



**FOSTERING EFFICIENCY OF COMPUTATIONAL RESOURCE
ALLOCATION - INTEGRATING INFORMATION SERVICES
INTO MARKETS**

Journal:	<i>17th European Conference on Information Systems</i>
Manuscript ID:	ECIS2009-0507.R1
Submission Type:	Research Paper
Keyword:	Grid Computing, Agent-based model, Information exchange, Marketplaces



FOSTERING EFFICIENCY OF COMPUTATIONAL RESOURCE ALLOCATION - INTEGRATING INFORMATION SERVICES INTO MARKETS

Borissov, Nikolay, University of Karlsruhe, Englerstrasse 14, 76131 Karlsruhe, Germany,
borissov@iism.uni-karlsruhe.de

Brunner, René, Technical University of Catalonia, 08034 Barcelona, Spain,
rbrunner@ac.upc.edu

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

Freitag, Felix, Technical University of Catalonia, 08034 Barcelona, Spain,
felix@ac.upc.edu

Navarro, Leandro, Technical University of Catalonia, 08034 Barcelona, Spain,
leandro@ac.upc.edu

Weinhardt, Christof, University of Karlsruhe, Englerstrasse 14, 76131 Karlsruhe, Germany,
weinhardt@iism.uni-karlsruhe.de

Abstract

The application of market mechanisms for the allocation of computing services is a demanding task, which requires bridging economic and associated technical challenges. Even if the market-based approach promises an efficient allocation of computing services, the wide heterogeneity of consumer requirements and the diversity of computational services on provider side are challenging the processes of finding, allocating, and using an appropriate service in an autonomous way. The focus of the most papers is mainly devoted to the optimization embedded in the allocation process itself. However, we think that the optimization process starts much earlier and contains the information gathering until the final market-based resource allocations.

In this paper we introduce an integrated framework for market-based allocation of computing services, integrating information retrieval of market information, prediction models, bidding strategies and market mechanisms. As proof-of-concept, we implemented a first prototype of the framework. Furthermore, we propose a methodology for evaluating strategic behavior in market mechanisms with bidding strategies using market information and statistical prediction techniques. First simulation results show strategic behavior in selected market mechanisms by applying the proposed techniques.

Keywords: Market Mechanisms, Market Information, Bidding Strategies, Evaluation.

1 INTRODUCTION

Grid and Cloud Computing become not only in science increasingly popular, but also in industry. Prominent examples are Sun Microsystem's Network.com, Amazons Elastic Compute Cloud (EC2) and its Simple Storage Service (S3) and SimpleDB Service. The companies frequently offer a fixed pay-per-use price for static resource configurations. Fixed prices can lead to inefficient utilization, reward and usability, as it does not reflect the dynamics of the market supply and demand. Efficient provisioning and usage of computational resources as well as pricing in vigorous environment like Grid Computing is not manually manageable. Such processes should be automated with no or minimal human interaction. Hence, market mechanism and strategic behavior play an important role for the design of the environment. Self-organization and automatic adaptation to the changing market conditions are key prepositions for efficient offering and consuming of computational resources.

An efficient allocation of consumers' jobs to providers' resources is a complex task, where agent decisions on resource provisioning, market-based allocation and information processing is included. Moreover, the wide heterogeneity of computational resources challenges the process of finding an appropriate resources for given consumer preferences. Since demand and supply of computational resources fluctuates in the course of time, information about current and future resource utilizations and prices are often not known a priori to the participants. In this case consumer and provider agents try to maximize their utilities by generating bids based on available information. This enables strategic behavior both on provider as well as on consumer side.

The knowledge about a market is essential for the design of efficient bidding strategies. Examples are computational approaches incorporating game theory that allow predicting the future through forecasting or learning rules on former or actual trading information. Bergemann's survey (Bergemann and Valimaki 2007) shows that the economic aspect of information acquisition in market mechanisms got more attention by the economic research community. Moreover, the study demonstrates the importance of the economic information disclosure for market participants. The need for this information lies in both being able to apply efficient economic strategies and to feed business models, which are behind these strategies.

The first contribution of this paper is to survey and bring together results for market-based allocation of computational resources. More specifically, to structure the allocation processes in individual steps, we go further in regard to specify mechanisms for information provision and retrieval, price prediction, bid generation and market-based resource allocations. Based on them, we propose an integrated system for market-based computational resource allocation, which organizes the discussed processes. Finally, we specify a common evaluation scenario and present first results.

Following the structure of the paper, section 2 presents promising market mechanisms and bidding strategies. Section 3 shows the role of market information and prediction models for improving allocation efficiency. Thus, section 4 shows an integrated model of a market-based allocation system. Evaluation methodology and first results are depicted in section 5 and related work in section 6.

2 MARKET-BASED RESOURCE ALLOCATION

Auction and strategy selection are closely connected in the sense that a given choice of strategy should affect the choice of auction, and vice versa. For example, some bidding strategies perform well in a Continuous Double Auction, but not in other auctions such as the Dutch auction. This also implies that the choice which auction to participate in depends on the available strategies. Other factors to take into account are the bidding rules – when, how and what to bid, current and average prices, transaction costs, and pricing schemata in the different auctions.

In this section, we discuss potential market mechanisms for allocation of computing services as well as potential bidding strategies enabling the automation of the provisioning and usage processes.

