CORE SERVICES FOR GRID MARKETS

Pablo Chacin*, Xavier Leon, Rene Brunner, Felix Freitag and Leandro Navarro
{ pchacin,xleon,rbrunner,felix,leandro } @ac.upc.edu
Computer Architecture Department,
Technical University of Catalonia, Spain

Abstract Markets are a powerful model for the coordination of distributed systems and, in particular, in the face of incomplete information and changing environments. The application of markets for the resource allocation in grid systems has recently been researched as an alternative to traditional approaches. However, the proper implementation of sophisticated markets capable of handling diverse trading models (various auctions types, bargaining) and structures (direct negotiation, brokering, etc.) requires a set of supporting services to provide participants a proper environment to engage in negotiations. Grid Market Middleware (GMM) is a framework that aims to ease the development of market based grid systems. In this paper we present its architecture, the services it provides and describe how they can be used to implement diverse market models. We also discuss our experience with the implementation of prototypes for various core services.

Keywords: Grid Market, Economic Models, Middleware.

*Corresponding author. Phone:+34(93)405 40 59, Fax:+34(93)401 70 55

1. Introduction

Markets are decentralized, goal-oriented mechanisms for allocating resources among competing interests while meeting some general goal, like the global utility of the users. A distributed system set up along market rules can adapt to changes in the environment, variations in the demand or supply of resource, or disturbances to individual members [10]. Many projects have used market approaches for resource allocation (See [24] for a recent survey).

Despite of this interest, there is a significant absence of general-purpose frameworks to develop market oriented systems. Most of the proposed frameworks are specific of a market model (e.g. a specific type of auctions) or problem domain (e.g. logistic optimization, robotics), or are not suited for large scale fully decentralized deployments.

The intent of the Grid Market Middleware project (GMM)[1] is to provide such a generic framework. More specifically, the GMM aims to address the problems of a general infrastructure for decentralized markets, a scalable architecture and a high level programming abstractions independent of the market model. This work is based on our experience implementing market based resource allocation middleware for Application Layer Networks and its integration with grid applications [5, 13], and in our ongoing effort to implement the foundations for a grid market in the context of the SORMA[2] and Grid4All[3] European projects.

In this paper we present the design of the GMM's core services and how they can be used to implement diverse market models. The document is organized as follows: section 2 presents the general model of a GMM based grid market. Section 3 presents the general architecture and introduces its main components. Section 4 shows how diverse market models can be implemented using the GMM. Finally, section 5 presents some related work and 6 presents our conclusions and future work.

2. Model

The general model of a GMM based grid market can be seen in figure 1. In this model, we differentiate three main components: (a) participants, which engage in economic exchanges to sell/buy/allocate resources, (b) the Economically Enhanced Resource Managers (EERM), which manage those resources and (c) Core Services that provides general functionalities for market oriented systems and a programming interface that facilitates the implementation of the participants and the EERM.

[1] http://recerca.ac.upc.edu/gmm
[2] http://www.sorma-project.eu
[3] http://www.grid4all.eu

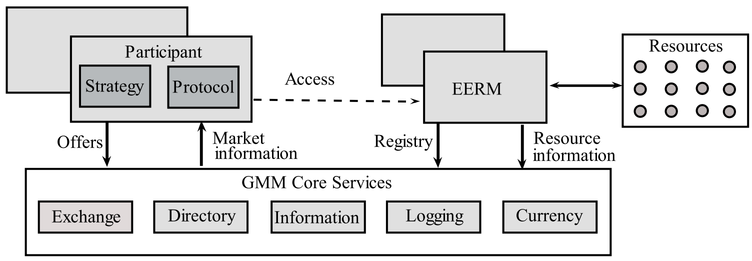

Figure 1: Model for a GMM based grid market.

Participants are agents (in the general sense) that behaves on behalf of resource providers, resource consumers, or meditate between them. Participants are responsible for gathering and evaluating market information and deciding their strategies to sell/buy (e.g. pricing). GMM does not make any special distinction among sellers and buyers. This is important to allow a participant to become an intermediary, buying resources and then reselling them and behave as a mediator (broker, arbiter) or a market maker which improves the scalability of the market and improves its liquidity [4].

Markets require a resource allocation mechanism [11], a protocol to allocate resources among participants. This mechanism can be embedded in the buyers and seller themselves (when they use direct bargaining or single side auctions for trading) or might also be run as a separate participant (when auctioneers are used). GMM does not enforce any mechanism, but it does enforce and tracks the primitive messages on which such mechanisms could be implemented.

