

2008

Reputation-Based Pricing for Grid Computing in E-Science

Arun Anandasivam

Information Management and Systems, University of Karlsruhe, Germany, anandasivam@iism.uni-karlsruhe.de

Dirk Neumann

Albert-Ludwigs-Universität Freiburg, dirk.neumann@vwl.uni-freiburg.de

Follow this and additional works at: <http://aisel.aisnet.org/ecis2008>

Recommended Citation

Anandasivam, Arun and Neumann, Dirk, "Reputation-Based Pricing for Grid Computing in E-Science" (2008). *ECIS 2008 Proceedings*. 103.

<http://aisel.aisnet.org/ecis2008/103>

This material is brought to you by the European Conference on Information Systems (ECIS) at AIS Electronic Library (AISeL). It has been accepted for inclusion in ECIS 2008 Proceedings by an authorized administrator of AIS Electronic Library (AISeL). For more information, please contact elibrary@aisnet.org.

REPUTATION-BASED PRICING FOR GRID COMPUTING IN E-SCIENCE

Anandasivam, Arun, University of Karlsruhe, Englerstr. 14, 76131 Karlsruhe, Germany,
anandasivam@iism.uni-karlsruhe.de

Neumann, Dirk, University of Freiburg, Platz der Alten Synagoge, 79085 Freiburg, Germany,
dirk.neumann@vwl.uni-freiburg.de

Abstract

One of the fundamental aspects for an efficient Grid usage is the optimization of resource allocation among the participants. However, this has not yet materialized. Each user is a self-interested participant trying to maximize his utility whereas the utility is not only determined by the fastest completion time, but on the prices as well. Future revenues are influenced by users' reputation. Reputation mechanisms help to build trust between loosely coupled and geographically distributed participants. Providers need an incentive to reduce selfish cancellation of jobs and privilege own jobs.

In this paper we present a reputation-based pricing mechanism for a simple, but fair pricing of resources. In e-Science researchers do not appreciate idiosyncratic pricing strategies and policies. Their interest lies in doing research in an efficient manner. Consequently, in our mechanism the price is tightly coupled to the reputation of a site to guarantee fairness of pricing and facilitate price determination. Furthermore, the price is not the only parameter as completion time plays an important role, when deadlines have to be met. We provide a flexible utility and decision model for every participant and analyze the outcome of our reputation-based pricing system via simulation.

Keywords: Grid computing, reputation, pricing, incentives.

1 INTRODUCTION

Grid computing is a promising paradigm for sharing IT-resources in large-scale geographically distributed systems through collaboration (Foster, Kesselman 2001). It enables to use software and hardware infrastructures from other institutes for an effective sharing of heterogeneous computing resources, data or even high-performance and complex services (Joseph et al. 2004). The collaboration in the particle physics Grid Community has been facilitated by virtual organizations (VO) like Atlas or LHCb (Berlich et al. 2005). Scientists are associated with virtual organizations, which is a loosely-coupled team of people working in the same or closely related projects. They work on the same infrastructure with an interoperable application environment. One of the fundamental aspects of an efficient Grid usage is the optimization of resource allocation among the participants. However, this has not yet materialized. Users tend to send one job several times to the Grid to assure that evaluable results will be returned. The redundant job submission blocks resources, which potentially could be allocated to other, more important jobs. Moreover, site administrators hesitate to always provide all the resources to the VO, in case an internal job is waiting. Then, the instant allocation of the internal job is preferred. This behaviour is comparable to free-riding behaviour in P2P-networks (Adar, Huberman 2000). By introducing prices for jobs the behaviour can be redirected avoiding free-riding. Apparently, we have a conflict in goals. Users are interested in satisfying their resource demand as quickly as possible while the overall goal is to provide a fair and efficient resource sharing. On the supply side providers try to get as much jobs as possible for the highest price they can achieve. Each user is a self-interested participant trying to maximize his utility whereas the utility is not only determined by the fastest completion time, but on the prices as well.

Self-interested agents can lead to inefficient market outcome. Enforcing authorities are not always able to detect and punish misbehaviour. Reputation mechanisms help to build trust between loosely coupled and geographically distributed participants (Resnick et al. 2000, Dellarocas 2003) to avoid a market of lemons (Akerlof 1970). Future revenues are influenced by users' reputation based on the behaviour of all participants. In Grid networks, the incorrect results returned by a finished job do not reveal information, whose fault it was. On the one hand, the provider could have aborted the job. On the other hand, the consumer could have made mistakes in programming the job. Thus, reputation mechanisms have to provide tailored metrics to rate provider and consumer. Providers need an incentive to reduce selfish cancellation of jobs, while consumers have to thoroughly analyze their jobs, before they submit them.