2.1 Market Mechanisms for Computing Services

Economic models for resource scheduling are widely explored in the literature (Wolski et al. 2001, Parkes et al. 2004, Lai et al. 2005, Nassif et al. 2007). According to the allocation modes, scheduling mechanisms can be distinguished into mechanisms, which execute periodically, called also “offline mechanisms”, and mechanisms which execute continuously, called also “online mechanisms”. Mechanisms like Vickrey, English, Dutch, double auctions (Grosu and Das 2006) and combinatorial mechanisms (Bapna et al. 2005) involve a scheduling problem that is complex. The complexity of such mechanisms drives research to look at allocation models, which can be adapted to real-world scenarios and requirements.

Following section presents promising market mechanisms, implemented within the integrated model (section 4) to evaluate market-based allocation of computing resources as well as strategic behavior using the market information from the proposed market information service.

2.1.1 Continuous Double Auction

The CDA has been widely employed in experimental economic studies, where different agent strategies (Gjerstad and Dickhaut 1998, He et al. 2003, Vytelingum et al. 2008) are investigated for applying automated bidding behavior by the provisioning and usage of resources. The matching in CDA is mostly based on a single value – price, which in a computing resource scenario incorporates the consumer’s preferences for a job i.e. his value v_j per time unit. Based on a preferred bidding strategy (see section 2.2) a consumer generates and submits a bid price ($b_j \leq v_j$) to the CDA market. Respectively, based on a preferred bidding strategy, a resource provider also generates an offer price ($o_i \geq v_j$) and submits it to the market. In the provider case v_j represents the reserve price (minimum price) for using the provider’s machine for a unit of time. In case of a match, the consumer executes immediately his job on the provider’s machine with a payment to the provider, calculated by the winning bid and offer prices. The market price of a match is calculated using the pricing schema *k-pricing* (Satterthwaite and Williams 1989): $p_m = kb_j + (1 - k) o_i$, with $k = 0.5$ in our case.

2.1.2 Decentralized Online Machine Scheduling

In the case of the Decentralized Local Greedy Mechanism DLGM (Heydenreich et al. 2006), each time a job j arrives on the consumer side, his bidding agent creates a request in the form $t_j = \{r_j, d_j, v_j\}$ with its release date r_j , duration d_j and valuation v_j , and reports this to all known providers. The valuation v_j expresses the costs of the job for waiting one additional time unit in the provider machine’s queue. Based on the received bids, the machines perform real-time planning based on a local scheduling policy – if job j has a higher priority value than $k \Leftrightarrow \frac{v_j}{d_j} \geq \frac{v_k}{d_k}$, then j is scheduled before job k in the waiting queue. Depending on the current local waiting queue, the machine i reports a tentative (ex-ante reported) completion time \tilde{C} and payment $\tilde{\pi}$ to the agent of job j . The payment $\tilde{\pi}$ contains the aggregated compensations to all job-agents whose jobs are currently waiting at machine i and are delayed due to allocation of j .

Upon receiving information about its tentative completion time and required payments, the job-agent makes a binding decision to queue at certain machine i , and pays $\tilde{\pi}$ to the delayed jobs. The decision on which machine to submit the job is taken based on the consumer’s utility function $u(j) = -\sum_i v_j * \tilde{C}_{j,i} + \tilde{\pi}_{j,i}$, which selects a provider machine’s offer i with the shortest weighted tentative completion time $v_j * \tilde{C}_{j,i}$ and tentative compensation payments $\tilde{\pi}_{j,i}$. The providers applying the DLGM mechanism do not behave strategically and do not get compensation for the use of their services. The payments are divided only among the consumers. Heydenreich et al. showed that the plain DLGM mechanism achieves a performance ratio of 3.281 against an optimal offline scheduling mechanism.

2.1.3 Fix-Price Technical Scheduling

In order to compare the outcomes of the introduced market mechanisms, we implemented FIFO as baseline scheduling mechanism. The agents submit their jobs to a central scheduler, which schedules and executes them through the FIFO scheduling policy on the various registered provider machines. For each machine the providers request a fixed reservation price.

2.2 Defining Bidding Strategies

Bidding strategies for market-based scheduling are well explored in the economic literature (Reeves et al. 2005, Li and Yahyapour 2006a, Vytelingum et al. 2008). Wellman et al. 2007 give an overview of the various agents and their strategies taken place in the trading agent competition. In general, bidding strategies can be classified into non-adaptive, where the generated bid price do not depend on past and market information e.g. Truth-Telling, and ZI as well as adaptive strategies, where the generated bid price depends on past and market information e.g. ZIP and GD.

Although a simple strategy, truth-telling is essential in case of strategy-proof mechanisms, where in such mechanisms this strategy guarantees to obtain optimal payoffs, no matter what strategies are adopted by the others. A comparison of state-of-the-art bidding strategies is evaluated by Das et al. 2001. The authors of Adaptive-Aggressiveness (AA) strategy (Vytelingum et al. 2008) describe and evaluate a novel bidding strategy, which implements short and long-term learning, considering and responding also to dynamic market fluctuations. Simulation results show that the AA outperforms the ZIP and GD (Tesauro and Bredin 2002) strategies.

