A Decentralized Grid Market Infrastructure for Service Oriented Grids*

Isaac Chao[a], Rene Brunner[a] , Felix Freitag[a], Leandro Navarro[a] Pablo Chacin[a]

Oscar Ardaiz[b]

Liviu Joita[c], Omer F. Rana[c]

[a] Computer Architecture Department, Polytechnic University of Catalonia, Spain

{ichao,rbrunner,felix,leandro,pchacin}@ac.upc.edu

[b] Department of Mathematics and Informatics, Public University of Navarra, Spain

oscar.ardaiz@unavarra.es

[c] School of Computer Science and the Welsh eScience Centre, Cardiff University, UK

{l.joita, o.f.rana} @cs.cardiff.ac.uk

Abstract. Service Oriented Computing has a deep impact on how IT infrastructures are conceived both in academia (e-science) and in industry (Service Oriented Architecture and commercial Web Services). Increasingly, economic models are being considered as suitable coordination mechanism for the management of service allocations to clients. However, few complete infrastructures have demonstrated the enabling of economic-based Service Oriented Grids (SOGs). We propose a complete infrastructure for economics-based SOGs and we demonstrate its application in a prototype. We conduct experiments showing that practical agent-based automatic and fair trading of services at stable prices can be achieved using the proposed infrastructure.

Keywords: Automatic Resource Allocation, Contract-Net, Decentralized Economic Models, Service Oriented Grids

1 Introduction

Grid Computing leverages the power of thousands of resources distributed across computers/supercomputers/clusters linked by networks (from intranets to the Internet). Through the concept of Virtual Organizations (VOs), the Grid enables the dynamic composition of such resources into interoperable services, which potentially multiply

* This work was supported in part by Spanish Government under Contract TIC2002-04258-C03-01, and the European Union under Contract CATNETS EU IST-FP6-003769

exponentially the VOs added value. However, given the unpredictability of the underlying platform (Internet), scalable realization of such synergies (in both physical and organizational levels) poses serious challenges to modern large scale distributed systems research. Contrarily to other distributed systems, Grids have many independent resource providers with varying access policies. In addition to large sizes, the diversity of policies leads to a very complex allocation task that cannot be handled manually by users. Automatic and adaptive resource management is the solution to these challenges.

The primary visions for Grid computing are utility computing infrastructure and Grid services/service providers. In utility computing, a third party service provider hosts and manages the Grid solution dedicated to serving a single organization or the needs of multiple ones. Customers only pay for the used resources. Grid services/service providers modularity enables the dynamic composition and coordination of e-services which can be exchanged or traded between Grid users or brokers, following the usage models from utility computing. These two features enable for the automatic trading of Grid Services in SOGs.

There has been an increase of interest in SOGs recently within the Grid community towards services that are often considered as a natural progression from component based software development, and as a mean to integrate different component development frameworks. A service in this context may be defined as a behavior that is provided by a component to be used by any other component. A service stresses interoperability and may be dynamically discovered and used [Foster, 2002, p. 37-46]. Decentralized Grid Markets based on agents have been proposed as suitable coordination mechanisms for Grids and SOAs [Eymann, 20005, p. 297-307]. The market here is nothing more than a communication bus – it is not a central entity of its own and does not participate in matching participants' requirements using some optimization mechanisms. Direct agent to agent bargaining allows participants to use the negotiation strategy more suitable to its objectives and current circumstances. Local bilateral bargaining also facilitates the scalability of the system and the quick adaptation to fluctuations in resource allocation dynamics. This enables for high scalability in both physical and organizational levels. These concepts have been capitalized in [Eymann, 2005b] for SOG purposes.

