On the Design of a Two-Tiered Grid Market Structure

Torsten Eymann[1], Dirk Neumann[2], Michael Reinicke[1],
Björn Schnizler[2], Werner Streitberger[1], Daniel Veit[2]

[1] University of Bayreuth, Chair for Information Systems,
Universitatsstrasse 30, 95440, Bayreuth, Germany
{eymann,reinicke,streitberger}@uni-bayreuth.de
http://wi.oec.uni-bayreuth.de
[2] University of Karlsruhe, Information Management and Systems,
Englerstrasse 14, D-76131 Karlsruhe, Germany
{neumann,schnizler,veit}@iw.uka.de
http://www.iw.uni-karlsruhe.de

Abstract Designing an end-to-end Grid market for computational utility with service requesters on the one side, and resource providers on the other, can become a complex issue. This paper motivates to break up into two different, though interdependent markets, in which both a market engineering process is carried out: a resource market to build basic servies out of computational and data resources, and a service market, which composes complex application services out of these building blocks.

The paper describes considerations and processes for designing dedicated market mechanisms for the resource management in Grids. The organization of these mechanisms can be either centralized (e.g. using an auctioneer) and decentralized (e.g. by bilateral bargaining processes).

Areas: Market Mechanisms, Computational Grid, Market Engineering, Catallaxy

1 Introduction

What makes economics so attractive for computing environments is that its central research question lies in the efficient allocation of resources, provided by suppliers and in demand by customers. Given a highly complex and dynamic Application Layer Network (ALN) infrastructure, applying economical principles to the resource management is seen as a promising approach, as they have the potential to achieve an efficient allocation [8].

ALNs comprise an abstract view on overlay networks (e.g. Peer-to-Peer networks, Grid infrastructures) on top of the TCP/IP protocol. Their common characteristic is the redundant, distributed provisioning and access of data, computation or application services, while hiding the heterogeneity of the service network from the user's view.

It appears that, by just implementing markets in computing environments, the satisfying ability of economics might be viable for creating cost-effective computer architectures. However, between the mostly descriptive economic concept and the normative technical implementation lies a fundamental realization gap. In essence, the gap reflects the problems in tailoring the market mechanism into a technical system, such that the peculiarities of the actors and resources are adequately addressed.

The paper at hand attempts to tackle this gap by designing dedicated market mechanisms for the resource management in ALNs. The organization of these mechanisms can be either centralized (e.g. using an auctioneer) and decentralized (e.g. by bilateral bargaining processes). The central auctioneer tends to be superior in terms of allocation and information efficiency [5] while the decentralized organization tends to be more advantageous regarding scalability issues. This trade-off is often stated, but not always confirmed [30,23]. Thus, we present design alternatives for ALNs for both types of mechanisms – the centralized and decentralized. Those mechanisms will be compared according to their economical and technical performance in future work.

In our approach, we follow the recommendations given by Hevner and coauthors [11], who elaborate on guidelines for design science in information systems. The key issue of our contribution is the design of the mechanisms we developed for the allocation of computational resources. Those mechanisms should be general enough such that they can be applied to any kind of ALNs.

This contribution is structured as follows: Section 2, analyses related approaches in market design for distributed systems. In section 3, the envisioned scenario is described. Consequently, in section 4 the general principle of engineering both, the centralized as well as the decentralized market mechanism for ALNs is described. The goal of section 4.1 is the outline of a possible centralized market mechanism for resource and service markets. This approach is opposed by a decentralized market mechanism for both markets in section 4.2. Finally, in section 5 conclusions are drawn from the design process and an outlook on the implementation of both market approaches in a large scale simulation is provided.

2 Related Approaches

The technical allocation of computational resources by maximizing a system specific objective function has been addressed as a research question in the past in several different areas: The main contributions stem from scheduling and technical resource allocation domain. Prominent examples for the definition of resource scheduling mechanisms for Computational Grids are [17,28]. However, in these approaches, the main focus is the technical feasible and local efficient allocation of the available resources. Less attention is dedicated to the economic efficiency.

The use of market mechanisms for allocating computer resources is not a completely new phenomenon. Essentially, the following short overview shows existing pricing market mechanisms, which are incorporated in the resource management system. The presented systems covers mainly auction technology (SPAWN, Bellagio, G-Commerce) or combinations of auction mechanisms with bargaining concepts (OCEAN). The SPAWN system provides a market mechanism for trading CPU times in a network of workstations [32]. SPAWN treats computer resources as standardized commodities and implements a standard Vickrey auction. However, the proposed approach is single-sided and favors monopolistic sellers or monopsonistic buyers in a way that allocates greater portions of the surplus. Installing competition on both sides is deemed superior, since no particular market side is systematically given an advantage.

Bellagio is intended to serve as a resource discovery and resource allocation system for distributed computing infrastructures. Users express preferences for resources using a bidding language, which support XOR bids. The bids are formulated in virtual currency. The auction employed in Bellagio is periodic. Bids from users are only accepted as long as enough virtual currency is left [2].