Although mainly designed for financial markets, these strategies are fundamental for the design and evaluation of bidding strategies and market mechanisms for Grid and Cloud services (Foster et al. 2008). Following sections present bidding strategies, which we aim to adopt in a Grid market scenario.

2.2.1 Zero-Intelligence Plus Strategy

Zero-Intelligence Plus (ZIP) agents are widely explored and become a popular benchmark for agents trading on continuous double auctions (Das et al. 2001). In Cliff 1997, the author showed that zero intelligence (ZI) agent strategy is not enough, since the bids are uniform generated between a given interval and not depend on current market information or past bids. They introduced ZIP agents, which use public market information to adapt the bid price of a certain market. Experiments (Das et al. 2001) show that ZIP agents perform better than (non-expert) human traders on CDA markets as well as it converges quickly to equilibrium price by high demand and supply.

In most experiments in the literature (Cliff 1997, Tesauro and Das 2001, Das et al. 2001), the ZIP agents receive information of all bids and offers, whether they are accepted or not. There are, however, experiments showing that ZIP based agents perform well also in auctions where the agents only receive information on the winning bids (Bagnall and Toft 2005). In our scenario we will adopt the ZIP strategy in the same way, considering only the winning bids and clearing price of the market.

2.2.2 Gjerstad-Dickhaut Strategy

The Gjerstad-Dickhaut, GD (Gjerstad and Dickhaut 1998, Tesauro and Bredin 2002) bidding strategy is also mainly developed and evaluated in CDA markets. Compared to ZIP, this mechanism is memory based, thus it is using historic market information of the last M bids and offers, H_M , in order to calculate a “belief” function $f(p)$ estimating the probability for a bid or offer to be accepted at price p .

$$\text{Provider belief function } f_o(p) = \frac{\sum_{\substack{o \in H_M \\ o \geq p}} i + \sum_{\substack{b \in H_M \\ b \geq p}} i}{\sum_{\substack{o \in H_M \\ o \geq p}} i + \sum_{\substack{b \in H_M \\ b \geq p}} i + \sum_{\substack{o \in A_{H_M} \\ o \leq p}} i}$$

$$\text{Consumer belief function } f_b(p) = \frac{\sum_{\substack{b \in H_M \\ b \leq p}} i + \sum_{\substack{o \in H_M \\ o \leq p}} i}{\sum_{\substack{b \in H_M \\ b \leq p}} i + \sum_{\substack{o \in H_M \\ o \leq p}} i + \sum_{\substack{b \in H_M \\ b \notin A_{H_M} \\ b \geq p}} i}$$

For the provider, $\sum_{\substack{o \in H_M \\ o \geq p}} i$ represents the number of accepted offers in H_M with offer price $o \geq p$,

$\sum_{\substack{b \in H_M \\ b \geq p}} i$ is the number of accepted bids H_M with bid price $o \geq p$ and $\sum_{\substack{o \in H_M \\ o \notin A_{H_M} \\ o \leq p}} i$ is the number of

unaccepted offers in H_M , $o \notin A_{H_M}$, with offer price $o \leq p$, respectively in opposite for the consumer. Interpolation is applied for prices at which no bids and offers are registered in H_M . The bid or offer price is calculated by the product $f(p) * g$, where g is the calculated gain for a trade at that price, $g = p - v$ for providers and $g = v - p$ for consumers, where v is the valuation.

Although mainly evaluated in CDA markets, GD is applicable as a state-of-the-art strategy in markets for computing resources to the estimation of bids, considering past bids and market information.

2.2.3 Q-Strategy

Q-Strategy (Borissov and Wirström 2008), is a novel bidding strategy adopting a reinforcement learning approach with an e-greedy selection policy. Using the Q-Strategy, an agent explores the environment, *exploitation phase*, with a probability of ϵ , learning available provider machines and the reward of executing a job on them. With a probability of $1-\epsilon$, *exploitation phase* the strategy exploits the collected knowledge (i.e. market information) for generating bids more intelligently.

Using Q-Strategy, the agents can bid and adapt also in markets, where market information is incomplete. Simulations show that Q-Strategy adapts to market dynamics (fluctuating prices and changing utilities from providers machines) in different market mechanisms and against bidding strategies like ZIP and Truth-Telling. A common drawback of reinforcement learning algorithms is that learning the optimal bid price needs “training time”. In worst case, Q-Strategy will perform worse at the beginning, but will converge to optimal values in the time (Watkins and Dayan 1992).

3 THE ROLE OF MARKET INFORMATION

Bidding strategies (section 2.2) and prediction techniques (sections 3.1.1 and 3.1.3) require market information in order to efficiently bid in selected markets. However, in a distributed environment like in open and wide Grid markets like in SORMA (<http://www.sorma-project.eu/>) such information is not locally accessible. (Bergemann and Valimaki 2007) demonstrate the importance of economic information disclosure and show the increased attention, paid to economic information acquisition. Traders require information that enables them to deduce entry prices for available markets and trading times. Centralized markets are able to furnish the current information necessary for simple or sophisticated bidding strategies - such as ZIP agents (Preist and van Tol 1998) or human traders. Therefore, this is an issue for distributed and segmented markets. Distribution and segmentation of markets result in loss of information such as prices, products and effective supply (Brunner et al. 2008). An efficient information system should allow participants to choose a compromise between exact global information and partial information. Furthermore, aggregated, anonymous and summarized information contribute to scalability and may in most cases be sufficient. We conclude the following hypothesis for the influence of information to distributed markets.

