Influence of Reputation in Revenue of Grid Service Providers

Mario Macías and Jordi Guitart
Barcelona Supercomputing Center
Jordi Girona 29
08034 Barcelona, Spain
{mario.macias,jordi.guitart}@bsc.es
http://www.bsc.es

Abstract

The trustworthiness of commodity providers is a key issue to motivate customers to participate in a market. Several Grid Market frameworks are being developed and they are not an exception in the need of demonstrate their trustworthiness to clients. This paper proposes a simple reputation mechanism for auctions within Grid Markets, which is easy to embed in old Grid solutions that do not provide reputation yet. Effectiveness of this mechanism is proved through simulation in different market scenarios, by showing the direct correlation between reputation and revenue of Grid service providers.

1 Introduction

Grid markets allow the participation of multiple agents from different organisations: any provider can join the market and sell its own resources or services to market clients. Before the sale is performed, both parts must negotiate and agree the terms of a contract: the Service-Level Agreement (SLA) [4]. But the fact is that this SLA can be violated by the service provider as a consequence of a system failure, an error in the negotiation (e.g. provider did not calculate well the expected system workload and it gets overloaded), or simply because the service provider is dishonest with the customers.

To avoid this, a mechanism for putting pressure on providers to fulfil the agreed SLAs is needed. The most usual mechanism is adding terms in the SLA for describing penalties, specifying the amount of money that the provider must pay to the customer if any of the SLA terms is violated [13]. This paper focuses on a complementary mechanism for enforcing SLA fulfilment: reputation and trust [7]. If a provider violates an agreed SLA its reputation decreases, and this information is accessible to market participants in future negotiations. We propose a simple reputation mechanism whose effectiveness is demonstrated by showing its influence in revenue of Grid market providers. A framework for simulating Grid resource auctions is created and the behaviour of the market is tested and compared between three scenarios: *demand excess*, *offer excess*, and *market equilibrium*. Taking the output data of the simulations, statistical correlation between reputation and revenue in resource providers is demonstrated.

The rest of the paper is organized as follows: Section 2 briefly describes the SORMA framework. Section 3 introduces our reputation mechanism. Section 4 describes the simulation scenario used to evaluate our mechanism, while Section 5 discusses the evaluation results. Section 6 describes the related work and Section 7 presents the conclusions of the paper.

2 Research Framework

This research paper is performed within the framework of Self-organising ICT Resource Management (SORMA) [3] project. It is an EU IST [2] funded project aimed at developing methods and tools for efficient market-based allocation of resources, using a self-organising resource management system and market-driven models, supported by extensions to existing Grid infrastructure.

Unlike traditional grid environments, tasks submitted to SORMA are matched with available resources according to the economic preferences of both resource providers and consumers and the current market conditions. This means that the classic Grid job scheduler, which relies on performance rules, is replaced by a set of self-organising, market-aware agents that negotiate SLAs to determine the best resource allocation to fulfil both performance and business goals.

In SORMA, all the components are classified into layers depending on their functionality. Figure 1 represents the

