# GRID RESOURCE NEGOTIATION: SURVEY WITH A MACHINE LEARNING PERSPECTIVE


Cyril Briquet
Department of EE & CS
University of Liège
Montefiore Institute, B37
B-4000 Liège, Belgium
email: `C.Briquet@ulg.ac.be`

Pierre-Arnoul de Marneffe
Department of EE & CS
University of Liège
Montefiore Institute, B37
B-4000 Liège, Belgium
email: `PA.deMarneffe@ulg.ac.be`



**ABSTRACT**

Grid computing can be defined as *coordinated resource sharing and problem solving in dynamic, multi-institutional collaborations* [1]. As more Grids are deployed worldwide, the number of multi-institutional collaborations is rapidly growing. However, for Grid computing to realize its full potential, it is expected that Grid participants are able to use one another resources. Resource negotiation (i.e. exchange or trading of resources between Grids) enables Grid participants to face an unstable request environment.

The aim of this position paper is to present a survey of the current state and challenges of resource negotiation research, with a Machine Learning perspective. We support the view that negotiation and learning are intrinsically linked. In particular, we show the expected benefits of integrating Machine Learning techniques with resource negotiation.



**KEY WORDS**

Grid Computing, Resource Negotiation, Machine Learning.


## 1 Introduction

Grid computing can be defined as *"coordinated resource sharing and problem solving in dynamic, multi-institutional collaborations"* [1]. The focus is the use of multiple resources to *"solve a single, large problem that could not be performed on any one resource"* [1]. A Grid Resource Management System (RMS) is further defined as *"the subsystem of a Grid that identifies requirements, matches resources to applications, allocates, schedules, monitors Grid resources over time in order to run Grid applications as efficiently as possible"* [1].

A key concept of Grid computing is the ability of a Grid (via its RMS) to *"negotiate resource-sharing arrangements among a set of participating parties (providers and consumers) and then to use the resulting resource pool for some purpose"* [2]. More precisely, an important objective of an RMS is to establish a mutual agreement between a resource supplier and a resource consumer by which the supplier agrees to supply a capability that can be used to perform some task on behalf of the consumer [3].

Resource negotiation (i.e. exchange or trading of resources between Grids) appears as an important feature to enable Grids to face sudden, transient spikes of computing requests from users. Note that we set our work within a resource negotiation context where each Grid RMS (1) makes its own resource negotiation decisions (i.e. autonomous resource negotiation) and (2) can act as both consumer and supplier of resources, which is different from a context where the objective of an RMS is either only to provision or only to supply resources.

Research about resource negotiation has been until now mainly devoted to resource negotiation protocols, so-called Service Level Agreements (SLA) and mechanisms aiming to balance resource trading so that the resource market established by a set of Grid participants reaches some equilibrium [4]. Research about negotiation objectives and some aspects of decision making in resource negotiation have received comparatively less attention so far.

In this position paper, we review the current state and challenges of resource negotiation research. We highlight the impact that basic Grid computing concepts have upon resource negotiation. We support the view that negotiation and learning are intrinsically linked. We therefore propose that resource negotiation, which is expected to become fully automated, will benefit greatly from an integration with Machine Learning techniques.

The rest of the paper is structured as follows: Section 2 motivates the need for resource negotiation research and states the constraints imposed by the very nature of Grid computing as well as the link between negotiation and learning; Section 3 presents the current state and challenges of resource negotiation resarch; and finally Section 4 concludes this position paper.

## 2 Motivation

Ian Foster proposed the well-known three-point checklist [5] to decide whether a given distributed system is a Grid: (1) coordination of resources not subject to centralized control, (2) using standard, open, general-purpose protocols and interfaces, (3) to deliver nontrivial qualities of service. Grid computing research has until now mainly focused on interoperability issues and useful features for Grid users. The first point in Foster's three-point checklist is now receiving more attention as production Grids are deployed and begin to interact.