Bidding strategies need information about a state of the market, commonly through the offered resource type and price dynamics in time. The main focus here is the treatment of market information by agent’s bidding strategies applying statistical price prediction techniques.

Statistics are used in markets to measure current conditions as well as to forecast financial or economic trends. Indicators are used extensively in technical analysis to predict changes in stock

trends or price patterns. Economic indicators quantifying current economic and industry conditions are used to provide insight into the future demand for commodity goods. In our approach, we use statistical prediction methods for price forecasting.

3.1.1 Trend Extrapolation

The trend extrapolation indicates new trends within the market. In computation markets this is important when newer or higher performance resources entered the market, which will start a decreasing trend for older resources. We will use different kind of moving averages that are indicators in technical analysis of market charts and show the average value of a security's price over a set period. Moving averages are used to emphasize the direction of a trend and to smooth out price and volume fluctuations, or noise. Most bidding strategies implement learning mechanisms, which use statistical prediction models. The ZIP bidding strategy, for example, uses a Widrow-Hoff gradient algorithm, a learning rule (Preist and van Tol 1998), which is a variation of the moving average rule.

In time series analysis, the *Moving Average* model is common approach to model univariate time series models. A *Simple Moving Average* (SMA) is the unweighted mean of the previous n data points. For example, a 10-day simple moving average of closing price is the mean of the previous 10 days' closing prices. If those prices are $P_M, P_{M-1}, \dots, P_{M-9}$ then the formula is $SMA = \frac{P_M + P_{M-1} + \dots + P_{M-9}}{10}$

In all cases a moving average lags behind the latest data point, simply from the nature of its smoothing. An SMA can lag to an undesirable extent, and can be disproportionately influenced by old data points dropping out of the average. This is addressed by giving extra weight to more recent data points, as in the weighted and exponential moving averages.

A SMA calculates the average of a whole period to predict the future price, which has poor ability to react to new market events. Therefore we will evaluate our approach with a *Weighted Moving Average* (WMA) to give more weight to recent prices: $\bar{x} = \frac{1}{n} \sum_{i=1}^n x_i = \left(\frac{1}{n}\right)x_1 + \left(\frac{1}{n}\right)x_2 + \dots + \left(\frac{1}{n}\right)x_n$, where $\left(\frac{1}{n}\right)$ are the weights, which sum to 1.

An Exponential Moving Average (EMA) is based on the WMA and applies weighting factors which decrease exponentially. The weighting for each older data point decreases exponentially, giving much more importance to recent observations while still considering older observations

$$EMA = \frac{p_1 + (1 - \alpha)p_2 + (1 - \alpha)^2 p_3 + (1 - \alpha)^3 p_4 + \dots}{1 + (1 - \alpha) + (1 - \alpha)^2 + (1 - \alpha)^3 + \dots}$$

3.1.2 Cyclic and seasonal component models

Cyclic and seasonal prediction models are applied to detect deviation and trends according to the properties of computational market, which follows periodic type of behaviour. This can be the property of different prices due to the demand on daytime or weekdays. We evaluated workloads (Feitelson 2008) resolving that e.g. in a university network the demand for computational resources is significantly higher during official working hours than during night times or weekends. In our analysis, we will use the autoregressive moving average model (ARMA) and the seasonal autoregressive moving average model (ARIMA).

Autoregressive Moving Average Model ARMA (p) refers to the autoregressive model of order p . The *ARMA(p)* model is defined as $X_t = c + \sum_{i=1}^p \varphi_i X_{t-i} + \epsilon_t$, where $\varphi_1, \dots, \varphi_p$ are the *parameters* of the model, c is a constant and ϵ_t is white noise. An autoregressive model is essentially an all-pole infinite impulse response filter with some additional interpretation placed on it. Some constraints are necessary on the values of the parameters of this model in order that the model remains stationary. For example, processes in the *AR(1)* model with $|\varphi_1| \geq 1$ are not stationary.

3.1.3 Grid specific models

As we investigate computational market we will also apply models, which are successful in predicting the demand of computational resources. Sandholm and Lai (Sandholm and Lai 2008) showed that a prediction based on Chebyshev is better suited for predicting performance guarantees, which need worst case scenarios than perfect forecasting.

4 AN INTEGRATED MODEL – INFORMATION AND MARKETS

We propose an integrated architecture for market-based allocation for computational resources. The components have clear separated capabilities and interactions to achieve an economically efficient allocation of applications to needed computing resources. The model is composed by different independent systems (Borissov and Wirström 2008, Brunner et al 2008), which allow a clear separation of code, functionality and expert domain for an easier development, maintenance and fault-detection. The *BidGenerator*, the *Market Information System*, the *Trading Manager* and the *Resource Manager* are the main components, building an infrastructure for market-based allocation of computing services.

Real applications of the proposed model are listed and described in detail in (Nimis et al. 2008). The application component expresses batch jobs or real (web or desktop) applications, which has to be executed on demanded computing resources. Therefore, the consumer of the application submits a resource request to the *BidGenerator*. On the other side, the resource manager manages the resources of the provider and is responsible for the execution of the allocated applications. In order to offer a free resource, the resource manager submits a request to the *BidGenerator* for the resource provision.