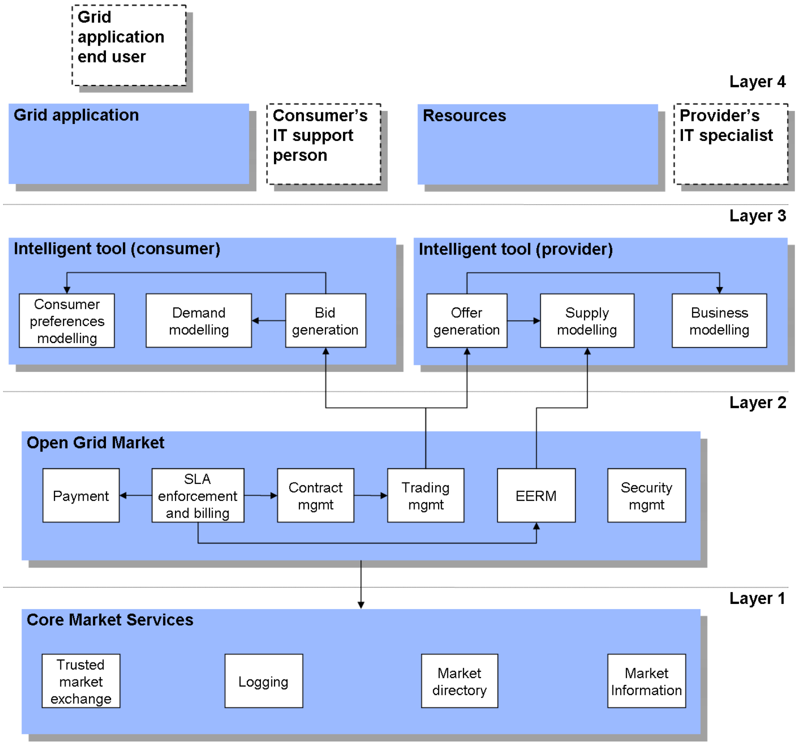

Figure 1. SORMA Architecture Logical View

current status of the logical architecture and describes what the system does in terms of its functional entities, their responsibilities and their dependencies. In Figure 1, boxes represent functional entities that in general (but do not have to) result in a corresponding component. Arrows depict dependencies, where an arrow from an entity A to an entity B means that entity A depends on entity B (in the sense that A receives input from B or uses B's services).

The remainder of this section describes the layers and associated components of the architecture to provide a high-level overview of the system design.

Layer 4: Grid Applications Layer 4 is the home of the Grid applications and Grid resources to be traded on the SORMA market. At the provider side a provider IT specialist makes use of the intelligent tools in layer 3 to model the provider's business strategies and the offered Grid resources. Grid resources in this context means a physical resource, a raw service and/or a complex service. On the consumer side it has to be distinguished between the Grid application's end user(s) and the consumer's IT support staff who will use the intelligent tools to model an application's resource requirements and the consumer's preferences.

Layer 3: Intelligent Tools The users (consumers and providers) are supported by intelligent tools for an easy access to the SORMA market. It makes easier the task of describe both the technical features and the business model of the traded resources and creates automatically both bids and offers to be sent to the SORMA market.

Layer 2: Open Grid Market Is the place where the offered resources/services are assigned to the Grid applications of the consumers, following certain market organizations. The **Trading Management** component is the access point for the consumers to the Open Grid Market where they can find the offered services and place their according bids. The **Contract Management** component transforms corresponding pairs of bids and offers to Service Level Agreements (SLA). The **SLA Enforcement and Billing** component is responsible for the surveillance and enforcement of the contracts it receives from the contract management. The **Economically Enhanced Resource Manager (EERM)** provides a standardized interface to typical Grid middleware (e.g. Globus Toolkit or Sun Grid Engine). The EERM can shield clients from resource platform specific issues and also enhance or complement the management functions provided by job scheduling and submission systems.

Layer 1: Core Market Services Standard Grid middleware does not provide all the infrastructure services necessary for an open marketplace. Layer 1 extends the standard Grid middleware by additional infrastructure services, as trusting, logging, market information systems, etc.

3 Adding a Reputation Mechanism to SORMA

In a Grid Market like SORMA, a malicious provider may offer some overestimated services to attract unwary customers; there is also the possibility that the service provider be not able to perform correct scheduling for their incoming tasks within its available resources. In both cases, the SLA agreed in negotiation time will be violated.

A penalty system for compensating customers could not be sufficient for some customers, because this mechanism is activated after the SLA is violated. Some customers may need to send some critical interactive tasks (e.g. real time video recognition for security surveillance) and prefer to prevent SLA violations rather than receiving an economic compensation derived from the bad operation of the purchased service. This paper proposes **reputation** as a mechanism to help market customers choosing a suitable service provider to fulfil their application requirements.