In this paper, we address the design of an economic SOG infrastructure and its evaluation based on the Grid Market Middleware (GMM), a resource allocation middleware which incorporates decentralized economic models [Ardaiz, 2005]. We show how the GMM can be used as a modular infrastructure for service oriented application of decentralized market models to Grid resource allocation. Integration of the GMM in a SOG is demonstrated by means of a prototype. A Data Mining application interfaces seamlessly with the GMM, which incorporates a decentralized economic algorithm based on the Contract-Net protocol. The GMM is offered as a Web Service via a convenient access point. For the application, the access to utility based computing is completely transparent and accessed just as a conventional Web Service. In addition, we conduct experiments for this SOG scenario, showing the fair automated trading of services as a result of fully decentralized negotiations. The economic algorithm, though very simple, is able to dynamically adapt to changes in offer and demand. This has several advantages over centralized resource trading mechanisms, which do not scale easily to large and dynamic environments.

2 Related Work

Economy based resource allocation has received a great deal of attention in the last years. The GridBus Project [Buyya, 2005] is a reference in SOG and utility based computing, and has proposed a great variety of market models and tools for the trading of Grid Resources. However, its strong emphasis on computational intensive Grids and the hierarchical nature of some of the proposed components, like the Grid Market Directory, diverges from the fully decentralized resource allocation mechanism proposed here. Centralized approaches exist such as [Schnizler, 2006], but scalability issues both in size and computational requirements further complicate its applicability to large size Grids. Tycoon [Lai, 2004] is a market-based system for managing compute resources in distributed clusters or Grids. It uses distributed auctions with users having a limited amount of credits. Users who provide resources can, in turn, spend their earnings to use resources later. Another approach which combines both economic-based Grid schedulers and adaptive learning agents is presented in [Li, 2006, p. 567-583]. In this paper, the authors use a learning mechanism such as Q-Learning in order to select the negotiation strategy applied, depending on the evolution of the environment and the negotiation partners.

A few papers address fully decentralized market mechanisms for computational resources. In [Despotovic, 2004, p. 7], a P2P double auctioning mechanism is proposed which builds on Zero-Intelligence agents (ZIP Agents) [Presit, 1998]. It was shown that the results with original ZIP agents in continuous double auction (CDA) depend strongly on the availability of the complete set of bid and offers coming from all buyers and sellers, and the commitment to winner-to-winner allocations. But a P2P or fully decentralized trading mechanism must be free of any central authority for scalability reasons. Interestingly, synchronization issues arise when trying to adapt the ZIP agent's to a P2P environment. Notably the hard constraints on information availability need to be relaxed. As a result, convergence and price stability of the markets are no longer assured. Another fully decentralized approach is the one adopted with the catallactic agents [Eymann, 2000]. In this approach bilateral negotiations are established between a set of learning agents, and the spontaneous coordination arises from both the bargaining and co-evolutionary learning processes.

However, none of these approaches provide the infrastructure for integrating explicitly the market based algorithms into service oriented Grids. Tycoon has been used mostly in a clusters environment, and GridBus is provided as complete software toolkit, not as a service. Our approach is to offer the economic algorithms as Web Services for a seamless integration in any SOG.

3 Service Oriented Grid Infrastructure

3.1 The Grid Market Middleware

The GMM provides the mechanisms to register, manage, locate and negotiate for services and resources. It allows trading agents to meet each other based on its requirements and engage in negotiations. Furthermore, the middleware offers a set of generic negotiation mechanisms, on which specialized strategies and policies can be dynamically plugged in. A detailed description of both the design and implementation of the GMM architecture can be found in [Ardaiz 2005].In this version of the GMM, we have plugged in as economic agents an implementation of the Contract-Net protocol, standardized by FIPA [FIPA, 2007]. The Contract-Net protocol has been developed to specify problem-solving communication and control for nodes in a distributed problem solver. Task distribution is affected by a negotiation process, a discussion carried on between nodes with tasks to be executed and nodes that may be able to execute those tasks. This protocol has been applied to many domains in multiagent systems [Paurobally, 2004]. An example of application of the protocol to Grid resource allocation can be found in [Chao, 2004].

