring nodes.

- A seller nodes checks all its incoming messages and if a match can be made it contacts the node the message originated from. If that node is still in Buyer mode both nodes go into the satisfied state. Otherwise the Sellers carries on as usual. After this all messages are forwarded to its neighbours.
- A satisfied node passes all its incoming messages.
- In the dynamic model a satisfied returns to its previous mode. It also changes the state of it peer node. Buyer nodes lower their price to a fraction of what they have just paid, $p'_b = (1-\Delta p)p_b$. Seller nodes increase their price in the same way $p'_s = (1+\Delta p)p_s$.

Initially the nodes set into Buyer or Seller mode and the prices are drawn uniformally from two overlapping intervals.

4 Results and future Work

Initially the implementations of the models have only been run without a state change of the nodes once they have reached a satisfied state.This showed that the nodes are able to set prices to ensure that the maximum number of nodes reaches the satisfied state. Also, it indicated that in such a world with little dynamics areas of the network can loose out as they fail to find matches before the TTL times out.

The most extension to our model is the use of different network types to see how this influences the behaviour of the system. Other options are for instance to allow the network to rewire itself by introducing new links between successfully linked peers. To preserve the node degree distribution we would drop an existing link for both nodes. The choice could be made by measuring which link has been least useful in achieving the new link. In the real world the network will also be dynamic with respect to the life of the nodes. New nodes will appear old ones disappear. Another problem is that nodes will have a reliability distribution attached to themselves.

Interesting system are the utilisation of node. We would want to see what parameters increase or decrease the utilisation of the nodes and the overall system. Does this depend on the TTL choice, the pricing mechanism or is it completely independent of both? We will also investigate the number of transactions and of course the development of the price, locally and systemwide.

In another improvement one would also need to make the system more realistic by requesting by requesting a number of resources and have the seller be able to deal with more than one request at a time. This also opens up the interesting question what a realistic distribution of computing power is.

And of course does money matter? How does the system cope when money is abandoned and requests are simply fulfilled when a request comes in. Does this matter to the system utilisation?

Acknowledgments

I would like to thank Maya Paczuski, Vishal Sood, Kim Baskerville, Pete Harrison and many other people for ideas and discussion (Some of whom should arguably be co-authors). I also would like to thank the Complexity Group at the University of Calgary for its hospitality and the support of EPSRC (grant EP/D061717/1).

References

[1] I. Foster and C. Kesselman, "Computational Grids," *Chapter 2 of "The Grid: Blueprint for a New Computing Infrastructure", Morgan-Kaufman*, 1999. `http://www.globus.org/research/papers/chapter2.pdf`.

[2] U. Harder, P. Harrison, M. Paczuski, and T. Shah, "A dynamical model of a GRID market," tech. rep., Department of Computing, Imperial College London, October 2004.

[3] J. Darlington, C. Richardson, and J. Cohen, "Global Open Grid." http://www.internetcentre.imperial.ac.uk/projects.

[4] I. E. Sutherland, "A futures market in computer time," *Commun. ACM*, vol. 11, no. 6, pp. 449–451, 1968.

[5] I. W. Cotton, "Microeconomics and the Market for Computer Services," *Computing Surveys*, vol. 7, no. 2, pp. 95–111, 1975.

[6] N. R. Nielsen, "The Allocation of Computing Resources – Is Pricing the Answer?," *Communications of the ACM*, vol. 13, no. 8, pp. 467–474, 1970.

[7] T. Eymann, B. Padovan, and D. Schoder, "The catallaxy as a new paradigm for the design of information systems," in *Proceedings of the 16th IFIP World Computer Congress, Conference on Intelligent Information Processing*, 2000.

[8] O. Ardaiz, P. Artigas, T. Eymann, F. Freitag, L. Navarro, and M. Reinicke, "The catallaxy approach for decentralized economic-based allocation in grid resource and service markets," *Applied Intelligence*, vol. 25, no. 2, pp. 131–145, 2006.

[9] U. Harder, P. Harrison, M. Paczuski, and T. Shah, "A dynamic model of a GRID market," in *Proceedings of Invited poster at ACM/IEEE Mascots 2004 - ISBN 0-7695-2251-3*, October 2004.

[10] K. Lua, J. Crowcroft, M. Pias, R. Sharma, and S. Lim, "A survey and comparison of peer-to-peer overlay network schemes," *IEEE Communications Surveys & Tutorials*, pp. 72–93, 2005.

[11] K. Aberer, P. Cudre-Mauroux, A. Datta, Z. Despotovic, M. Hauswirth, M. Punceva, and R. Schmidt, "P-grid: A self-organizing structured p2p system."

[12] K. Lai, B. A. Huberman, and L. Fine, "Tycoon: A distributed market-based resource allocation system," 2004.

[13] K. Lai, L. Rasmusson, E. Adar, S. Sorkin, L. Zhang, and B. A. Huberman, "Tycoon: an implementation of a distributed, market-based resource allocation system," 2004.

[14] M. Feldman, K. Lai, and L. Zhang, "A price-anticipating resource allocation mechanism for distributed shared clusters," 2005.

[15] K. Lai, "Markets are dead, long live markets," 2005.

[16] M. Rosvall and K. Sneppen, "Self-assembly of information in networks," *Europhysics Letters*, vol. 74, p. 1109, 2006.

[17] G. Ghoshal and M. E. J. Newman, "Self-organizing information networks that can be searched in constant time," 2006.

[18] A. Oram, ed., *Peer-to-Peer: Harnessing the Power of Disruptive Technologies*, ch. 14 (Performance) by Theodore Hong. O'Reilly and Associates, 2001.

[19] A. Bagnall and G. Smith, "A multiagent model of the UK market in electricity generation," *Evolutionary Computation, IEEE Transactions on*, vol. 9, pp. 522– 536, October 2005.

[20] D. Harbord, N. Fabra, and N.-H. von der Fehr, "Modeling electricity auctions," Game Theory and Information 0206001, EconWPA, June 2002. available at http://ideas.repec.org/p/wpa/wuwpga/0206001.html.

[21] B. A. Huberman and N. S. Glance, "Evolutionary Games and Computer Simulations," *Proc. Natl. Acad. Sci. USA*, vol. 90, pp. 7716–7718, August 1993.