CATNETS – Open Market Approaches for Self-organizing Grid Resource Allocation

Torsten Eymann, Werner Streitberger, and Sebastian Hudert

Bayreuth University, 95447 Bayreuth, Germany
`torsten.eymann@uni-bayreuth.de`
`http://www.wi.uni-bayreuth.de/`

Abstract. Grid computing has recently become an important paradigm for managing computationally demanding applications, composed of a collection of services. The dynamic discovery of services, and the selection of a particular service instance providing the best value out of the discovered alternatives, poses a complex multi-attribute n:m allocation decision problem. Decentralized approaches to this service allocation problem represent a flexible alternative to central resource brokers, thus promising improvements in the efficiency of the resulting negotiations and service allocations. This paper analyses the impact of the service density on the profit and market price estimation using a decentralized service allocation mechanism in a grid market scenario.

Keywords: Self-Organisation, Economic Resource Allocation, Grid Service Allocation.

1 Introduction

Grid computing represents a concept for coordinated sharing of globally distributed resources spanning several physical organizations [1]. Currently the idea of Service-Oriented Architectures (SOAs) underlie several of the current Grid initiatives and reflect the common approach to realize Grid computing infrastructures, where participants offer and request application services. SOA defines standard interfaces and protocols that enables developers to encapsulate resources of different complexity and value as services that clients access without knowledge of their internal workings [2]. Grid systems have therefore increasingly been structured as networks of inter-operating services that communicate with one another via standard interfaces. Such infrastructures of services provided to an a priori unknown set of consumers can be efficiently organized as markets, analogously to traditional service markets in real world economies. Grid computing can thus become an object of Economics research, and thus provide insights not only for computer scientists, but also for economists.

The design and construction of resource allocation schemes is a particular research topic that can be tried (and evaluated) in globally distributed, large-scale Grid environments. Apart from computable general equilibrium approaches (NP-complete and thus not feasible) and all kinds of auctions (a Grid eBay?),

D.J. Veit and J. Altmann (Eds.): GECON 2007, LNCS 4685, pp. 176–181, 2007.

it becomes also possible to investigate in self-organization approaches. Self-organization can be found everywhere in our world, e.g. biological evolution, social group behaviour, market dynamics phenomena and other complex adaptive systems.

This article describes an investigation in implementing a self-organizing Grid Market based on the ”‘Catallaxy”’ concept of F. A. von Hayek [3]. Catallaxy describes a ”‘free market”’ economic self-organization approach for electronic services brokerage, which can be implemented for realizing service markets within service-oriented grid computing infrastructures. In such infrastructures, participants offer and request actual application services and computing resources for providing such services, of different complexity and value - creating interdependent markets:

- a service market - which involves trading of application services, and
- a resource market - which involves trading of computational and data resources, such as processors, memory, etc.

The distinction between resource and service allows different instances of the same service to be hosted on different resources. It also enables the price of a given service to base on the particular resource capabilities that are being made available by the hosting environment. In such interrelated markets, allocating resources and services on one market inevitably influence the outcome on the other market. This concept of two interrelated markets takes the current Grid concept one step further.

This paper investigates the general outcome of decentral resource negotiations in Grid systems. For this purpose a particular Grid environment is implemented and used for the actual simulation runs. Using Grid simulation software, different economic settings are investigated. The simulation results are evaluated using a defined set of metrics. The paper concludes with discussing the resulting metrics.

2 Related Work

The use of market mechanisms for allocating computer resources is not a completely new phenomenon. Regev and Nisan propose within the scope of the POPCORN project the application of a Vickrey auction for the allocation of computational resources in distributed systems [5].

Buyya motivated the transfer of market-based concepts from distributed systems to Grids [6]. However, he proposed classical one-sided auction types, which cannot account for combinatorial bids. Wolski et al. compare classical auctions with a bargaining market [7]. As a result, they come to the conclusion that the bargaining market is superior to an auction based market. Eymann et al. introduce a decentralized bargaining system for resource allocation in Grids, which incorporates the underlying topology of the Grid market [8].

Subramoniam et al. account for combinatorial bids by providing a *tâtonnement* process for allocation and pricing [9]. Wellman et al. model single-sided auction protocols for the allocation and scheduling of resources under consideration of different time constraints [10]. Conen goes one step further by designing a combinatorial bidding procedure for job scheduling including different running, starting, and ending times of jobs on a processing machine [11]. However, these approaches are single-sided and favor monopolistic sellers or monopsonistic buyers in a way that they allocate greater portions of the surplus. Installing competition on both sides is deemed superior, as no particular market side is systematically put at advantage.