Why is resource negotiation an important component of a Grid RMS ? The main reason is that resource negotiation enables Grids to cope with environment instability in terms of requests. In a context of multiple autonomous agents each managing resources (i.e. RMS trading resources on behalf of the Grid participants), resource negotiation has to cope with environment instability in terms of resources as well. So, the negotiated access to additional resources will help to shape (i.e. stabilize) the resource environment of a Grid, so that it can focus on serving its request environment.

Orthogonally to these instability problems, Machine Learning is a domain fast gaining momentum to solve a wide range of problems. Machine Learning techniques, among other things, perform automatic classification of collected data, allowing better decision making based upon some automatically constructed data models. This will certainly be useful in the given context of autonomous RMS.

## 2.1 Impact of the VO Paradigm on Resource Negotiation

Ian Foster's proposed paradigm of Virtual Organization (VO paradigm [2]) has a deep impact on resource negotiation: we thus examine how it shapes the constraints framing resource negotiation. Four important consequences of the VO paradigm are: (1) autonomous nature of VO members, (2) scaling required by the involvement of multiple domains with their potentially conflicting objectives, (3) very dynamic nature of potential agreements due to the ever-changing resource environment and (4) necessary measurement of trust.

**Lack of control over managed resources.** Given the autonomous nature of VO members, the lack of control by an RMS over the resources of other domains is an important issue. It is now beginning to receive full attention. Indeed, *"the Grid resources are not subject to centralized control"* [2]. Interactions with the Grids owning the resources borrowed by an RMS are therefore an important aspect of Grid computing. A worthy approach is to adopt a Multi-Agent System (MAS) perspective where RMSs are considered as autonomous agents [6]. For a given RMS, the behaviour of other domains RMSs is unknown as each is autonomous: one way to obtain a model of their behaviour would be to learn them from externally observable data (e.g. their supply patterns). The management data of other domains RMSs will be unknown as well: by design, data about other RMSs will have to be observed rather than communicated, which has the side effect of considerably reducing the data privacy constraints to take into account.

**Multiple perspectives.** According to the VO paradigm, resources used by an RMS may come from multiple management domains. An RMS and the RMSs of other domains it interacts with may not have a single unified objective. They rather may have different, possibly conflicting objectives and requests. As the scale of inter-Grid resource negotiations grows, there will be a need to deal with more complex policies and more conflicting requests from the involved RMSs. It would therefore be interesting to use models of the resource negotiation behaviour of the other RMSs. Gil & al. state that *"without a knowledge-rich infrastructure, fair and appropriate use of Grid environments will not be possible"* [7].

**Ever-changing environment.** Following the VO paradigm, a Grid will operate onto an ever-changing resource base as members enter and leave the VO. Resource management techniques, including resource negotiation, must therefore be considered within the context for which an absolute knowledge of system state and an absolute control over resource policy and use are not possible. Observing the trading behaviour of other participants during a resource negotiation and then dynamically computing a model of resource availability should help to clarify and frame at any moment the availability of the resources of the other domains. Time series and other machine learning techniques would be good candidates and could use supply patterns as data attributes.

**Untested trust relationships.** With the VO paradigm, organizations trade resources with other organizations they may have never dealt with before. As asserted in Foster's Grid textbook [3], quantifying both expected trust of negotiation partners and cost and benefits of negotiations agreements will be essential to meet VO objectives. It is then important to somewhat remember the terms and outcomes of past interactions. To this end, it would be interesting to develop a measure of trust or perceived reliability based on the supply patterns of the negotiation partners. For example, clustering techniques would help to discover clusters of related behaviours so that a discrete scale of trust can be established.

## 2.2 Machine Learning and Grid Resource Negotiation

Machine Learning (ML for short) is aimed at identifying patterns in data. Management data (i.e. data about the Grid) can be used to steer adaptively the behaviour of the Grid. Nonetheless this data is not immediately available: it has to be extracted, mined, from raw data collected by observing the behaviour of the Grid.