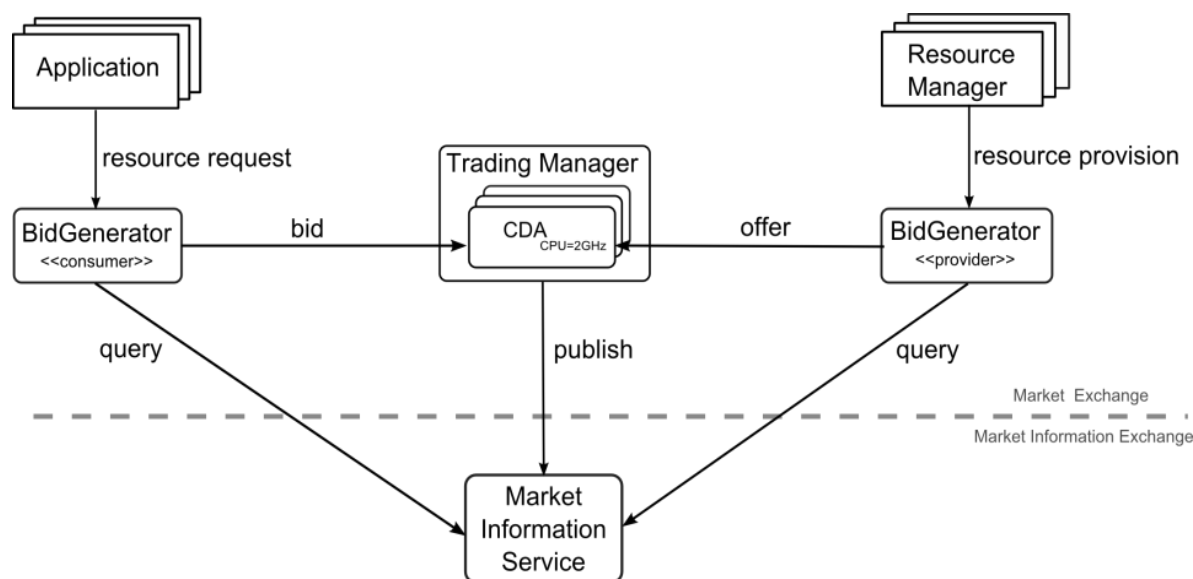


Figure 1. Integrated architecture of the related components.

The *BidGenerator* (Borissov and Wirström 2008) implements common and new bidding strategies (section), that autonomously generate and place the consumer bids and provider offers. The bids and offers are generated based on the consumer's preferences and the provider's business model, both initiated in the requests to the *BidGenerator*. A consumer specifies its preferences in form of a description of the required computing resources, duration of the job/application usage, valuation (maximum price) and preferred bidding strategy. A provider specifies its supply request in form of a resource description, duration of the provision, valuation (reserve price) and preferred bidding

strategy. Then the *BidGenerator* instantiates the selected bidding strategy and starts the bidding process.

The bids are submitted to the *Trading Management (TM)* component, which implements and runs market mechanisms (section 2.1) for technical and economic matching. The TM is a platform, which defines the interfaces and rules for implementing market mechanisms and the conversation protocol. In our case, the bids and offers (e.g. simplified example will be a CDA for CPUs with 2GHz), which are generated by the *BidGenerator*, are submitted to the TM via well defined interfaces of the selected market mechanism. When there is a match, the *BidGenerator* receives and informs the consumer or provider to execute its application on the allocated resource.

The *Market Information Service (MIS)* (Brunner et al. 2008) obtains economic data from the *Trading Manager* and provides it to the *BidGenerator*. The architecture has been designed to meet both the economic information requirements and that of scalability and robustness of distributed environment. Aggregation mechanisms are used to reach scalability in number of data and agent requests. Many of the introduced bidding strategies (Section 2.2) like ZIP, GD are exploiting prediction techniques (see Section 3), which require public market information. The MIS retrieves and aggregates public market information from the TM to provide them to the *BidGenerator*.

A short sample of products are commodity goods which are analysed commonly in markets, however, we seek to extend the model to more realistic goods for computing resources. We introduce a resource tuple with core description entities as baseline for representing and aggregating market information for computing resources.

Definition 1: We define a resource tuple $R = \{CPU, M, S, T_s, T_e, D\}$, where *CPU* is the processing power, *M* is the memory, *S* is the storage, T_s is the earliest start time, T_e is the latest end time and *D* is the duration of the job or resource reservation.

A demanded resource request can technically be allocated if each of the individual requirement is met e.g., minimum *CPU* of 2GHz, memory capacity of 1GB, storage of 10GB, etc. The same principle is used for the market information system, a bidding agent can search for the minimum, maximum or average price of *R*, where the individual requirements are met. This information will be used by the prediction mechanisms described in Section 3.1.

5 EVALUATION METHODOLOGY

To evaluate the integrated architecture in section 4, market mechanisms, bidding strategies and prediction techniques, we propose a common agent-based evaluation scenario and present first results.