3 Simulation Model

This section describes the Grid simulation model used to simulate the Catallactic free-market allocation approach. The CATNETS Grid simulator – an extension of the OptorSim Grid Simulator [15]– is used for simulation. The Grid network (GN) is defined by a connected non-oriented graph

$$GN = \langle S, L \rangle$$

with $S = n$ network sites and L a set of links which connect the sites with a bandwidth. The BRITE network generator is used to create the links between the sites [14]. Each site is characterized by a triple $\langle CSA_i, BSA_i, RA_i \rangle$ where CSA_i is a set of *Complex Service Agents (CSAs)*, BSA_i is a set of *Basic Service Agents (BSAs)*, and RA_i is a set of *Resource Agents (RAs)*. In every site there can be zero or more complex/basic service agents and zero or more resource agents.A node with no associated agents is a *router*.

Complex Service Agents. CSAs are entry points to the Grid system and are able to execute *Complex Services (CSs)* for Grid clients. A CS is defined as a set of *Basic Services (BSs)*. CSAs are not specialized: they accept any type of complex service request and take care of the execution of the component basic services. For simplicity reasons, a complex service requests always one basic service in the evaluated scenario. Several CSAs are available in the network, which enable parallel allocation and execution of BSAs.

Basic Service Agents. BSAs provide CSAs with the BSs they need to furnish their complex services to Grid clients. A predefined number of BSAs is available for selection of the CSAs.

Resources Agents. Resources have a *name* which is a unique identifier whose intended semantics is shared among all agents. Every resource is also characterized by a *quantity* whose value is a positive integer. RAs are “proxies” for aggregations of resources. Their task is to provide BSAs with resources needed for the execution of basic services. For simplicity reasons, a RA provides only of one unit of their resource and the BSA requests one unit from RA.

4 Simulation Scenario

The simulation scenario analyses the impact of the agent’s density on the outcome of the market. The total number of agents changes between the simulation scenarios from 30 and 60 to 300, while keeping the number of Grid sites fixed. In detail, the number of agents is equally split in CSAs, BSAs and RAs. For example, a total number of 60 agents means a set of 20 CSAs, 20 BSAs and 20 RAs. The agents are distributed over the 30 Grid sites using a uniform distribution. The bandwidth between the sites is set to 1 Gbit/s, which guarantees similar communication conditions.

All agents' price intervals are initialized in an interval between 80 and 180, with an interval length of 30. The lower bound for BSAs is drawn from a uniform distribution. Its upper limit is obtained as a simple addition of 30 to the lower limit value. The initial price interval for the CSAs is computed drawing a number from the random price interval and a subtraction of 30.

A CSA requests a basic service by broadcasting a call-for-proposal message to BSAs, which is received by all BSAs reachable within 2 hops in the given network topology.After a discovery timeout of 500 ms, the requesting CSA selects one BSA for negotiation. A best price selection policy picks the best offered proposal and starts the negotiation by iterative and bilateral message exchange, until one party accepts or rejects.

5 Evaluation

The simulation scenario is evaluated with two metrics on the population level. Figure 1 shows the profit and the estimated market price of the complex service and basic service agents during one simulation run in different settings. The profit is computed as the difference between negotiation price and the estimated market price. The goal of the agents is to optimize their profit. The estimated market price is computed as a weighted average of historical agreement prices.

The analysis concentrates on the evaluation of the service market; similar results are found at the resource market. In each simulation run, 1000 complex services issues requests, the delay between the requests was 1000 ms and the execution time of one basic service was set to a constant 1000 ms. Each diagram shows the BSAs as service providers and CSAs as service consumers for a population of 30, 60 and 300 agents.

In the smallest scenario, the profit of the agents converges fast to values near zero. Both, complex services and basic services trade very often and estimate the market price very well. They use their good market price estimation to optimize their trading strategy. None of the trading partners is able to make a large profit. Only at the beginning of the simulation the profit/loss peaks indicate wrong market price estimations due to insufficient knowledge on market prices. This effect increases with 60 agents and leads to high deviations in the 300 agent scenario, as the necessary amount of information needed for feedback learning decreases *per capita*. Only sellers are able to make a profit - they capitalize on the (far too low) market price estimations of the complex services and realize high profits. At the end of the simulation run, the buyers slightly increase their profits due to better market price estimations. The number of 1000 requests is not enough for all agents to learn the market price and optimize their behavior within a population of 300 agents.