ML techniques are useful tools to perform such data mining tasks [8]. These techniques operate on a set of examples (e.g. data items or objects) each composed of a set of attributes and output relationships between or about these examples. We can mention classification (put instances into some predefined classes), unsupervised learning (learn relationships between the attributes), clustering (discover clusters of examples that belong together), reinforcement learning (select an optimal decision. with a positive, implicit feedback signal only). ML techniques can also perform a regression (learn to predict a numeric quantity instead of a class). Another important ML technique is the analysis of time series, to predict the next element from a set of past elements.

The achievements of the ML research field are solid enough to enable the embedded use of ML techniques into other research fields. The idea of using ML techniques in Grid RMS is not new and has been put into practice successfully. For instance, Punch [9] uses instance-based learning, regression and nearest neighbour techniques to model and predict application performance. The Network Weather Service [10] uses time series to predict on the fly the state of the network. Some recent examples are related to resource allocation [11], [12]. However, integrating Machine Learning techniques with resource negotiation has not yet been systematically considered.

To close this short introduction to the use of ML in Grid computing, we explain why we support the view that negotiation and learning are intrinsically linked. Consider the following quotations: *"negotiation can be viewed as a distributed search through a space of potential agreements"* [13] & *"Machine Learning is searching through a model space"* [14]. We could then infer that negotiation is a particular form of distributed learning. Indeed, to negotiate (i.e. to have influence over a trading partner), an agent must be able to convince it to act in a particular way [13] by making proposals, trading options or offering concessions to come to a mutually acceptable agreement. It then appears that learning from past interactions with trading partners can be a key enabler of negotiation agents success.

### 2.3 Behaviour Analysis in Agent-based Resource Negotiation

In any agent-based resource negotiation scenario (including the chosen context of autonomous consumer/supplier RMS), a matter that any agent should pay attention to is the expected behaviour of other agents with whom negotiations are conducted. This issue goes much beyond those arising from the initial lack of trust between trading partners. Indeed, dealing with the ever-changing environment considered by the VO paradigm is an issue much wider than taking into account members entering and leaving the VO.

One must consider the fact that the openness of the VO paradigm comes with the darker side of potentially dealing with sometimes not-so-honest RMS. A naive or static resource negotiation service would exhibit behaviour patterns that could be taken advantage of by any unscrupulous RMS. The responsible behaviour of other RMS cannot be taken for granted, despite what might sometimes be suggested (*"service acquisition implies guarantee of service"* [3]). Instead, and well-known in the peer-to-peer community, the free-riding behaviour (consuming resources without ever paying back [3]) would alone justify the need for a resource negotiation service aware of its environment and in particular of behaviour patterns.

To this end, some authors propose to encode the *"quality of experience"* in resource negotiation [6]. Furthermore, it would be helpful to use behaviour engineering techniques (defined as the engineering of the resource negotiation behaviour of an RMS) able to limit the effect of unexpected or unfriendly behaviour of trading partners. The use of ML techniques into a resource negotiation service could support such behaviour engineering techniques. For example, an RMS could first compute models of other RMSs behaviour (e.g. based on observed supply patterns) and then compute for each of these a set of consistent models that have led to good outcomes (i.e. reliable supplying) during past interactions.

All RMSs of a given Grid could be perfectly honest or not yet equipped with behaviour engineering techniques (e.g. they respond neither to incentives nor to the behaviour of the cognizant RMS). In such an environment, one might question the use of a cognizant RMS relying on ML techniques. Negotiation rounds would be more simplistic because interactions would be limited by the capacities of the less cognizant RMS. However, as the cognizant RMS would be more aware of its environment, it would benefit from a better planning and a better capacity to profit from opportunities.