5.1 Agent-Based Simulation

Table 1 summarizes the bidding strategies and mechanisms to be evaluated in varying job workload settings. The trading object is a computing resource as a commodity good or resource bundle. Within the simulation we aim to measure the consumer U_c and provider U_p utilityutilities and overall welfare adopting various bidding strategies (see 2.2 and 3) in three market mechanisms (see 2.1). Furthermore we aim to measure the quality Q of the forecast, deviation of real $R(t)$ to forecast value $F(t - 1)$, applying the prediction techniques (section 3) in novel bidding strategies, Prediction Traders.

In order to compare the strategies and market mechanism, we define 6 settings, Table 2, whereas each setting comprised from 20 until 200 providers and consumers. Settings 1 to 3 are generated with a Poisson arrival process, suggested by (Feitelson 2008), by increasing the mean λ , generating 1000, 3000 and 5000 jobs. Settings 4 to 5 are real workloads, taken from (Feitelson 2008).

Configuration	Market Mechanism	Trading Object	Job Data	Metrics
Agents Zero-Intelligence Plus Truth-Telling Gjerstad-Dickhaut QStrategy	CDA, DLGM, FIFO	Commodity good, Resource bundle - {CPU, Memory, Storage}	Generated and real workloads	Consumers utility $U_c = \sum_c \sum_j u_c(j)$, Providers utility $U_p = \sum_p \sum_i u_p(i)$, Overall welfare $U_w = U_c + U_p$, Forecast quality $Q = R(t) - F(t - 1)$
Prediction Traders (section 3)				

Table 1. Overall simulation scenario.

The workloads *LPC EGEE* (Medernach 2005), *LLNL Atlas* and *LLNL Thunder* are real cluster usage workloads, each containing 234,889, 42,725 and 121,039 jobs, occupying settings 4, 5 and 6. The LPC log was chosen as a basis because it contains a large variety of jobs with different run-times, numbers of used CPUs, and varying submit and start times. Finally, the valuation of a job and the reserve price of a resource are taken from a normal distribution.

Provider machines	10, 20, 100	
Consumers	10, 20, 100	
Job-Valuation	Normal distribution	
Reserve price	Normal distribution	
Generated workloads	Job arrival	Poisson process $P(\lambda)$
	Job duration	Normal distribution
	Setting 1 to 3	1000 to 5000 jobs
Real workloads	Setting 4 to 6	LPC EGEE, LLNL Atlas, LLNL Thunder

Table 2. Simulation settings.

5.2 Simulation Results

First results are shown in Table 3. The table shows the evaluation of consumer strategic behavior using Truth-Telling and Q-Strategy (2.2.3) in a centralized CDA (2.1.1), decentralized DLGM (2.1.2) and in a FIFO (2.1.3) mechanisms. The providers are acting strategically only in CDA, in DLGM and FIFO they are offering their resources for free.

Setting	Strategy	$DLGM_\mu$	$DLGM_\zeta$	CDA_μ	CDA_ζ	$FIFO_\mu$	$FIFO_\zeta$
1	Truth-Telling	-110	272	$-7,92 \cdot 10^{-4}$	95	-140	8
1	Q-Strategy	-175	258	$-10,33 \cdot 10^{-4}$	94	-135	8
2	Truth-Telling	-213	285	$-11,95 \cdot 10^{-4}$	94	-277	8
2	Q-Strategy	-392	266	$-14,63 \cdot 10^{-4}$	93	-265	8
3	Truth-Telling	-404	286	$-7,89 \cdot 10^{-4}$	87	-549	8
3	Q-Strategy	-901	265	$-23,22 \cdot 10^{-4}$	91	-525	8
4	Truth-Telling	-1104	647	$-9,91 \cdot 10^{-4}$	392	-4532	400
4	Q-Strategy	-1172	581	$-11,04 \cdot 10^{-4}$	314	-4444	400

Table 3. Consumer strategic behavior in three mechanisms.

Each line in the table represents the evaluation of a selected bidding strategy in a selected market for one setting. The workloads – setting 1-3 – are created with a Poisson distribution with means $\mu =$

{0.1, 0.3, 0.5} generating 751, 1502 and 3004 jobs. Setting 4 is the real cluster workload – LPC-EGEE, taken from a Parallel Workload Archive (Feitelson 2008). The first two columns represent the setting (corresponding to those of Table 1) and the evaluated strategy. The next two columns represent the average utility per job $\mu = \frac{-\sum_i v_j * \tilde{c}_{j,i} + \tilde{\pi}_{j,i}}{n}$ as well as the standard deviation ς of job budget and actual payment in DLGM and the last four columns represent the same for CDA and FIFO.

The results show that Truth-Telling strategy achieves the highest average utility for all settings in both DLGM and CDA, which follows that bidding truthfully in DLGM (2.1.2) and CDA (2.1.1) can only increase the consumer utility. Understating the truthful valuation in lower bid results in a poorer “job priority” by DLGM and thus job can be displaced by other jobs which have higher priority. In case of CDA, the prices depend on the current demand and supply, bidding a lower price instead of the truth valuation increases the risk of no allocation by the mechanism. However, in CDA, the matching is based only on the price without considering the “priority” of a job as with DLGM, and thus achieves very low utility compared to DLGM. Furthermore, CDA-provider machine agents do not maintain a priority queue of the submitted job bids and by an allocation the job is immediately submitted and executed on the provider machine. A provider submits an offer as soon as he becomes idle. Thus each time the agents are competing by adapting their job bids based on the used strategy. In overall, the DLGM mechanism outperforms FIFO in all four settings using the Truth-Telling bidding strategy and achieves the highest common wealth for all participating consumer agents.