The Contract-Net protocol starts with a task announcement phase by the initiator (the buyer), which can be answered by one or more participants (the sellers). This announcement is carried out by a groupcast of a call for proposals (CFP). After conclusion of this period, the initiator selects from the set of collected proposals the best one, informing the winner. In top of this protocol, we apply a simple offer/demand-based economic algorithm: The sellers will answer the CFPs which meet its current selling price. If the CFP does not meet its requirements, the seller will lower its expectations and it will decrease the selling price. As for the buyers, if a seller rejects the CFP, then it will lower its expectation by increasing the offer in the next CFP. Both the buyers and the sellers will increase their expectations in case of receiving offers/bids which meet their expectations. The price updating is done at fixed small price steps.

In the context of the GMM, the buyer agents are called ComplexServices (CSs) and the seller agents BasicServices (BSs). CSs aggregate BSs from the market. As BSs and CSs get involved in trading, the price will evolve by offer and demand, with dependence on the limited CS budget and the limited resources which can be sold by the BS. If a BS is selling its resources, it will increase the prices to test to what extend it can increase the profit. When it no longer sells, it will lower the price until it becomes competitive again. The same mechanism applies for the CSs. The dedicated resource model translates for the market concerns in a limited offer by the BSs. Once the BS has sold its resource to a CS, it cannot accept more bids from other CSs' CFPs until the moment when the client of the awarded CS ends the execution of the sold Data Mining service in the resource. This peculiarity is what makes this scenario novel and interesting with respect to other previous research on decentralized economic algorithms based on simple learning agents in idealized markets [Preist, 1998], [Despotovic, 2004].

3.2 GMM as a Service

In a SOG infrastructure, the GMM is exposed to be accessed by applications trough a convenient access point, a Web Service. Figure 1 describes the main steps in the interaction trough the access point. When a client issues a request, the application determines which Grid services are required to fulfill it. These Grid services represent either software services (e.g. a data processing algorithm) or computational resources. The application translates these requirements into a standardized WS-Agreement [WS-Agreement, 2007]. The application invokes the access point and passes the corresponding WS-Agreement request. This is in turn parsed and processed at the access point, which instantiates the GMM with the required economic agents to fulfill the client request.

The GMM searches among the available service providers. When a suitable service provider is found, the application requirements are negotiated within the middleware by agents who act on behalf of the service providers as sellers and the application as buyers. Once an agreement is reached between the trading agents, a Grid service instance is created for the application and a reference is returned to the application/client, which can invoke it. The server-side infrastructure is deployed by a set of scripts which allow for the bootstrapping of BSs in available resources. The scripts perform the automatic deployment and configuration of the BSs, which are then ready to be contacted by CSs. Services offered by BSs for clients executions are also deployed and exposed in Apache Tomcat application servers. Complemented by the access point, this comprises a complete infrastructure for economic-based SOGs.

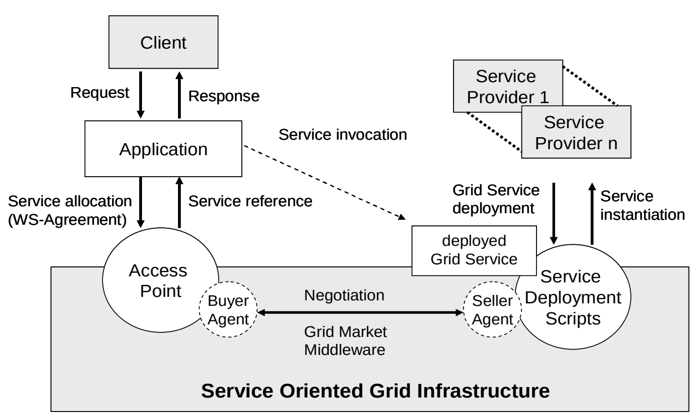

Figure 1. Service Oriented Grid (SOG) infrastructure.