## 3 State and Challenges of Resource Negotiation Research

We now examine the state of resource negotiation research and the challenges that need to be tackled. Automatic resource negotiation research has been previously classified into 3 topics [13]: (1) negotiation protocols, (2) negotiation objects (i.e. what is negotiated) and (3) decision making models. The first topic can be seen as the *how* (from a communication perspective) of resource negotiation, the second topic would be the *what* and the third topic would also be the *how* (from a processing perspective).

We argue that a fourth aspect, the *when/why* (i.e. negotiation objectives), of resource negotiation should also be taken into account. Accordlingly, this section reviews the four highlighted topics.

### 3.1 Negotiation Objects

An RMS seeks to stabilize its resource environment through resource negotiation so as to exhibit a more predictable behaviour. It is necessary that its trading partners produce *"commitments (contracts) about the willingness to provide a service and the characteristics, or quality, of its provision"* [6]. As might be expected, the contract is an important concept in resource negotiation. In practice, a contract defining what resources are supplied and on what terms can be detailed by a so-called Service Level Agreement [15].

To enforce the terms of a contract resulting from resource negotiation, there must be some form of contract monitoring [3], either centralized or autonomous. Monitoring the enforcement of contracts allows to dynamically renegotiate or terminate them if they are breached or if resource requirements of one of the trading partners change before contract completion [3]. An example of such a recent architecture for resource usage SLA specification and enforcement is GRUBER [16]. Data mining techniques can be used to extract useful patterns from the data produced by the monitoring activities.

### 3.2 Negotiation Objectives

Not far from concerns about negotiation objects are the concerns about negotiation objectives. Indeed, however closely related these two topics might seem, negotiations objectives should be distinguished from negotiation objects. Indeed, studying *what* resources can be negotiated (which is the purpose of negotiation objects) is different from studying *when* and *why* these should be negotiated (which is the purpose of negotiation objectives).

The focus of an RMS can be application performance [3], system performance, user satisfaction [17] or VO administrator satisfaction, maybe further than classic performance metrics such as average resource utilization, average response time, average job completion, average job re-planning, workload completion time [18]. Whatever the focus, in the long-term a multicriteria approach should prevail and take all of them into account, as the objective of the RMS is basically to automate scheduling and resource management to *"minimize stakeholders' interventions"* [1]. A multicriteria approach seeks a *"compromise solution to increase the level of satisfaction of many stakeholders"* (i.e. VO participants and administrators) *"and combine different points of view"* [1].

To service incoming requests, the RMS has to produce resource requirements. It also has to set resource negotiation objectives, given both the produced resources requirements and the stakeholders-composed RMS focus. In this perspective, the resource negotiation service can be said to be *responsive*. There is however another perspective to be considered: an RMS can perform some resource trading without having any incoming request to service. In this perspective, the resource negotiation service can be said to be *proactive*.

Negotiation service proactiveness can be useful to accumulate Access Potential for use at a later time. In other words, the resource negotiation service can proactively acquire resources for some time when it has predicted these would be needed soon (i.e. using the Access Potential). It can also proactively supply resources because it has predicted these would not be needed for some time (i.e. building Access Potential) . . . hoping there will be a payback later when most needed.

Therefore, while conflicting interests from multiple Grid participants may be hard to manage, heterogeneity of focus (i.e. different objectives) in a set of Grid participants has some advantage, after all. With homogeneous needs and assets and when there are tight deadlines to be met, load balancing between Grid participants naturally emerges but little resource trading takes place because each keeps all its resources committed to its own use first. On the other hand, if there are different focuses, overall system utilization and application performance can both be high. When a Grid participant has to meet strict deadlines, it can use previously built Access Potential (e.g. using external resources) in order to reach high application performance. When it has few demanding requests to service, it can build Access Potential (e.g. lending its resources) in order to reach high system utilization.