The EERM provides the capabilities to access grid resources from the grid market. It registers resources in the market and provides information about its availability and relevant performance metrics, which is integrated with the market information. EERM also serves as a gateway to access resources, verifying that the intended access are backed by a previous agreement between the parties (provider and consumer). The EERM is mostly platform dependent and is not provided by the GMM. However, the interfaces to integrate them are provided by the GMM.

3. Architecture

GMM's architecture has been designed under two guiding principles: (a) integrate under a common framework the services related to information gathering and dissemination in decentralized market and (b) take advantage of the functionalities provided by P2P overlays to organize a distributed system and allow efficient communications and decentralization. The resulting architecture is arranged in four main layers:

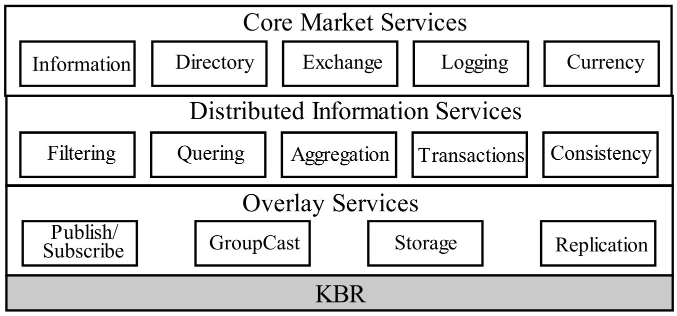

Figure 2: Architecture of the GMM

- **Core Market Services:** market specific services that supports the development of participants, enabling them to engage in negotiations for resources.
- **Distributed Information Services:** generic services that allows an efficient management of information in fully decentralized deployments. Offers services for processing queries and their responses, filtering messages, aggregating information and ensuring consistency and transactional access to critical data.
- **Overlay Services:** provides sophisticated communication and cooperation mechanisms like publish/subscribe, group casting, distributed storage of data items (DHT), and replication of critical elements .
- **KBR:** offers a Key Based Routing mechanism to communicate different nodes based on logical, location independent keys [7].

In this paper we focus on the components of the Core Services Layer, which are explained below:

Exchange Service. Provides a trading infrastructure designed to support different market-based systems. It defines a set of primitive messages that can be used to implement complex negotiation protocols and provides transparent message routing among participants. The utilization of these market primitives allows the Market Exchange to accommodate different market models, even simultaneously and ensures that the infrastructure can apply rules for routing, validation and security. It also frees the participants from the need to validate each received message.

Market Directory Service. Provides a decentralized, market wide registry for participants of the market (providers, consumers) and the resources/services been traded. It offers a generic framework on which diverse specialized strategies can be plugged to adapt and optimize the discovery to the characteristics such as the density of resources and its distribution in the network [12], and support diverse query languages.

Market Information Service. Provides current aggregated information and historical statistics of market indicators, under publish/subscribe and query interfaces. Examples of such information are the level of activity in a market (indicated by the number of available products and the volume of traded products) and the current maximum, minimum or average price for a resource type. Trader agents will rely on this information to create or adapt their internal strategy like choosing the right market to negotiate for resources, finding the optimal time to enter or create markets, and to propose competitive bids.

Logging Service. Keeps a registry of the transactions occurred during a particular negotiation, and maintains the negotiation state across restarts of the Exchange Service. It can be used for accounting, dispute resolution and security purposes. To achieve these goals, it offers a secure and reliable storage of the messages that were exchanged during a negotiation. Transactions can be retrieved using the negotiation id or the specific transaction id as a key. More complex queries are not currently considered but can be implemented by indexes which can be stored internally in the DHT Layer (following a model similar to that of [9]) or in external repositories like relational data bases.

Currency Service. Is a distributed banking service for the Grid which enables users to perform and receive payments for resources usage and sharing using a virtual currency (g-currency), without incurring in the cost of real payment mechanisms (feeds and taxes). It can be used to control participants behaviors, offering incentives for providers to share their resources and for client to give a reliable valuation of their jobs. It also serves as an overall regulation system, by restricting users with a limited purchase power leading to price contention during high peak demand periods

4. Implementing Markets

The Market Exchange Service is organized around the basic concept of Exchange Sessions, which are a gathering mechanism that allow participants to engage in Negotiations to interchange Offers about the session's underlying subject and potentially reach Agreements. An Offer represents the terms that the participant proposes for the negotiation. An Agreement is the confirmation that two parties agree on their mutual offers. Notice that session's subject are abstracts, and can refer to anything, from resources to be allocated to tasks to be executed.