4 Prototype Application

As a proof of concept of the system model, we provide an application of the GMM with Contract-Net to an existing decentralized "free-market" prototype, the Catallactic Data Mining application [Joita, 2006, p. 146-153]. The basic problem addressed by the data mining process is one of mapping low-level data (which are typically too voluminous to understand) into other forms that might be more compact (for example, a short report), more abstract (for example, a descriptive approximation or model of the process that generated the data), or more useful (for example, a predictive model for estimating the value of future cases). At the core of the process is the application of specific data-mining methods for pattern discovery and extraction. This process is often structured into a discovery pipeline/workflow, involving access, integration and analysis of data from disparate sources, and to use data patterns and models generated through intermediate stages. Selection and conversion of datasets as well as the execution of the data-mining algorithm itself are the typical steps required. In the Catallactic Data Mining services prototype, two Data Mining Services encapsulating data conversion and algorithm execution are combined in a workflow achieving a solution to the overall problem.

Consider a scenario where a client issues sequential requests for Data Mining services. The CSs try to map the incoming workflows to available set of services. The BSs, try to sell their services to the CSs which are instantiated after successful negotiation upon the client request. Figure 2 shows a scenario with two service types in the service market, as well as the pipeline of services to be traded in the market.

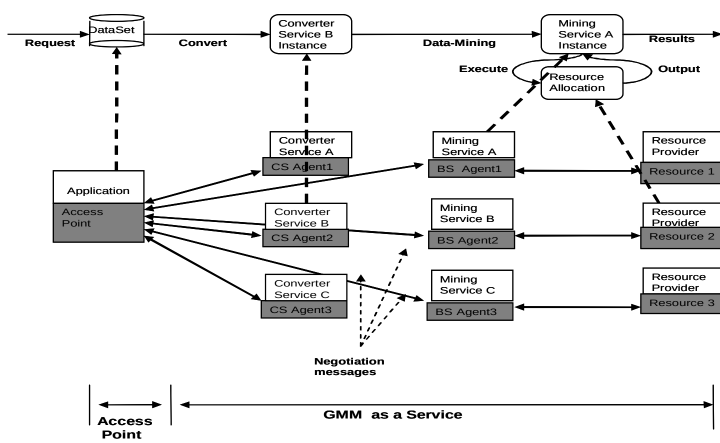

Figure 2. Prototype Application using the SOG infrastructure.

5 Experiments

5.1 Experimental setup

The goal of the experiments is to show the performance of the GMM as an automated economic-aware resource management tool by means of the the DataMining Grid prototype application. We evaluate the ability of the Contract-Net based negotiation protocol for stabilizing fair prices in the Grid service, trading in different scenarios. We setup controlled experiments by deploying several instances of the GMM in a Linux server farm. Each machine has a 2 CPU Intel Xeon at 2.80GH and 2 GB of memory. The nodes in the farm are connected by an internal Ethernet network at 100Mps. The topology is a mesh: All interconnected. CFPs are transmitted via group-cast to all the nodes in the destination groups (in our scenario CFPS are groupcasted from CSs to BSs).

We deploy the GMM in 4 nodes. Two nodes host a BS each and the Data Mining Web Service and other two nodes host the CSs, access points and clients. The Web Services are exposed in Tomcat servers. Access for execution of these Web Services on the resource node is what is traded between BSs and CSs. The experiments consist in launching 2 clients concurrently, which use each one of the CS as broker. Each client makes 100 requests to the CS in intervals of 2 seconds. Whenever a CS wins a bid with a BS, it invokes the Data Mining Service in the selected node, and the resource in the corresponding node gets locked for the duration of the service execution. We measure the selling prices of the BSs and observe the proportion of successful CFPs issued by the CSs.

5.2 Experimental results and evaluation

We have two different scenarios in the dedicated resource model. The demand indicates the rate at which new CFPs are issued. If the proportion demand rate/Data Mining Web Service execution time is lower than 1, then the resources are potentially able to handle all the demands. In the contrary case, the demand exceed the offer of resources, hence several CFPs will be disregarded even in the case they meet the pricing criteria of the BS.