The problem of connecting the focus of an RMS to resource requirements and negotiation objectives is beginning to gain some attention. See for example the study of resource requirements translation across abstraction layers [19] and the existing GrADSoft system [20] (where the scheduler and resource negotiator are merged). The problem of managing Access Potential is only beginning to become explicitely recognized as a very important topic to be addressed for Grid computing to live up to its expectations. Much research is needed for this problem, for it has received limited attention so far [19].

Once chosen, the RMS focus has to be translated into lower level requirements. Coupled with actual request data, the RMS focus has to be translated first into resource requirements, then into negotiation objectives. In this case, the translation is done to allow the use of Access Potential. The RMS focus alone has also to be translated into negotiation objectives at the negotiation service level to allow the building of Access Potential. This is a task which could benefit from ML modelling techniques by using predictive translation models. Finally, we argue that research is needed to study the translation into negotiation objectives of the RMS focus, both alone and with associated request data, and always within a multicriteria approach.

### 3.3 Negotiation Protocols

An important aspect of resource negotiation is of course the standardization of the exchange of messages between trading partners so that negotiation data can be transformed, composed/decomposed, managed like any other resource and dynamically modified as agreements are breached or not concluded [6]. This calls at least for the establishment of a standard negotiation protocol, possibly to be added into common Grid toolkits, such as the Globus Toolkit.

SNAP (Service Negotiation and Acquisition Protocol) is a protocol which provides lifetime management and at-most-once creation semantics for remote SLAs (Service Level Agreement) [15]. It follows a classic client-server RPC pattern. WS-Agreement is a protocol from a protocol stack proposed by the Global Grid Forum to allow resource suppliers and consumers to negotiate resources by means of SLAs [21]. WS-Agreement seems to be currently considered as an important step towards an automated resource negotiation service [6].

Negotiation protocols essentially define the structure and not the content of negotiation agreements. Once a protocol is chosen, its users would tend to keep it constant over time. Therefore, ML techniques would not be helpful for this aspect of resource negotiation.

### 3.4 Decision Making Models

Another important aspect of resource negotiation is its decision making process: it consists essentially of selecting trading partners and agreeing to, refusing or proposing negotiation agreements. The trend currently dominating research related to resource negotiation decision making models is market-driven or market-based negotiation, also called computational economy. It stems from the observation that some incentive must be offered to all resource suppliers to sustain the interest of the RMS to do resource negotiation on a regular basis. It is then only natural to propose a *"Grid economy as a model for managing and handling requirements of both Grid providers and consumers"* [22].

A market mechanism can be defined as a kind of competitive balance protocol that adjusts the price of a valuable resource given the demand for it, until demand matches supply [4, 23]. There are many design choices [24, 25] involved in a market mechanism. The interest of market-based negotiation is mainly twofold: enable negotiation agents to find the desired resources at the lowest cost possible and to stabilize the price of traded resources. With few exceptions [26], most computational economy research has considered a centralized organization. Given the partial similitary with resource location, a peer-to-peer approach would allow a decentralized organization with its expected benefits.

A common problem in market-based negotiation is that *"most research needing cost has [...] assumed it could be supplied by some oracle agent"* [27]. What it means is that market-based mechanisms may be efficient at stabilizing the resource market but do not yet sufficiently explain how to take into account the valuation of the resources by the Grid policies (i.e. market-based mechanisms simply suppose that if a resource is important for a Grid, its RMS will demand much of it). This problem can certainly be related to the definition of negotiation objectives given in this paper.

In market-based negotiation, another common problem is that each agent seeks to maximize its own utility or benefits in a short-term perspective only, or to minimize risks associated with the considered agreement [28]. While this approach is certainly worthy in totally unstable resource markets, without structure or repeating patterns, it does not directly promote the building of trust and lasting trading relationships which could bring more benefit in the long-term. Many Grid participants will certainly benefit from advances in the study of long-term trading relationships, including departments of the same university or subsidiaries of the same global corporation, as their resource markets will, over time, exhibit trading behaviour patterns that can be taken advantage of.