Other reputation mechanisms rely on the feedback given by customers. Taking the most of SORMA architecture, reputation feedback is sent by a neutral entity: *SLA Enforcement* [10] component, which continuously watches the correct fulfilment of the agreed SLAs by using monitored data of the resources 2. The purpose of the neutrality of this component is to prevent incorrect or malicious reputation feedback from the clients because the reputation is calculated in base to objective parameters.

The reputation component assigns a reputation indicator between 0 (worst) and 1 (best) to each provider. Each time a provider violates an SLA, its reputation is decreased proportionally to its category (this paper considers Gold, Silver and Bronze providers as described in Section 4) and the seriousness of the violation, which ranges from 0 (most catastrophic) to 1 (less serious).

In further auction processes, reputation is multiplied to the amount of money that a customer is willing to pay for the usage of the resource. This means that customers will only pay the 50% of the usual price for a resource whose provider has a reputation of 0.5, and the 100% for a resource whose provider has a reputation of 1. This system will automatically discard low-reputation providers, since they will practically never reach an agreement with customers, who are offering much less money for the service than providers ask for.

Each time an SLA is violated, the reputation of the provider is updated with the next formula:

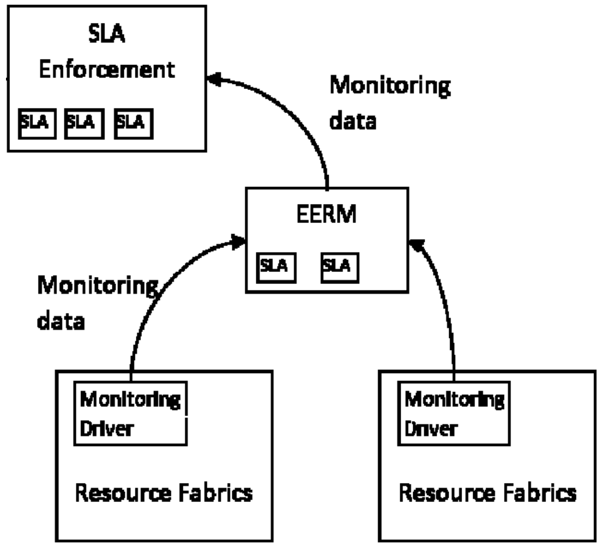

Figure 2. SLA Enforcement watches the fulfilment of the SLAs with the help of EERM, by using the monitoring data of the resources

$$R_t = R_{t-1} * (F_q + (1 - F_q) * S)$$

where R_t is the reputation at time step t, S the seriousness of the violation, and F_q is the factor for the reputation reduction in function to offered Quality of Service q. F_q can have values from 0 to 1: it will be greater for low-class providers and smaller for high-quality providers. This formulation assumes that high-quality providers are less allowed to have service failures and customers should be more tolerant with low-quality providers. In our experiments, $F_{gold} = 0.5$, $F_{silver} = 0.75$ and $F_{bronze} = 0.85$.

Service providers must also have the possibility to restore their reputation. Each time one of them fulfils correctly an SLA, the next formula is applied:

$$R_t = \begin{cases} R_{t-1} + \alpha & (R_{t-1} + \alpha < 1) \\ 1 & \text{(otherwise)} \end{cases}$$

where α is the reputation recovery rate. At greater values of α, reputation will be recovered more quickly.

It is also interesting to notice that the proposed mechanism for reputation is easily embeddable in Grid markets that do not provide this kind of functionality, because the simplistic conception of it: it only implements a service for updating the reputation and another to retrieve it, and the data stored is independent of the market data structure, since it is only needed an unique identifier for each provider and a float for its reputation. Actually, SORMA project does not provide any kind of reputation solution.

4 Experimental Environment

A simplified SORMA architecture has been simulated (see Figure 3). Both **Clients** and **Resource Providers** send bids and offers respectively to **Trading Management** component, which assigns clients to providers using the *English auction* mechanism: an initial price for the resource is given, and every client who wants to access it increases the given price which is willing to pay. At the end, the highest bid obtains the resource.