In this paper we present a reputation-based pricing mechanism for a simple, but fair pricing of resources. In e-Science researchers do not appreciate idiosyncratic pricing strategies and policies. Their interest lies in doing research in an efficient manner. Consequently, in our mechanism the price is tightly coupled to the reputation of a site to guarantee fairness of pricing and facilitate price determination. Furthermore, the price is not the only parameter as completion time plays an important role, when deadlines have to be met. In (Dellarocas 2003) one of the research questions for reputation mechanisms is how they affect the behaviour of participants in a community. We provide a flexible utility and decision model for every participant and analyze the outcome of our reputation-based pricing system via simulation. The goal is to enhance collaboration and trust in the Grid community (Ranganathan et al. 2004).

2 SCENARIO

The community in the particle physics Grid is based on trust between the institutes. Institutes are comprised of several researchers. Researchers typically analyze data and need thus huge amounts of computation power to run simulations and calculations. Therefore, they share resources with other institutes and have the option to schedule their jobs either on the local machine or to send it to an

external site. It depends on the queue length estimation when the job will be finished. Researchers submit their jobs to external sites and anticipate their job will be finished within the expected timeline. Obviously, the jobs of internal users are always more important than jobs of external users due to selfish manner. Thus, sites tend to cancel current running jobs, if there is an important internal job to process. Nowadays, there are no incentives why running jobs of external users should not be cancelled as the institutes do not gain from this job. Users do not have to reciprocate (e.g. payment) to use the Grid resources. Consequently, job requesters will inefficiently consume the offered resource by sending jobs redundantly to the Grid. Figure 1 illustrates the advantage for the consumer sending his job J10 redundantly to several machines. On two of the three foreign machines the job is not executed properly. These jobs are of no value for the consumer. He expects that the other two jobs will deliver valuable results. Without compensation payment he always has the incentive for redundant job submission to avoid the risk of job loss. However, J10 delays other jobs (e.g. J11), which could be more important.




Figure 1. Four jobs are sent to the Grid network. Only two jobs will deliver valuable results.

The goal of the system is to achieve a fair allocation of resources and enforce obedient behaviour of the participating agents. Generating a reliable Grid platform in the e-Science community for distributed resource sharing and collaboration requires mechanisms to solicit and predict resource contributions of individual users (Buragohain et al. 2003). Buragohain suggested an incentive mechanism for P2P file sharing, where users with higher service provision have a higher probability to be accepted by others for downloading files. Every user has costs for offering files and he can gain from offered files by other users. The approach in this paper is to differentiate between the services a user provides. A user with a high contribution is more likely to be accepted than a user with a low contribution. This analysis framework can be used to identify the benefit of participating in the network after the transaction. Buragohain's game theoretic analysis framework is not applicable for Grid Computing, since it does not consider that CPU sharing is not reusable for one timeslot. In P2P networks a file can be downloaded by several peers at the same time (parallel resource usage). Subsequently, the assignment of resources does not only depend on his provision level, but also on the available resources at the requested timeslot. Moreover, we do not reject requests based on probability. Jobs, which have been submitted to a site, must be accepted.

The introduction of payment can solve the problem of inefficient resource usage. Every user has to pay a certain amount of money to receive resources. The amount of money has to be limited, since real money in scientific Grid networks is undesirable. Instead virtual currencies or credits can be implemented like Karma or Nuglets (Vishnumurthy et al. 2003; Buttyan and Hubaux 2001). This induces further problems as virtual credits are used to price resources. Users or site administrators

have to decide, how to determine the price of a resource depending on capacity, demand and availability of resources. This entire process requires time, which distracts from research. Another option is to have a fixed price for resources, e.g. 1 credit per CPU/minute. The prices need not be determined dynamically and an incentive is provided to stop overconsumption of resources. However, fixed price fail to set incentives to behave compliantly as site administrator can cancel job at any time. They will accept a short-term loss in payment made by the current running job. In this case the own job, which has to be finished before the deadline and has a higher valuation than the fixed payment, will replace the running job. In this paper a fixed-price scheme is extended by a reputation mechanism to enhance incentives for collaboration in the Grid network. Prices usually reflect the supply and demand. The proposed pricing is advantageous as it reflects the service level. An automatic adaptation of the price according to users' behaviour allows setting the right incentives for collaborative work resulting in an improved exchange of resources. Furthermore, a decision framework for cancelling a job is presented to depict the scenario in a scientific Grid network.