Challenges in decision making models include achieving a fully decentralized organization, taking into account the valuations of the resources traded by the agents (i.e. the so-called resource accounting problem) and considering the benefits or utility of long-term trading relationships. The first challenge is related to taking into account the autonomous nature of the RMS negotiation agents. The second and third challenges can both make good use of ML techniques. For example, the long-term behaviour (in terms of supply patterns) and perceived motivations of other RMS negotiation agents could be learnt better with each resource exchange.

## 4 Conclusions

Resource negotiation is an important capacity of a Grid RMS that enables Grids to cope with environment instability in terms of requests. Resource negotiation allows Grids to also cope with environment instability in terms of both resources availability and trading partners behaviour.

Supporting the view that negotiation and learning are intrinsically linked, we have motivated the integration of resource negotiation with Machine Learning techniques. This motivation comes from core Grid computing concepts, the intrinsic nature of negotiation and the expected negotiation environment.

We have then reviewed the current state and challenges of resource negotiation research, including negotiation objects, negotiation objectives, negotiation protocols and decision making models. As we have pointed out, with the exception of negotiation protocols, nearly all resource negotiation challenges may be more easily met with the help of Machine Learning techniques. This is only natural when considering that it will be expected of a Grid RMS to exhibit more autonomous behaviour as Grid computing matures. Considering this perspective, this position paper can be viewed as an extension, focused on negotiation, of Foster, Jennings & Kesselman recent call for a tight integration of Grid computing and Multi-Agent Systems [6].

## Acknowledgements

We want to thank Sébastien Jodogne and Raff Brancaleoni for their discussions, support and technical proofreading. We also want to thank Claire Kopacz for her spelling check, Krzysztof Rzadca for interesting references and an anonymous reviewer for useful suggestions.

## References


[1] J. Nabrzyski, J.M. Schopf & J. Weglarz (Editors), *Grid Resource Management: State of the Art and Future Trends* (Boston: Kluwer Academic Publishers, 2003).

[2] I. Foster, C. Kesselman & S. Tuecke, *The Anatomy of the Grid: Enabling Scalable Virtual Organizations*, Int. J. Supercomputer App., 15(3), 2001.

[3] I. Foster & C. Kesselman, *The Grid 2: Blueprint for a New Computing Infrastructure* (San Francisco: Morgan Kaufmann, 2004).

[4] R. Buyya, D. Abramson & S. Venugopal, *The Grid Economy*, Proc. of the IEEE, Special Issue on Grid Computing, 93 (3), New York, USA, 2005.

[5] I. Foster, *What is the Grid ? A three Point Checklist*, GridToday, July 2002.

[6] I. Foster, N.R. Jennings & C. Kesselman, Brain Meets Brawn: Why Grids and Agents Need Each Other, *Proc. Int. Conf. on Autonomous Agents and Multi-Agent Systems*, New York, USA, 2004.

[7] Y. Gil, E. Deelman, J. Blythe, C. Kesselman & H. Tangmunarunkit, *Artificial Intelligence and Grids: Worflow Planning and Beyond*, IEEE Intelligent Systems, Special Issue on E-Science, 19(1), 2004.

[8] T.M. Mitchell, *Machine Learning* (New York: McGraw-Hill, 1997).

[9] N.H. Kapadia, J.A.B. Fortes & C.E. Brodley, Predictive Application-Performance Modeling in a Computational Grid Environment, *Proc. 8th IEEE Int. Symposium on High Performance Distributed Computing*, Redondo Beach, California, USA, 1999.

[10] R. Wolski, N. Spring & J. Hayes, *The Network Weather Service: A Distributed Resource Performance Forecasting Service for Metacomputing*, J. Future Generation Computing Sys., 15(5-6), 1999.