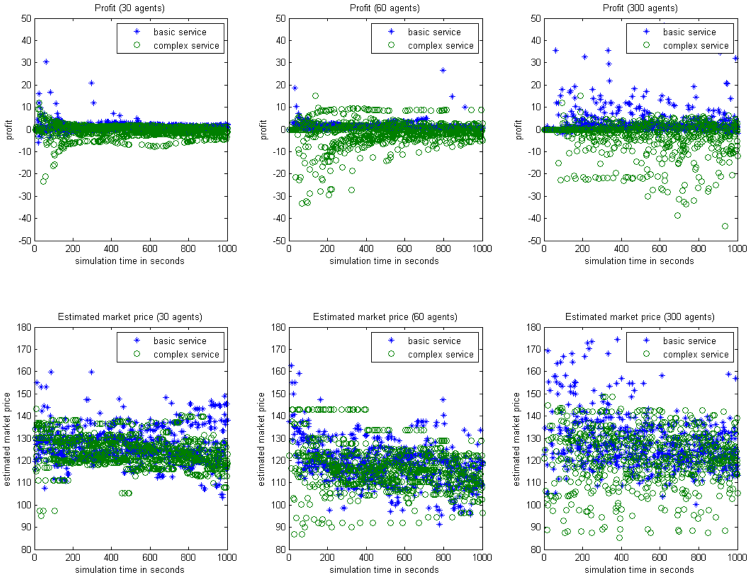

Fig. 1. Profit and estimated market price for 1000 requests and 30, 60 and 300 agents

6 Conclusion and Outlook

The paper presented an evaluation of a Catallactic free-market strategy, where no central resource broker exists. Different scenario settings are evaluated dependent on the service/agent density. In scenarios with a high number of agents, the Catallactic strategy is slow in converging to a stable market price estimation and profit of the market participants, because the individual agents interact not enough to arrive at a stable market price estimation. Further experiments increasing the number of requests have to be performed to analyze its impact on the agent's profit.

Acknowledgements

This work has partially been funded by the EU in the IST programme "Future and Emerging Technologies" under grant FP6-003769 "CATNETS".

References

1. Foster, I., Kesselmann, C.: The Grid: Blueprint for a New Computing Infrastructure. Morgan Kaufmann Publishers Inc. San Francisco, CA (1999)
2. Foster, I.: Service-Oriented Science. Science Journal 308, 814–817 (2005)
3. von Hayek, F.A.: The Use of Knowledge in Society. American Economic Review XXXV 4, 519–530 (1945)

4. Eymann, T., Sackmann, S., Mueller, G., Pippow, I.: Hayek's Catallaxy: A Forward-looking Concept for Information Systems. In: Proc. of American Conference on Information Systems (AMCIS'03) (2003)
5. Regev, O., Nisan, N.: The POPCORN Market: Online Markets for Computational Resources. In: Decision Support Systems, vol. 28, pp. 177–189. Elsevier Science Publishers, Amsterdam (2000)
6. Buyya, R., Stockinger, H., Ghiddy, J., Abramson, D.: Economic Models for Management of Resources in Peer-to-Peer and Grid Computing. In: Proceedings of the International Conference on Commercial Applications for High Performance (2001)
7. Wolski, R., Brevik, J., Plank, J.S., Bryan, T.: Grid Resource Allocation and Control using Computational Economies. In: Berman, F., Fox, G., Hey, A. (eds.) Grid Computing: Making The Global Infrastructure a Reality, pp. 747–772. John Wiley & Sons Publishers, West Sussex (2003)
8. Eymann, T., Reinicke, M., Ardaiz, O., Artigas, P., de Cerio, L.D., Freitag, F., Messeguer, R., Navarro, L., Royo, D.: Decentralized vs. Centralized Economic Coordination of Resource Allocation in Grids. In: Proceedings of the 1st European Across Grids Conference (2003)
9. Subramoniam, K., Maheswaran, M., Toulouse, M.: Towards a Micro-Economic Model for Resource Allocation in Grid Computing Systems. In: Proceedings of the 2002 IEEE Canadian Conference on Electrical & Computer Engineering (2002)
10. Wellman, M., Walsh, W., Wurman, P., MacKie-Mason, J.: Auction Protocols for Decentralized Scheduling. Games and Economic Behavior 35, 271–303 (2001)
11. Conen, W.: Economically Coordinated Job Shop Scheduling and Decision Point Bidding - An Example for Economic Coordination in Manufacturing and Logistics. In: Proceedings of the 16th Workshop on Planen, Scheduling und Konfigurieren, Entwerfen (2002)
12. Parkes, D., Kalagnanam, J., Eso, M.: Achieving Budget-Balance with Vickrey-Based Payment Schemes in Exchanges. In: Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence, pp. 1161–1168 (2001)
13. Biswas, S., Narahari, Y.: An Iterative Auction Mechanism for Combinatorial Logistics Exchanges. In: Proceedings of the 9th International Symposium on Logistics (2004)
14. BRITE: Boston university Representative Internet Topology gEnerator, `http://www.cs.bu.edu/brite/`
15. OptorSim: `http://sourceforge.net/projects/optorsim`