3 REPUTATIONS-BASED SCHEDULING AND PRICING

The main idea has been derived from (Jurca and Faltings 2005), where the authors propose a reputation-based pricing for services in P2P networks based on the provided quality of service. Deadlines and completion time are not considered in their utility function and thus not suited for Grid. Our utility function comprises these parameters. We adapted the online scheduling mechanisms from (Porter 2004) and (Heydenreich et al. 2006). Porter's utility function is not based on the length of the job, but on the valuation for the job. The mechanism contemplates when and how a job has to be submitted and users can report true or false values for job length or job valuation. We assume that a long job has a bigger impact and a higher risk to be cancelled than a short job and the values of a job are reported truthfully. Typically, long jobs comprise high effort in programming and thus the impact of the results is high. Due to their long running time the probability increases that the job will be cancelled. Moreover, the formal analysis of Porter is based on mechanism design for a single machine, whereas in our case we consider m machines in a simulation. Heydenreich et al. propose a mechanism called Decentralized LocalGreedy Algorithm (DLGM). There is no central planner to allocate jobs to the different nodes. Instead, jobs ask for the completion time and payment on each machine and decide on which machine they want to be scheduled. Jobs can report a value for their job and get a higher priority and be executed earlier than previously allocated jobs. Deadlines of the jobs are not taken into account. Furthermore, the option that a user (or the machine owner) can cancel the current running job was not analyzed. It makes a new option for decision available.

In e-Science Grid users are researchers who are sending jobs to the Grid consume subsequently from other Grid research institutes. In our model we will consider sites as consumer and provider. Other papers (i.e. Kwok et al. 2005) distinguish between provider and consumer as two different persons/institutes, whereas in our case the decision model is depending on both roles (i.e. Buragohain et al. 2003). Thus, sending and receiving jobs has an impact on the decision for a site in both roles. To distinguish between the provider and consumer role we will name the consumer as jobs and provider as machines. Jobs and machines can belong to the same user. For simplicity, this model implies without loss of generality that one site has only one user, where every user has a reputation. The calculation of the reputation value is not fixed to a certain scheme. Promising examples for reputation mechanisms are (Alunkal et al. 2004) for Grid networks as well as (Xiong and Liu 2004) and (Kamvar et al. 2003) known from P2P networks. Similar to the mentioned examples we assume that users in Grid report the feedback truthfully and act rationally.

3.1 Parameter

Preferences of a user are expressed by the utility function. It is crucial for defining the relation between the loss and the benefit of a job on a machine at a certain time. Besides obvious and essential characteristics of a utility function further requirements have to be met:

1. A job which is completed after the deadline has a value of zero. It does not have any value for the user, if the job is finished too late.
2. The risk of a job cancellation increases, if the provider has a lower reputation.
3. The cancellation of a job must have a direct impact on the future income.

We consider a scenario with a set of Gridagents $A = \{1 \dots a\}$, who participate in the network by providing resources and submitting jobs. Resources are homogeneous. Every agent a can send one or more jobs j_a to the Grid in one timeslot $T = \{1 \dots t\}$. The jobs are defined by a processing time $p_j > 0$ (runtime of job), and a deadline $d_j > 0$ (when the job should be finished). The incoming jobs are always able to meet the deadline, if they are instantly started. Every job requests the machines in the network ($M = \{1 \dots m\}$ with $a=m$) for their queue time $q_{jm}(t)$. This approach is comparable to the DLGM setting (Heidenreich et al. 2006). The completion time $C_j(m)$ is defined as $C_j(m) = p_{j_x m} + q_{jm}(t)$, where $p_{j_x m}$ denotes the remaining time for the current running job j_x from agent x on machine m . Every machine has a reputation value $r_m \in [0, 1]$.