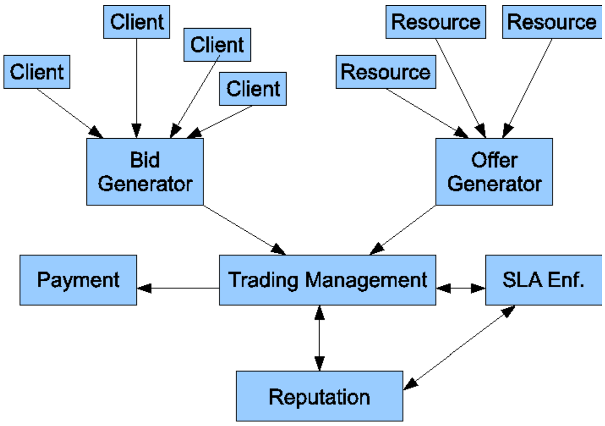

Figure 3. Architecture of the SORMA simulation

When the resource is obtained, the client pays to **Payment** component for it. In execution time, **SLA Enforcement** watches the resources to ensure that they are fulfiling the contract specifications in the SLAs. In the experiments in this paper, the SLA is simply defined by a given performance during 5 time slots.

There are three types of resources depending on the given quality of service: *Gold*, *Silver* and *Bronze*. Despite the prices vary in a given range, the average price of Gold providers is the higher, and Bronze providers are the cheapest ones. However, belonging to a low category does not imply that the price cannot be higher than the average of upper categories, because this is a free market.

Gold providers give the best guarantees for the correct SLA fulfilment (less than 2.5% of the SLAs are violated), Silver providers violate less than 5% of SLAs and Bronze providers violate less than 8.5% of SLAs. When a provider violates an SLA, the seriousness of the violation is calculated (in our experiments it is a random value), and the provider has to pay a penalty proportional to the seriousness and the paid amount for the service. This penalty is calculated according to the following formula:

$$Penalty = Price * Q * S$$

Table 1. Resource average characteristics

	Gold	Silver	Bronze
Probability to appear	0.2	0.3	0.5
Price range	[150,600]	[100,400]	[50,200]
Failure rate	2.5%	5%	8.5%

where $Price$ is the amount of money that the customer has paid to the provider and Q is 1, 2 or 3 if the provider is Bronze, Silver, or Gold class respectively.

All the data concerning resources (prices, reliability, and category) is generated randomly within some ranges specified for each service category (see Table 1). The simulation is performed through 1200 time slots. In each time slot, all the free customers and providers perform auctions to buy or sell resources for executing tasks. Each task needs 5 time slots, and after this time, the resource becomes free and ready to be auctioned. If a customer does not find a resource, increments a 10% his maximum amount of money to pay in the next time slot. If after 5 time slots the resource has not found any resource, it leaves the market.

In order to keep the specified customers/providers ratio, some new customers arrive to the market at each time slot. The behaviour of the market is simulated in three different scenarios: *excess of demand* (5 customers per each provider), *excess of offer* (3 providers per each customer), and *equilibrium* (approximately the same number of customers and providers).

5 Experiments Results

5.1 Demand Excess

In this scenario, there are too few providers and only the customers who are willing to pay the most can access the resources, which have a variable range of prices in function of their category. Figure 4(a) shows the prices of each of the 100 simulated resources, which are ordered by category (from Bronze to Gold) and by price. The local maximums delimit the different resource categories: clients from 0 to 35 are Bronze, from 36 to 75 are Silver, and clients from 76 to 100 are Gold. Prices and category are generated randomly, but the higher is the category, the more expensive tend to be the prices (see Table 1).