[11] A. Galstyan, K. Czajkowski & K. Lerman, Resource Allocation in the Grid Using Reinforcement Learning, *Proc. Int. Conf. on Autonomous Agents and Multi-Agent Systems*, New York, USA, 2004.

[12] C.W. Khuen, C.H. Yong & F. Haron, A Machine Learning Method for Resource Allocation in Multi-Agent Negotiation System, *Proc. GridAsia*, Biopolis, Singapore, 2005.

[13] N.R. Jennings, P. Faratin, A.R. Lomuscio, S. Parsons, C. Sierra & M. Woolridge, *Automated Negotiation: Prospects, Methods and Challenges*, Int. J. Group Decision and Negotiation, 10(2), 2001.

[14] L.A. Wehenkel, *Automatic Learning Techniques in Power Systems* (Boston: Kluwer Academic Publishers, 1997).

[15] K. Czajkowski, I. Foster, C. Kesselman, V. Sander & S. Tuecke, SNAP: A Protocol for Negotiating Service Level Agreements and Coordinating Resource Management in Distributed Systems, *Proc. 8th Workshop on Job Scheduling Strategies for Parallel Processing*, Edinburgh, Scotland, 2002.

[16] C. Dumitrescu & I. Foster, GRUBER: A Grid Resource SLA Broker, *Proc. Euro-Par*, Lisboa, Portugal, 2005.

[17] K. Kurowski, J. Nabrzyski & J. Pukacki, User preference driven multiobjective resource management in Grid environments, *Proc. 1st Int. Symposium on Cluster Computing and the Grid*, Brisbane, Australia, 2001.

[18] C. Dumitrescu & I. Foster, Usage Policy-based CPU Sharing in V.O., *Proc. 5th IEEE/ACM Int. Workshop on Grid Computing*, Pittsburgh, USA, 2004.

[19] A. Dan, C. Dumitrescu, M. Ripeanu, Connecting Client Objectives with Resource Capabilities: An Essential Component for Grid Service Management Infrastructures, *Proc. 2nd Int. Conf. On Service Oriented Computing*, New York, USA, 2004.

[20] F. Berman & al, *The GrADS Project: Software Support for High-Level Grid Application Development*, Int. J. High Perf. Computing App., 15 (4), 2001.

[21] HPC4U Project deliverable, *Negotiating on QoS aspects, version 1.2*, `http://www.hpc4u.org/public/docs/HPC4U_D12_NegotiatingQoS.pdf`, December 2004.

[22] R. Buyya, *Economic-based Distributed Resource Management and Scheduling for Grid Computing*, PhD Thesis, Monash U., Melbourne, Australia, 2002.

[23] P. Tucker & F. Berman, *On Market Mechanisms as a Software Technique*, Tech. Report, TR CS96-513, University of California, San Diego, USA, 1996.

[24] C.S. Yeo & R. Buyya, *A taxonomy of market-based resource management systems for utility-driven cluster computing*, Tech. Report, GRIDS-TR-2004-12, GCDS Laboratory, University of Melbourne, Melbourne, Australia, 2004.

[25] R. Buyya, D. Abramson & J. Giddy, Economy Driven Resource Management Architecture for Computational Power Grids, *Proc. Int. Conf. on Parallel and Distributed Processing Techniques and Applications*, Las Vegas, USA, 2000.

[26] M.P. Wellman, W.E. Walsh, P.R. Wurman, J.K. MacKie-Mason, *Auction protocols for decentralized sched.*, Games and Economic Behavior, 35, 2001.

[27] J.M. Schopf & B. Nitzberg, *Grids: The Top Ten Questions*, Tech. Report, CS Dept. TR #CS-00-05, Northwestern University, Evanston, USA, 2000.

[28] D. Irwin, J. Chase & L. Grit, Balancing Risk and Reward in Market-Based Task Scheduling, *Proc. Int. Symposium on High Performance Distributed Computing, HPDC-13*, Honolulu, USA, 2004.