Porter proposed a utility function based on a hard deadline. Every user has an expectation about the latest finishing date of a job. The job is worthless, if it is finished after the deadline and it is thus cancelled (Porter 2004). We do not abort jobs, if they are waiting in the queue and probably will not match the deadline. We assume that jobs can still match the deadline, if preceding jobs are cancelled and replaced by shorter jobs. Then, the completion time will be reduced. The option of cancellation is only considered, if a new job cannot be finished before the deadline. The valuation for the job's laxity is determined by the parameter V_j . Thus, the utility is defined as

$$U_{jm}(t) = \mu\left(D_j \geq \frac{1}{\sqrt{r_m(t)}} C_{jm}(t) + t\right) * V_j - \pi_{jm}(t).$$

We use the notation according to Porter, where $\mu(\cdot)$ is an indicator function, which returns 1, if the argument is true, and zero otherwise (requirement 1). The deadline should be bigger than the sum of the completion time and the current timeslot. Furthermore, every machine is evaluated by the risk of job cancellation. We introduce a risk factor $\frac{1}{\sqrt{r_m(t)}}$ to fulfil requirement 2. If the machine has a high reputation, the job will more likely be finished before deadline. $\pi_{jm}(t)$ is the total payment job j pays to the machine m . We propose a reputation-based pricing, which enables a direct price determination based on the reputation of the provider. A provider with a higher reputation will consequently receive a higher income per timeslot. If the reputation decreases, the price will decrease, too. Let the price per timeslot be $v_{jm}(t): [r^{min}, r^{max}] \rightarrow \mathbb{R}$ (requirement 3). In the remainder of this paper it is simplified to $v_{jm}(t) = r_m(t) \in [0, 1]$. The total payment of a user to the machine m is $\pi_{jm}(t) = p_j * v_m(t)$. A consumer can rate the provider depending on, if the job was cancelled, finished too late or successfully returned. The provider is not rated negative, if the deadline was matched, although the promised finishing date was delayed. At the beginning of the allocation the price is set accordingly to the reputation. The utility of a job can be positive or negative, since the payment can be higher than the valuation of a job. On contrary, the definition of DLGM only allows negative utility.

3.2 Sellers' and buyers' action space

When a job is created, the agent has the action space S for the job with $S = \{run\ job\ on\ own\ machine, run\ job\ on\ foreign\ machine, cancel\ running\ job\ of\ other\ agent\}$. Usually, the third option is the best, if no reputation and prices are considered, since the agent is not punished for his misbehaviour. We only consider the option to cancel the running job. The replacement of a job in the queue is not taken into account in this setting.

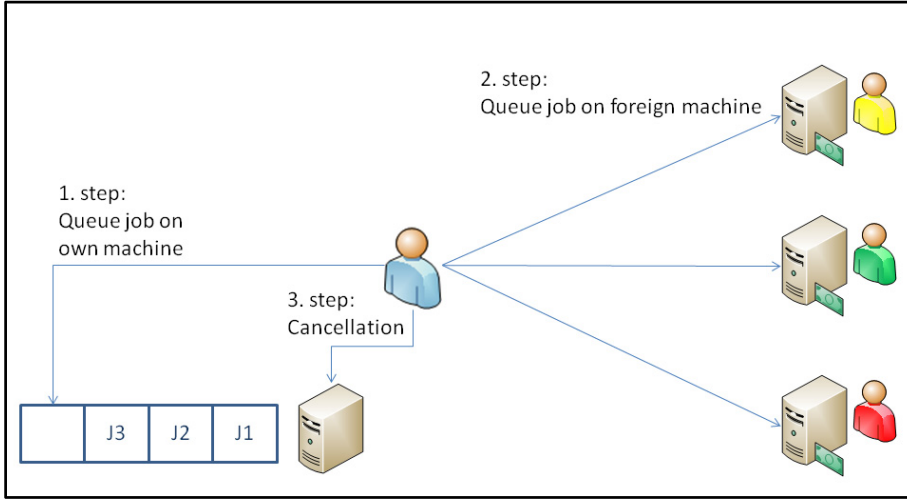


Figure 2. Decision process of the user

The decision process is as follows. At first, the agent calculates the utility, if his job is scheduled on his machine. Although the agent does not have to pay himself ($\pi_{jm}(t) = 0$), the agent has opportunity costs, since no other foreign jobs can run on this machine and the income is missing for this period of time. In the next step he analyses whether a better utility can be gained by running the job on a foreign machine. The completion time has to be lower, because the price decreases the utility ($\pi_{jm}(t) \geq 0$) compared to scheduling on his machine. As third option users have the ability to cancel running jobs on their own machine while processing their job instantly (figure 2). This is a big advantage, when queues are very long due to high demand and the completion time of a job extends the deadline on all machines, e.g. it does not get finished within time. We assume that a user would never cancel his own job on his machine. From the consumer perspective it is only attractive for the consumer to schedule his job on another machine, once the current running job is not from another provider. Otherwise, it is reasonable to cancel the running job, because he will obtain the lowest completion time (without reputation and payment). On the one hand, payments decrease the benefit of scheduling the job on another machine, since for the own site the user is not required to pay. Consequently, it is less attractive to send jobs to others. On the other hand, the cancellation of jobs results in a negative outcome for the machine owner, because he will not receive any payments and he will be punished by a lower reputation. Queues on other machines may comprise fewer jobs and thus attract users to schedule their jobs on other machines.