The programming model for the Market Exchange Service is an asynchronous, event based model. It is based on two interfaces: the session handler and the negotiation handler. The combination of session and negotiation level handlers offers a versatile framework to implement diverse negotiation control policies. The participants obtain from an external source a reference to the session and join it specifying a negotiation handler that will handle this negotiations events.

In the next sections we show how this programming model can be used to implement two very different markets: a direct bargaining and an auction. These examples were chosen to show the flexibility of the model.

4.1 Implementing bargaining

This scenario (see figure 3a) is the simplest one: one of the participants starts a one-to-one negotiation with another participant creating a Session with an initial offer. The owner and the participant then exchange a series of offers until they both agree on the terms. The session handler can be used to implement an admission policy, to control, for example, that there is only one active negotiation at a time.

4.2 Implementing a double auction

In a double side auction (see figure 3b), the auctioneer doesn't make offers, but is just a facilitator that matches the offers sent by participants. These participants play complementary roles, such as sellers and buyers (as in the example below) or task requester and task executors. This market is therefore a many-to-many negotiation. At the closing time, the auctioneer looks at all the offers and matches them according to an optimization criterion. When it finds a pair of matching offers, it exchanges the offers of the participants. For example, it sends the sellers offer to the buyer and vice versa. This causes two agreements to be formed between seller and buyer (one seller-buyer and the other buyer-seller), but those agreements are identical to any purpose.

5. Related Work

There are, to our knowledge, few generic frameworks for developing grid markets. GridBus [3] is framework which provides a set of tools for the development, execution and management of grid applications using economic mechanisms. The main difference with GMM is that GridBus relies on the meta-scheduller model, on which users submits jobs (using a portal) to a economic scheduler which takes the allocation decisions based on the user supplied information, including budget. This model is more appealing for commodity markets. GMM, on the other hand, supports this model (even when it does not provides a meta-scheduler) but also allows a more open setups on

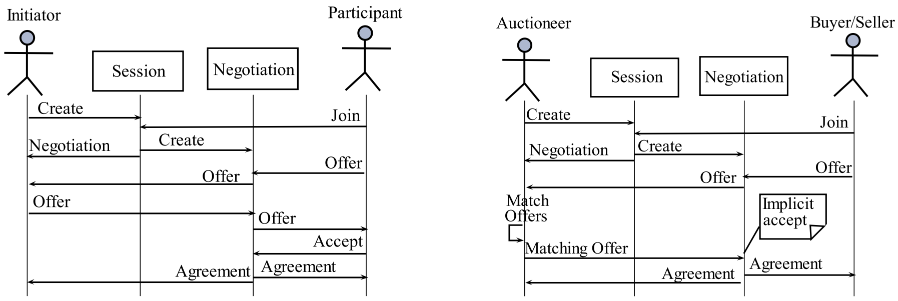

(a) Bargaining (b) Double Auction

Figure 3: Implementation Markets

which users submits jobs through agents that negotiate in one or more markets, better accommodating non commodity markets.

There has also been an important work in the development of market based agents in general [22] and for grid resource negotiation in particular [18–19]. However, these works generally focused on how to achieve desirable characteristics on individual agents (the strategies) or the negotiation protocols (the market mechanisms), but played little attention to the infrastructure actually needed to run a market for grid resources. GMM is intended to provide such infrastructure and facilitate the pass from prototype to actual utilization in real grid environments of such agents.

With respect of the individual components of the GMM, we briefly summarize some relevant work below.

Exchange. OCEAN [15] is a software infrastructure to automated commercial resource exchange over the Internet. It offers an optimized P2P search protocol to find a set of potential sellers and the automatic negotiation of those resources based on the rules defined in a XML format. The ability to define negotiation rules is a remarkable characteristic of OCEAN that allows the adaptation of the economic model to diverse applications. PeerMart [8] implements double auction over a P2P overlay network to distribute the auctioneer onto a clusters of peers, each being responsible for brokering several services. Its main limitation is the tightly integration of auctions model in the framework.