It can be seen that all the resources from the same category have similar revenue (Figure 4(b)), because there are much more customers than resources, and there is always a high probability to find a customer willing to pay a high amount of money for accessing the resource. This causes that providers with less reputation (Figure 4(c)) have also a high demand, and low reputation does not decrease the benefit. Furthermore, providers who violated an SLA have

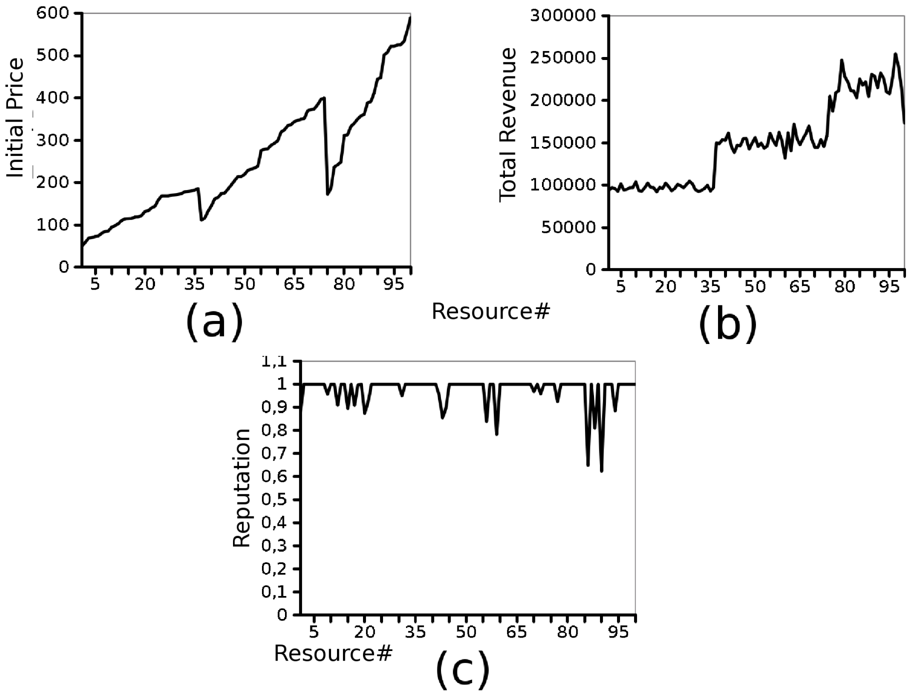

Figure 4. Results for the demand excess scenario for each resource: a) initial price for the auction, b) total revenue after the 1200 time slots simulation, and c) reputation of provider

more opportunities to restore their reputation. The plots in Figure 4 and the next ones are variable because the randomness of the parameters.

As a conclusion, this scenario is quite beneficial for all the providers, since they can increase enormously their prices and they continue getting clients. For customers, this is a very bad scenario, since a small part of them have to pay a lot of money to buy some resources, and the rest cannot have access to them. Taking into account the Offer and Demand Law [1], this scenario should lead to a drastic reduction of the number of customers which access SORMA (because they do not find what they want), and an increment of the number of providers (they want to have big profits).

5.2 Offer Excess

Offer excess scenario is simulated through 200 resources (the double that in the demand excess scenario). During the simulation, there are about three service providers for each customer. In consequence, only resources at affordable prices are frequently used and have some economic benefit (Figure 5(b)). Despite this, their economic benefit is quite low (about 50% of the benefit in the demand excess scenario). This is due to the auctions between customers, which are less aggressive than in previous scenario: there are less competing *rich* consumers, and they have more possibilities to choose.

Reputation of resources in offer excess scenario (Figure 5(c)) is generally lower than in demand excess one (Figure 4(c)). This is because resources do not have the same opportunities to restore their reputation in each scenario, since customers have more options to choose, and they will be reticent to buy low reputation resources.

This scenario is very good for customers; all of them can find what they want, and at low prices; but is very inefficient for providers. In this situation, the numbers of providers would be reduced (since they do not have economic benefit from their participation into SORMA), and the number of customers would be increased (for them, SORMA is a good place to find cheap resources).

5.3 Equilibrium Market

In a dynamic market, where offer and demand is continuously adapting to market scenario, previous two scenarios would lead to an equilibrium state, where the proportion between providers and resources allows to the first ones having an economic benefit, and to the others finding resources at reasonable prices. In this experiment, 100 providers (Figure 6(a)) are sold to an approximately equal number of customers in each time slot.