By missing the deadline the results of a job create no value. Henceforth, the provider faces two effects: *payment* and *reputation loss*. Payment loss arises by replacing the current running job by a new job and delaying other jobs in the queue. Delay can result in missing the deadline. Finished jobs beyond deadline are not being paid. The loss is calculated by summing up the expected payment for all delayed jobs including the cancelled job: $l_m^{delayLoss}(t) = \sum_{k=1}^Q \pi_{km}(t) * \mu(D_j < \hat{C}_{jm}(t) + t)$.

Delayed jobs and the cancelled job have the opportunity to rate the provider. Apparently, they will submit a negative rating and the provider will face a reputation loss. The number of negative ratings is $R_m^{neg}(t) = \sum_{k=1}^Q \mu(D_j < \hat{C}_{jm}(t) + t)$. Depending on the reputation mechanism the negative ratings will lower the reputation of agent a possessing machine m . Then, the agent has to collect $R_m^{pos}(t)$ positive ratings to regain his former reputation. $R_m^{pos}(t)$ is the number of required jobs, which rate the machine positively. $R_m^{neg}(t)$ and $R_m^{pos}(t)$ need not be equal, i.e. in asymmetric reputation mechanisms, where it is more difficult to receive a good reputation than a bad reputation. Next, it has to be analyzed how long it will take to receive the required jobs. As there are jobs already in the queue meeting the deadline the number of required jobs for obtaining the old reputation value is $\hat{J}_m(t) = R_m^{pos}(t) - \sum_{k=1}^Q \mu(D_j \geq \hat{C}_{jm}(t) + t)$. Since the machine will have a lower income due to the reputation loss and thus a lower price, the compensation is based on the current queued jobs and the incoming rate $\hat{\lambda}$ of jobs on the machine m in the future timeslot. Let the prospective jobs arrive according to a predefined distribution and have a processing time equalling the mean \bar{p} . The incoming rate of jobs is derived from the history. We assume that jobs will arrive according to former income rate. We use the exponential smoothing method to predict the jobs arriving in the future. Subsequently, the number of expected jobs arriving in timeslot $t+1$ is $J_m(t+1) = \alpha * y_t + (1 - \alpha) * J_m(t)$. The required number of timeslots to restore the old reputation encompasses the duration of all jobs in the queue and the expected runtime of future jobs: $t_m^{required} = \bar{p} * \hat{J}_m(t) + \sum_{k=1}^Q p_k$. The jobs, which arrive before the reputation is restored, create a loss for the provider, since they have to pay a lower price.

The reputation value including $R_m^{neg}(t)$ ratings $r_m^{neg}(t)$ and the current value is $r_m^{pos}(t)$. We average the price the incoming jobs have to pay until $t_m^{required}$ by $v_{jm}(t) = \frac{r_m^{pos}(t) - r_m^{neg}(t)}{2}$. Knowing the number of expected jobs and the expected payment, the expected loss can be determined, if the agent cancels the running job: $l_{jm}^{repLoss}(t) = t_m^{required} * v_{jm}(t)$.

The total utility for cancelation is $U_{jm}^{cancel}(t) = \mu(D_j \geq p_j(t) + t) * V_j - l_m^{delayLoss}(t) - l_{jm}^{repLoss}(t)$.

Users can decide whether the utility gained by the cancellation exceeds the utility of a regularly scheduled job. In our simulation we only consider this option, when the job will fail the deadline due to large queues.

3.3 Reputation mechanism

The selection of a reputation mechanism is crucial for the mechanism. Different reputation mechanism will influence the user's decision function. The authors of (Marti and Garcia-Molina 2006) provide a taxonomy for identifying properties of reputation mechanisms to choose the right mechanism for the right setting. Although the taxonomy was mainly developed for P2P networks, it is applicable for Grid networks as well. They distinguish between *information gathering*, *scoring and ranking* and *response*. The first category comprises the precondition to create identities, the information sources and the level of information detail. Scoring and ranking defines the input data and the output data of a reputation mechanism. Response is the action a user can take or the action space a user is limited to, i.e. users who have a low contribution level also have a low download capacity in P2P networks.