Directory. The Grid Market Directory [25] is a SOAP based service which relies on a centralized repository and offers a simple matchmaking capability suited for undifferentiated commodities. The GMM's Market Directory, on the other hand, is fully decentralized and is designed to support non-commodity, specialized resources. DyMRA [14] considers a decentralized, DHT based, Market Directory but it does not naturally supports complex multi-attribute and range queries needed in grid environments, even when those can be achieved extending its DHT's basic mechanisms [17].

Information. The Market Information Service integrates a publish-subscribe model with aggregation, filtering of routing messages, an important issue addressed in modern P2P systems. Systems like SDIMS [23] and Willow [16] uses DHTs specifically designed for this purpose. GMM's MIS, in the other hand, is built on top of a generic, unmodified DHT, which facilitate porting it to other DHT infrastructures. It also provides flexibility through a simple API that lets applications control propagation of queries and the aggregation of data.

Logging. [21] introduce publish/subscribe transactions for atomic production, delivery, and processing of asynchronous event notifications. As a general transaction management service, it incurs in a considerable overhead to maintain the atomicity of transactions. GMM's Logging, in the other hand, has been optimized to maintain a reliable log of activities under the control of a single Exchange Service instance. Distributed Log Service (DiLoS) [1] assign log files within a given scope (for type of service, geography, hierarchy, etc.) to a service instance, which stores it and disseminate to other instances in other scopes for redundancy. Its main drawback is the static organization of the services in scopes, while the GMM's Logging Service uses DHT's intrinsic characteristics to fulfill this requirements transparently and dynamically.

Currency. In PeerMint [8] and Karma [20] the bank role is distributed among a set of untrusted peers which conform the whole system, using protocols to overcome malicious peer's interferences and achieve consistency. Both systems suffers of performance and scalability issues due to the overhead generated by their reliable protocols. GridBank [2] is a centralized bank extensible to whatever kind of payment deployed in e-commerce. It implements a resource usage mechanism by means of RURs (Resource Usage Record) which helps in the task of accounting what resources have been consumed by which users. Its scalability and efficiency is questionable due to its centralized architecture.

6. Conclusions and Future Work

The Grid Market Middleware provides a number of core services c for the building of sophisticate market mechanisms in Grids. The GMMs architecture has designed to take advantage of the features of overlay networks to achieve scalability of several services and for supporting decentralized market models. In this paper we have shown how different market models such as bargaining and auctions can be implemented upon the GMM.

We have implemented prototypes of various GMM services (exchange, currency, information, directory) and our initial experiences show that they greatly simplifies the implementation of diverse market based resource allocation mechanisms. The prototypes were initially implemented separately, as part of the diverse projects we have been engaged. As we have learned more about both their particular traits and common ground, we have started to integrate them under the architecture shown in this paper. However, the prototypes still are not fully integrated. For example, they use diverse overlays, each suited to the service's particular needs. As has also been noted in [6], one of the main challenges is how to implement the services using only a generic overlay infrastructure, without harming its performance.

Acknowledgments

This work was supported in part by the European Union under contracts Grid4All EU IST-FP6-034567 and SORMA EU IST-FP6-034286, and the Spanish government under grant MEC P2PGRID .

References

[1] C. de Alfonso, M. Caballer, J. V. Carrin and V. Hernndez. Distributed General Logging Architecture for Grid Environments. *In High Performance Computing for Computational Science Conference: Selected and invited papers*, LNCS 4395, 2007.

[2] A. Barmouta and R. Buyya. Gridbank: a grid accounting services architecture for distributed systems sharing and integration. *In Parallel and Distributed Processing Symposium*, April 2003.

[3] R. Buyya and S. Venugopal. The Gridbus toolkit for service oriented grid and utility computing: an overview and status report. *1st IEEE International Workshop on Grid Economics and Business Models (CECOM)*, 2004

[4] Bai, X., Sivoncik, K., Turgut, D., Boloni, L. Grid Coordination with Marketmaker Agents. *International Journal of Computational Intelligence*, 3(2), 2006

[5] P. Chacin, F. Freitag, L. Navarro, I. Chao, O. Ardaiz. Integration of Decentralized Economic Models for Resource Self-Management in Application Layer Networks. *Second IFIP TC6 International Workshop on Autonomic Communications*, Athens, Greece, Octuber 3-5, 2005