Figure 6(b) shows a sensibly higher benefit for both cus-

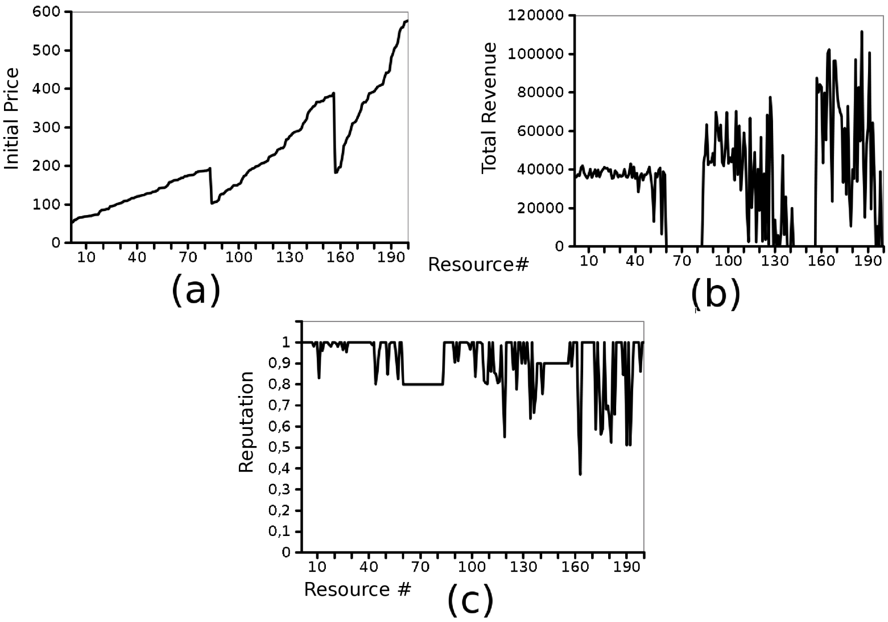

Figure 5. Results for the offer excess scenario for each resource: a) initial price for the auction, b) total revenue after the 1200 time slots simulation, and c) reputation of provider

tomers and providers: customers usually find what they need, and providers are economically profitable, except the most expensive ones.

Figures 6(b) and 6(c) have emphasized with circles three points where the graph falls down dramatically for Gold providers, and these points coincide for both reputation and revenue. This makes suspect that there is some relation between reputation and revenue. In the next section, all the simulation data will be analyzed to show if this relation really exists.

5.4 Influence of Reputation in Revenue

To calculate the true relation between reputation and revenue, this study used the *Pearson Correlation Coefficient (PCC)* (equation 1) in all the scenarios and resource categories separately.

$$PCC = \frac{\sum XY - \frac{\sum X \sum Y}{N}}{\sqrt{\left(\sum X^2 - \frac{(\sum X)^2}{N}\right)\left(\sum Y^2 - \frac{(\sum Y)^2}{N}\right)}} \quad (1)$$

PCC calculates the relation between X and Y data sets, each one with N elements. Its value is in the range from +1 (perfect linear relationship) to -1 (perfect negative linear relationship). Table 2 shows the PCCs between revenue and reputation for Gold, Silver and Bronze providers in each of the simulated market scenarios (demand excess, offer excess and equilibrium). Providers without revenue (nor usage) have been removed from the data sets, since this data would deform PCC calculations.

As it could be expected, Table 2 shows that the weakest correlations are given in **demand excess scenario** by the reasons explained in Section 5.1. In the other two scenarios, it seems logical that revenue in Gold and Silver providers has more relation with reputation than in Bronze providers, because customers accept that low-quality providers can have a greater failure rate.

A striking fact in Table 2 is that correlation between revenue and reputation in *equilibrium market* is greater than in the *offer excess* scenario. It seems against the common sense, due to the possibility to be more selective in offer excess scenario, and choose preferably the high reputation providers. The logical explanation for this is that since customers have more possibilities to choose in offer excess scenario, there are other parameters with more importance, e.g. the resource price.