In the particle physics Grid we deal with registered identities, which have to be certified by a certification authority. This process is only for authorising the user to participate in the network. The site administrator is unaware of whom the current running job belongs to. Jobs are mapped by a proxy identifier, which is managed by the resource broker. Thus, the circumvention of the reputation mechanism like whitewashing and Sybil attacks (Cheng and Friedman 2005) are impossible in Grid networks, since every certificate application is thoroughly analyzed by several authorities. However, the consumer is unknown to the provider due to anonymity. Information sources can either be local

reputation or global reputation. In this paper we restrict our attention to global reputation, since we have centralized authority to gather and disseminate this information and we assume to have agents rating honestly (Marti and Garcia-Molina 2006; Sonnek and Weissman 2005). The level of information detail will be reduced to aggregated information about former behaviour to keep it as simple as possible. The reputation will be represented by a single value. Strategic behaviour based on the former actions is not taken into account. The input data for computing the reputation value weigh current ratings higher than old ones. (Jurca and Faltings 2005) preferred an even simpler mechanism by deriving the reputation value as follows: $\frac{\text{number of good ratings}}{\text{number of total ratings}}$. We apply this simple reputation mechanism as in (Jurca and Faltings 2005; Sonnek and Weissman 2005) and adapt it with a straightforward decay function (Azzedin and Maheswaran 2002; Alunkal et al 2004). The rationale behind this is that older reputations are less important (Zacharia et al 1999). The quasi-decay function takes only the last g ratings and weights them equally. This approach induces that the order of the submitted rating is essential. Thus, ratings, which came in first, will be excluded first after g ratings.

4 SIMULATION AND IMPLEMENTATION

In this section, we present the first results from our simulation. Our results show that reputation-based pricing gains higher exchange of resources and less cancellation of jobs than fixed price schemes.

4.1 Setting

Currently the particle physics researchers do not have to pay for the usage of Grid resources. We therefore set up a scenario without payment and a scenario including reputation based pricing. The agents are fully trusted at the beginning. They start with a reputation value $r_m(0) = 1$. The simulation is round-based and comprised 1000 rounds in each 20 runs. Every round jobs were created according to a Poisson distribution with $\lambda=0.25$ for each agent. 50 agents were providing their resources and interacting with each other. Jobs had a completion time and stayed on a machine until the job was done or cancelled. The job processing time was derived from a truncated normal distribution with mean = 3 and deviation = 2. Only positive durations were allowed. Every job had a deadline which was uniformly distributed between 10 and 40 timeslots. The valuation of a job for the job owner was uniformly distributed between 1 and 10. A job was of no value, if it was finished after deadline. To benchmark the reputation-based pricing we used a fixed price scheme with $p = 0$, $p = 1.0$ to demonstrate the effect of prices on the scheduling outcome. A fixed price with $p = 0$ represents the current Grid where no payment is necessary. The highest payment in the reputation-based scenario is $r^{max} = 1$.

4.2 Results

The goal of the proposed mechanism is to show the effect of reputation-based pricing and the benefit for the particle physics Grid. One metric is to view the amount of cancelled jobs. The less jobs are cancelled, the higher the trust in the network. Table 1 illustrates the results. The reputation-based pricing enforces site administrator not to cancel jobs, since their behaviour is documented by the reputation. For the fixed price scenarios there is negligible discrepancy between each setting. Overall, there were about 900 attempts to cancel a job for fixed prices and about 750 for the reputation-based pricing. It encompasses the attempts to cancel the own jobs as well, which was excluded in our setting. Consequently, one third of the jobs were cancelled in the fixed price setting and only 5% in the reputation-based pricing. This is a significant reduction and a strong enhancement of trust in the Grid network.

Another metric considers the number of jobs, which have not met their deadline. Grid users rely on the jobs they sent to other sites. When jobs regularly do not meet the deadline due to cancellation users will distrust other sites, because participation in the Grid will not be individual rational (Schnizler et al 2006). Accordingly, jobs waiting in the queue may not meet their scheduled deadlines as in the meantime the machine has cancelled the running job in favour of a longer own job. As shown in table 1 there is a discrepancy between the reputation-based pricing and fixed pricing. The jobs of a user running on his machine were always lower than foreign jobs for all four schemes. Job cancellation affected queued jobs including own jobs. The loss of own jobs by failing the deadline was taken into account. Therefore, users avoided to cancel jobs, when cancellation had an impact on too many own jobs in the queue. Looking at all foreign scheduled jobs 98.1% of the jobs were finished successfully with the reputation-based pricing scheme. The fixed price schemes are above 93%. The discrepancy results from previous two metrics. The overall utility gained from the reputation model was 3% higher than the fixed price scheme. A better utility could be achieved by enhancing the agents with more intelligent tactics. The selection of sites can be differentiated according to value of a job, reputation, deadline and payment in a more detailed approach. Weighing these four parameters can attain a better outcome for all agents. Learning algorithms have to be applied in the next step (Erev and Roth 1998; Tran and Cohen 2003).