In a first experiment, we set up a demand rate lower than 1. Resources therefore are able to cope comfortably with the demand. We set up two CSs with initial bidding prices of 75, and 2 different BSs, with initial offer prices of 60 and 90 respectively. In Figure 3(left diagram) it can be seen that the price quickly stabilizes to a “fair” value around 72. This result holds true independently of the initial prices on the CSs. As for the initial prices of the BSs, the BS1 starting with a lower price trades more in the beginning, but trying to increment its surplus it soon reaches the equilibrium price. As for the BS2 starting with a higher price, it does not trade in the first CFPs which leads to a continuous price dropping towards the equilibrium price. After stabilization, the CSs get their CFPs for resources granted, provided the bid equals at least the offered price, which is true during most of the experiment.

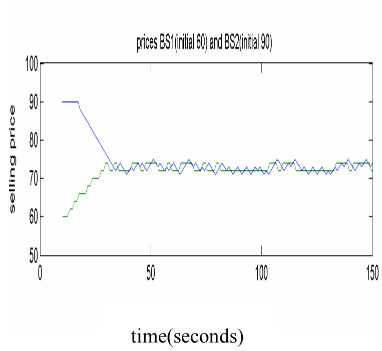

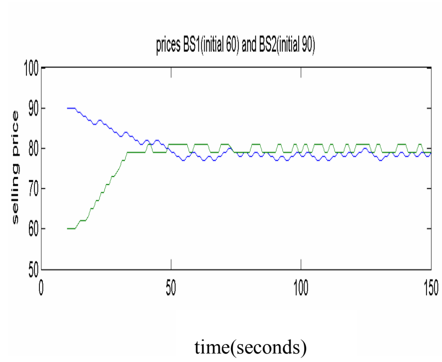

Figure 3. Evolution of prices vs time for a low (left) and high (right) demand rate.

In a second experiment, we set up a high demand. In this case the resources of the providers are not enough to completely meet the demand. This does however not make the prices increasing indefinitely, since the successful trades make the CSs react trying to decrease bids. We set up agents with the same initial prices as in the previous example. The price stabilizes also quickly, but in this case to a higher price at around 80, due to the resources scarcity (Figure 3, right diagram). This result holds true independently of the initial prices on the CSs. In this scenario, the CSs gets just half of their CFPs for resources granted, but this is evenly distributed between the two CSs, which is a fair resource share.

The experiments demonstrate how a simple decentralized economic algorithm can be plugged into the GMM infrastructure. These results extend state-of-the-art economic algorithms literature in both theoretical and practical aspects. Though the results needs further validation, the contract-Net based simple offer/demand economic algorithm achieves in simple scenarios stable fair trading of resources without the need of complete information while biding, as happens in alternative algorithms such as ZIP agents. This has interesting implications for scalability. As for practical usage of the prototype, the results correspond to a dedicated resources scenario where a limited pool of resources (once per BS) is available in the market, which maps well to several realistic scenarios in High Performance Grid Computing.

6 Conclusions

We have shown a complete infrastructure for economic-based SOGs and we have demonstrated its application in a Data Mining SOG prototype. The proposed infrastructure provides both the scripts for automatic bootstrapping of traded Grid services and the agents selling the services at the Grid service provider side, as well as the Web Service access point for the seamless usage of the SOG infrastructure by clients.

The modularity of the infrastructure which enables the GMM to plug in different decentralized economic algorithms has also been demonstrated. We have complemented previous decentralized economic models implemented in the GMM economic algorithms layer with a new one based on the Contract Net protocol.

The performance results show that agent-based automatic fair trading of resources at stable prices can be achieved using the GMM. In that sense, resource scarcity translates in the market to the stabilization of higher selling prices for the resources, as would be desired in a fair market. The economic mechanism does not impose any constraint on system size or computational requirements, hence enabling for high scalability in both physical and organizational dimensions.

Future work comprises the inclusion of more complex workflows and its integration/evaluation in the architecture, increasing the size of the test bed, and test the infrastructure with additional prototype applications.

7 References

[Ardaiz, 2005] O. Ardaiz, P. Chacin, I. Chao, F. Freitag, L. Navarro, An Architecture for Incorporating Decentralized Economic Models in Application Layer Networks, Internacional Workshop in Smart Grid Technologies, Utrecht, Holanda, July 25 - 29, 2005.