6 Related Work

Grid Market reputation mechanisms have been studied recently by some authors. Abawajy and Goscinski [5] describe a Grid information service with reputation management facility, based on the concept of dynamic trust

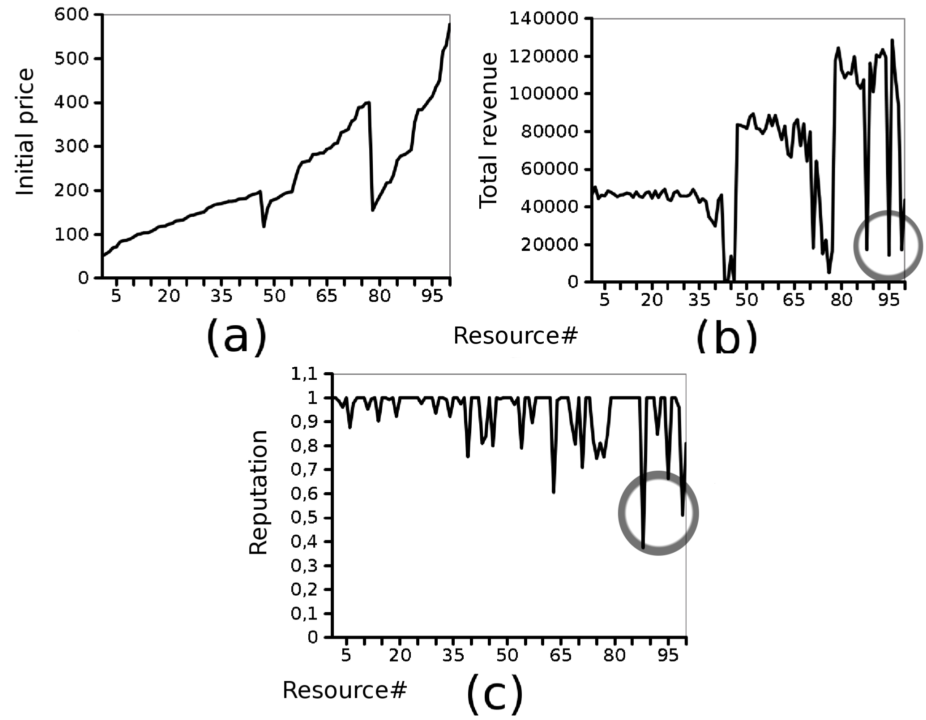

Figure 6. Results for the equilibrium scenario for each resource: a) initial price for the auction, b) total revenue after the 1200 time slots simulation, and c) reputation of provider. The circled spikes make suspect about the correlation between revenue and reputation

Table 2. Correlation between reputation and revenue

	Gold	Silver	Bronze
Demand Excess	0.001	0.167	0.143
Offer Excess	0.418	0.420	0.357
Equilibrium	0.883	0.633	0.550

and reputation adaptation based on community experiences. *GridEigenTrust framework* [6] proposes an algorithm for evaluating Grid reputation by combining eigenvectors to compute reputation, and integration of global trust. The main differences with this paper are the complexity of their theoretical proposals, and that these papers are focused mainly on experimental simulations to demonstrate how a simple-but-effective policy can also work.

Chen et al. [8] experimentally compare low-information, high-information and self-reporting reputation mechanisms. The results indicate that players strategically reacted to the reputation mechanisms, with higher information mechanisms increasing market efficiency. Sonnek et al. [11] perform wide experiments to compare some existent reputation mechanisms in quantitative terms. Weng et al. [12] consider the possibility of malicious referrers who report inaccurate testimonies and proposes a system to mitigate them. This proposal is not needed in a system like the one described in this paper, since reputation feedbacks are given by a neutral entity in the SORMA system. Lu et al. [9] propose mechanisms where reputation value of resource provider reflects the reliability of its resources: their mechanisms decide the trades between resource providers and consumers in terms of both reputation and price, and solve the information asymmetry problems.