	Number of cancelled jobs	Deadline not matched		Finished foreign scheduled jobs	Utility
		Own jobs	Foreign jobs		
Reputation-based pricing	35	6	22	98.1%	103,42%
Fixprice = 0	301	33	74	94.6%	100,00%
Fixprice = 0.5	309	28	64	93.3%	100,16%
Fixprice = 1	328	22	55	93.1%	100,32%

Table 1. Simulation results.

4.3 Application

The reputation-based pricing mechanism is currently implemented in a billing infrastructure called Billing the Grid (BtG). The goal of this infrastructure is to provide a reputation and billing mechanism for particle physics scientists. The introduction of incentives in the scientific Grid will enable an efficient utilization of existing resources. Consumers have the incentive to avoid the submission of redundant jobs as they have to pay for each job. Site administrators have the incentive to keep jobs running and avoid system downtime, since it results in payment loss. This infrastructure provides a graphical user interfaces based on the Gridsphere framework (Novotny et al 2004).

The portlets allow users to submit jobs, get detailed information about the sites and to rate users according to their behaviour (figure 3). The portlets send jobs to and receive information from the middleware gLite¹. Users describe their job requirements in a JDL-file (job description language). The resource broker matches the requirements with the available resources and sends the job to the according site. Consumers still have the ability to choose a certain site for their job. BtG supports the user's decision by providing additional information about the sites reputation. Beside the reputation-based payment model and the global reputation, users can maintain a local reputation table. In case, a specific type of job is not able to run on a particular machine, they can rate the site with a low reputation. Next time, this machine can be avoided, although it has a good global reputation. The graphical user interface eases the management of finished jobs, favourable sites and credit account as

¹ <http://www.glite.org>

well as rating of jobs. Next steps are to test the platform in cooperation with particle physics scientists and compare the performance and behaviour of the users with the simulated data.




Figure 3. Screenshot of the BtG portlet for job submission.

5 CONCLUSION

Grid computing facilitates IT resource sharing among distributed organizations. Particle physics Grids are currently running inefficiently due to redundant job submissions. We proposed an approach to overcome this problem by introducing a reputation-based pricing. This pricing scheme has the advantage of setting incentives for providing resources and consuming them prudently. But, the effort for determining the right price at the right time is avoided, since scientists prefer rather to concentrate on research topic than to make economic decisions. In fact, the dynamic of prices are based on the reputation of the users, which enforce them to behave cooperatively and elude the price specification. We presented a decision model to denote the behaviour in scientific Grids and showed in a simulation that reputation-based pricing can improve collaborative interaction. These results strengthen our approach and illustrated future avenues. A deeper analysis of agent-based modelling and decision will allow a more realistic simulation. Different strategies can be tested within the decision model. Moreover, appropriate reputation mechanisms can be analyzed regarding attacks and incentives in this special decision framework.

Acknowledgement

This work was supported by a grant from the Ministry of Science, Research and the Arts of Baden-Wuerttemberg (AZ: 23-7532.24-38-14/1) to the authors.

References

- Adar, E., and Huberman, B. A. (2000). Free riding on gnutella. *First Monday* 5(10).
- Akerlof, G. (1970). The market for lemons: Quality uncertainty and the market mechanism, *Quarterly Journal of Economics* 84, 488-500.