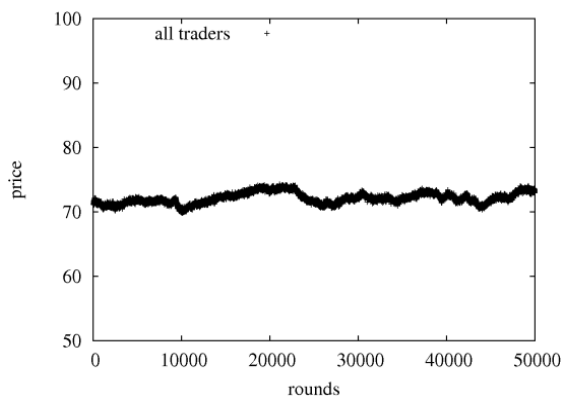


Figure 2. ZIP agents trading in a CDA with market information.

Figure 2 contains results from simulation of that market in which agents with a ZIP strategy are competing so sell resources as commodity goods. It depicts that the simulations run stable over a period of 50000 trading rounds with 2000 buyers and sellers.

6 RELATED WORK

The influence of information to market-based resource allocation mechanisms expands over different research areas. In Section 3, we explained related and potential market-based resource allocation methods and bidding strategies for these mechanisms. We summarize explicitly the information retrieval process and the individual prediction models in Section 4. This related work section describes work concerning the overall evaluation process of different prediction and learning mechanisms based on historical information.

Sandholm and Lai (Sandholm 2008) apply different mechanisms to predict the future demand of computational resources. They conclude to deduct the price of the resources from different real workload traces; however, markets with real market are not analyzed. The goal is to predict high peaks and though allow the consumers to avoid these to get lower prices or other benefits. This is especially important as sophisticated strategies change the markets behaviors. Moreover, different market mechanisms result in different peaks and distribution of the allocations. (Cardosa and Chandra 2007) analyze statistical aggregation for resource allocation. This information retrieval aggregates

historical data, which builds the basis for the prediction mechanisms. They provide further breakup of a commodity good like analyzed in most market-allocation mechanisms into resource bundles, however, economic mechanisms are not considered. A market-based scheduling mechanism for the allocation of time-specific resources is shown by (MacKie-Mason et al. 2004). It evaluates different prediction mechanisms to govern the time slot in competitive bidding process. They started with the combination of price prediction and bidding strategies to conclude that price prediction improves performance in a decentralized market-based scheduling environment.

Another survey and evaluation of market-based prediction mechanism is (Wellman et al. 2002), which analysis the prediction mechanisms of all participants of a trader competition. The evaluation compares price prediction mechanisms such as historical averaging, machine learning, and a competitive economy analysis in regard to their relative prediction accuracy. Wellman et al. conclude that a combination of historical data with current information produce the best prediction in a competitive analysis. In our work we will analyze and compare processes as entities in competition.

We can conclude from the presented related work that many mechanisms for the statistical prediction of prices exist for markets. Furthermore, different learning mechanisms and trading strategies are introduced, which are shown in previous sections. The related work, stated in this section, began with the comparison and analysis of these techniques. However, to our knowledge an overall approach combining the different areas of efficient information retrieval, statistical prediction, learning rules, trading strategies, and different market mechanisms is still an open issue.

7 CONCLUSIONS

In this paper, we proposed and analyzed alternative approaches to improve information flow within markets. Those approaches are measured with respect to the impact endogenous trading information has on the allocation quality. More precisely, we introduced promising market mechanisms, bidding strategies and prediction techniques, which have been evaluated and selected regarding their potential to optimize the resource allocation. Afterwards, we explained how these methods can be technically combined, integrated and evaluated within the allocation and information gathering process. For an easier and complete evaluation process, we developed a prototype for the proposed framework and showed baseline experiments with the proof-of-concept implementation. Positive integration results were obtained within an implemented market system for computing resource allocation.

As part of our analysis, we found that related work often proposes mechanisms in isolation for each topic, like the market mechanism, bidding strategy or information treatment, without regarding the complete interaction of components. Moreover, the focus of these mechanisms is not on the optimization of computational resource allocation. With the prototype, we show the integration of the approaches in a real framework, enabling autonomous provisioning and usage of computing resources. We identified evaluation scenarios and metrics for the impact of information in order to determine suitable combinations of market mechanisms, bidding strategies and related prediction techniques.

The future focus is on the evaluation of the proposed market mechanisms, bidding strategies and related prediction techniques from consumer to provider side. Applying the evaluation model with the proposed metrics will enable to measure the impact of real systems and identify optimal configurations of market mechanisms and bidding strategies with related prediction techniques to enable efficient market-based allocation of computing resources.

References

- Bagnall, A.J., Toft, I. (2005). Autonomous adaptive agents for single seller sealed bid auctions. *Journal of Autonomous Agents and Multi-Agent Systems*
- Bergemann, D., Valimaki, J. (2007). Information in mechanism design. 9th World Congress of the Econometric Society, volume Cambridge University Press 2007 of Chapter 5, pp. 186– 221.