[6] Y. Chawathe, S. Ramabhadran, S. Ratnasamy, A. LaMarca, S. Shenker, J. Hellerstein. A case study in building layered DHT applications. *Proceedings of the 2005 Conference on Applications, technologies, Architectures, and Protocols for Computer Communications*, 2005

[7] F. Dabek, B. Zhao, P. Druschel, J. Kubiatowicz, I. Stoica. Towards a Common API for Structured Peer-to-Peer Overlays. *2nd International Workshop on Peer-to-Peer Systems*, LNCS 2735, 2003

[8] D. Hausheer, B. Stiller. PeerMart: The Technology for a Distributed Auction-based Market for Peer-to-Peer Services. *40th IEEE International Conference on Communications (ICC 2005)*, Seoul, Korea, May 2005.

[9] R. Huebsch, J.M. Hellerstein, N. Lanham, B.T. Loo, S. Shenker and I. Stoica. Querying the internet with PIER. *Proceedings of the 29th international conference on Very large data bases*, 2003

[10] B. A. Huberman and T. Hogg. Distributed Computation as an Economic System. *Journal of Economic Perspectives*, 9(1):141-152, 1995.

[11] Hurwicz, L. The Design of Mechanisms for Resource Allocation. *The American Economic Review*, 63(2), 1973

[12] A. Iamnitchi, I. T. Foster On Fully Decentralized Resource Discovery in Grid Environments. *Proceedings of the Second International Workshop on Grid Computing*, p.51-62, 2001

[13] L. Joita, O.F.Rana, P. Chacin, I. Chao, F. Freitag, L. Navarro, O. Ardaiz. Application Deployment on Catallactic Grid Middleware. *IEEE Distributed Systems Online* 7(12), Dec. 2006

[14] D. Lzaro, X.Vilajosana and J. M. Marques. DyMRA: Dynamic Market Deployment for Decentralized Resource Allocation. *On the Move to Meaningful Internet Systems Workshops*, LNCS vol 4805, 2007

[15] P. Padala, C. Harrison, N. Pelfort, E. Jansen, M. P Frank and C. Chokkareddy. OCEAN: The Open Computation Exchange and Arbitration Network, A Market Approach to Meta computing. *In proceedings of the International Symposium on Parallel and Distributed Computing (ISPDC'03)*, Oct 2003

[16] R. van Renesse and A. Bozdog. Willow: Dht, aggregation, and publish/subscribe in one protocol. *In Proceedings Third International Workshop on Peer-To-Peer Systems*, LNCS 3279, 2004.

[17] J. Risson and T. Moors. Survey of research towards robust peer-to-peer networks: search methods. *Computer Networks: The International Journal of Computer and Telecommunications Networking*, Vol 50(17). 2006.

[18] Kwang Mong Sim. From market-driven agents to market-oriented grids (position paper). *ACM SIGecom Exchanges, Vol. 5(2)*. 2004.

[19] H. Tianfield. Towards agent based grid resource management. *Proceedings of Fifth IEEE International Symposium on Cluster Computing and the Grid (CCGrid'05)*. 2005

[20] V. Vishnumurthy, S. Chandrakumar, and E. Sirer.Karma: A secure economic framework. *In Proceedings of the Workshop on the Economics of Peer-to-Peer Systems*, June 2003.

[21] L. Vargas, L.I. W. Pesonen, E. Gudes, J.Bacon. Transactions in Content-Based Publish/Subscribe Middleware. *27th International Conference on Distributed Computing Systems Workshops*, 2007

[22] . M.P. Wellman.Market-Aware Agents for a Multiagent World. *Proceedings of the 8th European Workshop on Modelling Autonomous Agents in a Multi-Agent World*. LNCS Vol. 1237. Springer, 1997.

[23] P. Yalagandula and M. Dahlin. A scalable distributed information management system. *SIGCOMM Computer Communications Review*, 34(4):379-390, 2004.

[24] C. S. Yeo and R. Buyya. A Taxonomy of Market-based Resource Management Systems for Utility-driven Cluster Computing. *Software: Practice and Experience (SPE)*, 36(13):1381-1419, Wiley Press, New York, USA, Nov. 2006.

[25] J. Yu, S. Venugopal and R. Buyya. A Market-Oriented Grid Directory Service for Publication and Discovery of Grid Service Providers and their Services. *The Journal of Supercomputing*, 36(1), Springer, 2006.