[Buyya, 2005] Buyya, R. Abramson, D. Venugopal, S, The Grid Economy, Proceedings of the IEEE pages 698- 714 Volume: 93, Issue: 3 ISSN: 0018-9219, 2005

[Chao, 2004] Isaac Chao,Ramon Sangüesa and Oscar Ardaiz, Design, Implementation and Evaluation of a Resource Management Multiagent.System for a Multimedia Processing Grid, Workshop on Grid Computing and Its Application to Data Analysis (GADA) On the Move to Meaningful Internet Systems 2004:

[Despotovic 2004] Z. Despotovic, J.-C. Usunier, K. Aberer: Towards Peer-To-Peer Double Auctioning, Proceedings of the 37th International Hawaiian Conference on System Sciences (HICSS), Waikoloa, Hawaii, USA, 2004

[Eymann, 2000] T. Eymann, B. Padovan, and D. Schoder. The catallaxy asa new paradigm for the design of information systems. In16th IFIP World Computer Congress, Beijing, China, August2000.

[Eymann, 20005] Torsten Eymann, Michael Reinicke, Werner Streitberger, Omer Rana, Liviu Joita, Dirk Neumann, Björn Schnizler, Daniel Veit, Oscar Ardaiz, Pablo Chacin, Isaac Chao, Felix Freitag, Leandro Navarro, Michele Catalano, Mauro Gallegati, Gianfranco Giulioni, Ruben Carvajal Schiaffino, Floriano Zini, "Catallaxy-based Grid Markets", Multiagent and Grid Systems Issue: Volume 1, Number 4 / 2005 Pages: pp. 297 - 307 Special Issue: Smart Grid Technologies & Market

[Eymann, 2005b] A hayekian self-organization approach to service allocation in computing systems. ADVANCED ENGINEERING INFORMATICS , 19 (3) : 223-233. ISSN: 1474-0346

[FIPA, 2007] FIPA webpage: http://www.fipa.org/

[Foster, 2002] Foster, I.; Kesselman, C.; Nick, J.M.; Tuecke, S., Grid services for distributed system integration, Computer, Vol.35, Iss.6, Jun 2002, Pages:37-46

[Joita, 2006] Liviu Joita, Omer F. Rana, Pablo Chacin, Isaac Chao, Felix Freitag, Leandro Navarro, Oscar Ardaiz (2006) - "A Catallactic Market for Data Mining Services", in International Journal of Future Generation Computer Systems (FGCS) - Grid Computing:Theory, Methods & Applications, Volume 23, Issue 1, January 2007, ISSN 0167-739x, pp. 146-153

[Lai, 2004] Kevin Lai, Bernardo A. Huberman and Leslie Fine, "Tycoon: A Distributed Market-based Resource Allocation Systems", Technical Report arXiv:cs.DC/0404013, April 5, 2004

[Li, 2006] J.Li, R.Yahyapour. Learning-Based Negotiation Strategies for Grid Scheduling In: IEEE International Symposium on Cluster Computing and the Grid (CCGrid 2006), pages 567-583, IEEE Press, 2006

[Paurobally, 2004] Shamimabi Paurobally, Jim Cunningham, and Nicolas R. Jennings. Verifying the contract net protocol: A case study in interaction protocol and agent communication language semantics. In Proceeding 2nd International Workshop on Logic and Communication in Multi-Agent Systems, 2004

[Preist, 1998] C. Preist, M. van Tol, Adaptive agents in a persistent shout double auction. In Proceedings of the First international Conference on information and Computation Economies (Charleston, South Carolina, United States, October 25 - 28, 1998). ICE '98. ACM Press, New York, NY, 11-18

[Schnizler, 2006] B. Schnizler, D. Neumann, D. Veit, C. Weinhardt (2006) Trading Grid Services - A Multi-attribute Combinatorial Approach European Journal of Operational Research, in Press

[WS-Agreement, 2007] Web Services Agreement Specification (WS-Agreement), https://forge.gridforum.org/sf/projects/graap-wg