- Borissov, N., Wirström, N. (2008). Q-Strategy: A Bidding Strategy for Market-Based Allocation of Grid Services, *On the Move to Meaningful Internet Systems: OTM 2008*, p. 744-761
- Brunner, R., Freitag, F., Navarro, L. (2008) Towards the Development of a Decentralized Market Information System: Requirements and Architecture. *Parallel and Distributed Computing in Finance (PDCoF'08)*. In *Proceedings of the 22nd IPDPS*, pp. 1-7.
- Cardosa, M., Chandra, A. (2008). Resource Bundles: Using Aggregation for Statistical Wide-Area Resource Discovery and Allocation. *28th Int. Conference on Distributed Computing Systems*, 760-768.
- Das, R., Hanson, J., Kephart, J., Tesauro, G. (2001). Agent-human interactions in the continuous double auction. *Proceedings of the International Joint Conference on Artificial Intelligence 26*
- Feitelson, D. (2008). *Workload Characterization and Modeling*. www.cs.huji.ac.il/labs/parallel/workload/
- Foster, I., Zhao, Y., Raicu, I., Lu, S. (2008). Cloud Computing and Grid Computing 360-Degree Compared, *Grid Computing Environments Workshop GCE'08*, p. 1-10
- Gjerstad, S., Dickhaut, J. (1998). Price formation in double auctions. *Games and Economic Behavior* 22(1)
- Grosu, D., Das, A (2006). Auctioning resources in Grids: model and protocols. *CONCURRENCY AND COMPUTATION*, 18(15)
- He, M., Leung, H., Jennings, N. (2003). A fuzzy-logic based bidding strategy for autonomous agents in continuous double auctions. *IEEE Transactions on Knowledge and Data Engineering* 15(6) 1345-1363
- Heydenreich, B., Mueller, R., Uetz, M. (2006). Decentralization and Mechanism Design for Online Machine Scheduling. *Algorithm Theory – SWAT 2006*, p. 136-147
- Lai, K., Rasmusson, L., Adar, E., Zhang, L., Huberman, B. (2005). Tycoon: An implementation of a distributed, marketbased resource allocation system. *Multiagent and Grid Systems* 1(3) p. 169-182
- Li, J., Yahyapour, R. (2006). Learning-based negotiation strategies for grid scheduling. *Proceedings of the Sixth IEEE International Symposium on Cluster Computing and the Grid*, v. 00, p. 576-583
- MacKie-Mason, J., Osepashvili, A., Reeves, D., Wellman, M. (2004). Price Prediction Strategies for Market-Based Scheduling. *18th International Conference on Automated Planning and Scheduling*.
- Medernach, E. (2005), *Workload Analysis of a Cluster in a Grid Environment*, *LECTURE NOTES IN COMPUTER SCIENCE*, volume 3834, Springer, p. 36- 61
- Nassif, L., Nogueira, J., de Andrade, F.: Distributed resource selection in grid using decision theory. *7th IEEE International Symposium on Cluster Computing and the Grid (2007)* p. 327-334
- Nimis et al. (2008). SORMA-Business Cases for an Open Grid Market: Concept and Implementation, *In proceedings of the 5th int. workshop on Grid Economics and Business Models*, p. 173-184
- Parkes, D., Singh, S., Yanovsky, D.: Approximately efficient online mechanism design. *Proc. 18th Annual Conf. on Neural Information Processing Systems (2004)*
- Reeves, D., Wellman, M., MacKie-Mason, J., Osepashvili, A. (2005). Exploring bidding strategies for market-based scheduling. *Decision Support Systems* 39(1)
- Sandholm, T., Lai, K. (2007). A statistical approach to risk mitigation in computational markets. *16th international Symposium on High Performance Distributed Computing*, ACM, NY, 85-96.
- Satterthwaite, M.A., Williams, S.R., (1989). Bilateral trade with the sealed bid k-double auction: Existence and efficiency *Journal of Economic Theory*, Volume 48, Issue 1, June, p. 107-133
- Tesauro, G., Das, R.: High-performance bidding agents for the continuous double auction. *Proceedings of the 3rd ACM conference on Electronic Commerce (2001)*
- Tesauro, G., Bredin, J.L. (2002). Strategic sequential bidding in auctions using dynamic programming, *First international joint conference on Autonomous agents and multiagent systems*, p. 2, p.591-598
- Vytelingum, P., Cliff, D., Jennings, N (2008). Strategic bidding in continuous double auctions. *AI*
- Watkins, C., Dayan, P.: Q-learning. *Machine Learning* 8(3) (1992), p. 279-292
- Wellman, M., Greenwald, A., Stone, P.: *Autonomous Bidding Agents: Strategies and Lessons from the Trading Agent Competition*. MIT Press (2007)
- Wellman, M P., Reeves, D. M., Lochner, K. M., Vorobeychik, Y. (2003). Price prediction in a trading agent competition. *Journal of Artificial Intelligence Research*, Volume 21, pp. 216-217.
- Wolski, R., Plank, J., Brevik, J., Bryan, T. (2001). Analyzing market-based resource allocation strategies for the computational grid. *Int. Journal of High Performance Computing Applications* 15(3)