- Alunkal, B. K., Veljkovic, I., von Laszewski, G., and Amin, K. (2004). *Reputation-Based Grid Resource Selection*.
- Azzedin, F. and Maheswaran, M. (2002). Evolving and managing trust in grid computing systems. *Proceedings of the Canadian Conference on Electrical and Computer Engineering*. 1424-1429.
- Berlich, R., Kunze, M. and Schwarz, K. (2005). Grid computing in Europe: from research to deployment. *Proceedings of the 2005 Australasian workshop on Grid computing and e-research 44*
- Buragohain, C., Agrawal, D. and Suri, S. (2003). A Game Theoretic Framework for Incentives in P2P Systems. *P2P '03: Proceedings of the 3rd International Conference on Peer-to-Peer Computing*. Washington, DC, USA: IEEE Computer Society.
- Buttyn, L. and Hubaux, J. P. (2001). "Stimulating cooperation in self-organizing mobile ad hoc networks," *ACM Journal for Mobile Networks (MONET), Special issue on Mobile Ad Hoc Networks*.
- Cheng, A., and Friedman, E. (2005). Sybilproof reputation mechanisms. *P2PECON '05: Proceeding of the 2005 ACM SIGCOMM Workshop on Economics of peer-to-peer systems* 128-132.
- Dellarocas, C. (2003). The digitization of word-of-mouth: Promise and challenges of online feedback mechanisms. *Management Science* 49 (10) 1407-1424.
- Erev, I. and Roth, A.E. (1998). Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria, *The American Economic Review* 88(4) 848-881.
- Foster, I., Kesselman, C. and Tuecke S. (2001) . The Anatomy of the Grid: Enabling Scalable. Virtual Organizations. *International Journal of High Performance Computing Application*, 15(3) 200-222.
- Heydenreich, B., Mueller, R. and Uetz, M. (2006). Decentralization and Mechanism Design for Online Machine Scheduling. *Proceedings of the Scandinavian Workshop on Algorithm Theory (SWAT)* (S. 136-147). Springer.
- Joseph, J., Ernest, M. and Fellenstein, C. (2004). Evolution of grid computing architecture and grid adoption models. *IBM Systems Journal*, 43 (4) 624-645.
- Jurca, R. and Faltings, B. (2005). Reputation-based pricing of P2P services. *P2PECON '05: Proceeding of the 2005 ACM SIGCOMM workshop on Economics of peer-to-peer systems* (S. 144-149). New York, NY, USA: ACM Press.
- Kamvar, S. D., Schlosser, M. T. and H. Garcia-Molina (2003). EigenRep: Reputation Management in peer-to-peer Networks. *Proc. of the Twelfth International World Wide Web Conference*, Budapest, Hungary.
- Kwok, Y., Song, S. and Hwang, K. (2005). Selfish grid computing: game-theoretic modeling and NAS performance results. *CCGRID '05: Proceedings of the Fifth IEEE International Symposium on Cluster Computing and the Grid (CCGrid'05) - Volume 2* (S. 1143-1150). Washington, DC, USA: IEEE Computer Society.
- Marti, S. and Garcia-Molina, H. (2006). Taxonomy of trust: Categorizing P2P reputation systems. *Computer Networks* , 50 (4), 472-484.
- Novotny, J. Russell, M. Wehrens, O. (2004). GridSphere: an advanced portal framework. *Proceedings of the Euromicro Conference*, 412- 419.
- Porter, R. (2004). *Mechanism design for online real-time scheduling*.
- Ranganathan, K., Ripeanu, M., Sarin, A. and Foster, I. (2004). Incentive mechanisms for large collaborative resource sharing. *CCGRID '04: Proceedings of the 2004 IEEE International Symposium on Cluster Computing and the Grid*, 1-8. Washington, DC, USA: IEEE Computer Society
- Resnick, P., Zeckhauser, R. Friedman, E. Kuwabara, K. (2000). Reputation systems. *Comm. ACM* 43(12) 45-48.
- Schnizler, B., Neumann, D. , Veit, D. and Weinhardt, C. (2006). Trading grid services - a multi-attribute combinatorial approach. *European Journal of Operational Research*, forthcoming.
- Sonnek, J. D. and Weissman, J. B. (2005). A quantitative comparison of reputation systems in the grid., 242-249.
- Tran, T. and Cohen, R. (2003). Modelling Reputation in Agent-Based Marketplaces to Improve the Performance of Buying Agents. *Lecture Notes in Computer Science (2702)* 273-282.

- Vishnumurthy, V., Chandrakumar, S., and Sirer, E. (2003). *KARMA: A Secure Economic Framework for Peer-to-Peer Resource Sharing*.
- Xiong, L. and Liu, L. (2004). Peertrust: Supporting reputation -based trust in peer-to-peer communities. *IEEE Transactions on Knowledge and Data Engineering (TKDE), Special Issue on Peer-to-Peer Based Data Management*.
- Zacharia, G., Moukas, A., and Maes, P. (1999). Collaborative reputation mechanisms in electronic marketplaces., *Track8*, S. 7.