7 Conclusions and Future Work

This paper has introduced a proposal for a very simple reputation mechanism that can be easily embedded in Grid scenarios which do not provide reputation: in the SORMA example scenario, only is needed to add a service call from *SLA Enforcement* component to update the reputation data, and another one from *Trading Management* to retrieve the reputation in negotiation time. The reputation mechanism will help legacy Grid clients to acquire better resource providers in a transparent way, without need of modification. This affirmation is demonstrated through a simulation experiment and the analysis of its result data.

Several economic scenarios have been considered: offer and demand excess, and equilibrium market. It is demonstrated that the equilibrium market is the optimal scenario, because almost all the entities obtain benefit from their participation in the market. The simulations are pretty simple, but enough to demonstrate the importance of reputation when negotiating for resources and how the reputation makes providers to adjust their category and prices to their true reliability.

Future work will include simulating a market where providers adapt their number and prices and customers adapt their bids in function of the current demand/supply proportion. After that, it should be checked how the whole market evolves from disequilibrium to equilibrium state. In addition, the reputation mechanism should be also tested in an real SORMA implementation, and verify the behaviour or real market participants. This paper has only focused in provider-side reputation. Future work will also include the study of reputation not only in Grid providers, but also in Grid Market customers.

References

[1] Economics basics: Demand and supply (http://www.investopedia.com/university/economics/economics3.asp).

[2] Information Society Technologies Programme (http://www.cordis.lu/ist).

[3] Self-organizing ICT Resource Management (SORMA) (http://www.sorma-project.eu).

[4] Web Services Agreement specification (http://www.ogf.org/documents/GFD.107.pdf).

[5] J. H. Abawajy and A. M. Goscinski. A Reputation-Based Grid Information Service. In *International Conference on Computational Science (ICCS'06)*, volume 3994 of *Lecture Notes in Computer Science*, pages 1015–1022. Springer, 2006.

[6] B. Alunkal, I. Veljkovic, G. von Laszewski, and K. Amin. Reputation-Based Grid Resource Selection. In *Proceedings of the Workshop on Adaptive Grid Middleware (AGridM 2003)*, New Orleans, LA, USA, September 28, 2003.

[7] A. Josang, R. Ismail, and C. Boyd. A Survey of Trust and Reputation Systems for Online Service Provision. In Elsevier Science Publishers B. V., editor, *Decision Support Systems*, volume 43, pages 618–644, Amsterdam, The Netherlands, 2007.

[8] K. Chen and T. Hogg and N. Wozny. Experimental Study of Market Reputation Mechanisms. In *Proceedings of the 5th ACM conference on Electronic Commerce (EC'04)*, pages 234–235, New York, NY, USA, 2004. ACM.

[9] W. Lu, S. Yang, L. Guo, and R. Zhang. Reputation-aware Transaction Mechanisms in Grid Resource Market. In *Proceedings of the 6th International Conference on Grid and Cooperative Computing (GCC'07)*, pages 154–159, Xinjiang, China, August 16–18, 2007. IEEE Computer Society.

[10] M. Macias, O. Rana, G. Smith, J. Guitart, and J. Torres. Maximizing Revenue in Grid Markets using an Economically Enhanced Resource Manager. *To Appear in Concurrency and Computation: Practice and Experience*, 2008.

[11] J. Sonnek and J. Weissman. A Quantitative Comparison of Reputation Systems in the Grid. In *Proceedings of the 6th IEEE/ACM International Workshop on Grid Computing*, pages 242–249, Seattle, Washington, USA, November 13–14, 2005.

[12] J. Weng, C. Miao, and A. Goh. A Robust Reputation System for the Grid. In *Proceedings of the 6th IEEE International Symposium on Cluster Computing and the Grid (CCGrid'06)*, pages 548–551, Singapore, May 16–19, 2006.

[13] C. S. Yeo and R. Buyya. Service Level Agreement based Allocation of Cluster Resources: Handling Penalty to Enhance Utility. In *Proceedings of the 2005 IEEE International Conference on Cluster Computing*, pages 1–10, Boston, Massachusetts, USA, September 26–30, 2005.