G-Commerce provides a framework for trading computer resources (CPU and hard disk) in commodity markets and Vickrey auctions [33]. While the Vickrey auction has the aforementioned shortcomings in grid, the commodity market typically works with standardized products. Additionally, the commodity market cannot account for the complementarities among the resources, as only one leg of the bundle is auctioned off, exposing the bidder to the threshold risk.

OCEAN (Open Computation Exchange and Arbitration Network) is a market-based infrastructure for high-performance computation, such as Cluster and Grid computing environments [1,19]. The major components of the OCEAN's market infrastructure are user components, computational resources, and the underlying market mechanism (e.g. the OCEAN Auction Component). The auction mechanism implemented in the OCEAN framework can be interpreted as a distributed sealed-bid continuous double-auction [1]. A trade is proposed to the highest bidder and the lowest seller. Afterwards, the trading partner can renegotiate their service level agreements. The renegotiation possibility one the one hand allows to cope with multiple attributes and with the assignment of resources to time slots.

With Application Layer Networks and Computational Grids entering the commercial sector, the economically efficient allocation of resources becomes increasingly important. Hence, several approaches have been made in the past few years applying market mechanisms in order to allocate computational resources. Cheliotis and co-authors [6] analyze in their contribution, which concepts from financial markets can be applied in resource allocation. They refer to ten lessons learnt from finance. One challenge is the price formation mechanisms that are "easy to implement but difficult to design". In principal, these market mechanisms can be designed using a centralized matching instance like an auctioneer or matchmaker (e.g. [31]) or a decentralized bilateral negotiation (e.g. [7]).

In [8] Eymann et al. describe the over-all setup of the two markets. Readers who wish an in depth insight into the technical concepts applied in this work may be referred to this article.

3 Preliminary Design

Before we tailor mechanisms for ALNs, it appears reasonably to reduce the complexity of the design problems. Referring to the guidelines of Hevner et al., design-science research often simplifies a problem by explicitly representing only a subset of the relevant concepts or by decomposing a problem into simpler sub-problems [11]. In our preliminary design phase we decompose the allocation problem of ALNs into two logically distinct, though interdependent markets.

This approach is demand promising as it reflects adequately the situation in ALNs, where participants offer and request application services, on the one hand, and computing resources of different complexity and value, on the other hand. End-users prefer to acquire services or service level agreements, which aggregate many different resources that are necessary to achieve the service or to maintain the service level. From the service provider perspective, it is also necessary to trade with resources. Either unused or idle resources can be contracted out on the market or resources necessary to fulfill services are bought. To simplify matters, the paper breaks these two complex, interdependent logical areas into to two different, though interdependent markets (c.f. figure 1) in which both a market engineering process is carried out (compare [18]):

(1) A resource market – which involves trading of computational and data resources, such as processors, memory, etc, and
(2) A service market – which involves trading of application services.

This artificial decomposition can be described best by referring to a scenario. Suppose there is a set of basic services (e.g. services to create a PDF or to convert a MP3 file), a set of complex services demanding these services for a specific job (e.g. an application wants to create a PDF file), and a set of resource services capable providing computational resources for executing these services (e.g. a processor, main memory, and a hard disk for creating the PDF file). The nice feature of our two-tiered market structure is that an agent requesting a service is unaware of the resources the requested service needs to be carried out.

Unlike the approach by [4], who introduce market and pricing mechanisms for computing capacities, in our approach demand is supposed to be certain. Bhargava and Sundaresan experiment with pricing and market mechanisms under the assumption of uncertain demand. In practice this means, that neither participants which request services to be completed nor, consequently, participants which demand computational resources will step back from contracts that have been signed once.

The distinction into two different markets has also other ramifications than reducing complexity. Services and resources are two different good types. Designing market mechanisms for both markets is aggravated by the fact that both markets have different properties. While resources are rather standardized but highly complex to handle (resource isolation problem referring to the constraints that some resources can only be coupled if they operate on the same computer), services are not standardized but less demanding in terms complexity. According to those properties it appears promising to employ a centralized mechanism for the standardized resource market and a decentralized mechanism for the not standardized service market.

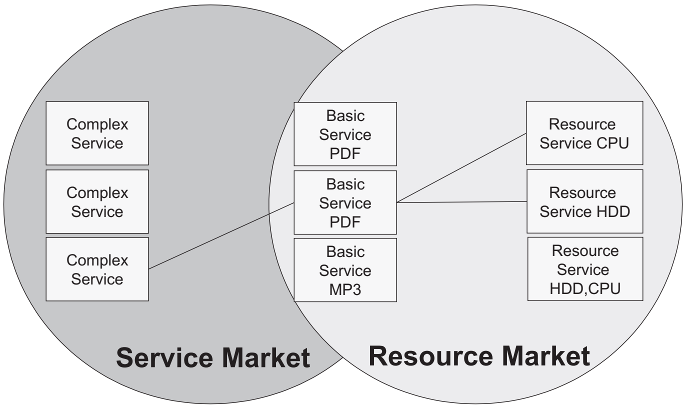

Figure 1. CATNETS Scenario: Service Market and Resource Market

4 Two Tiered-market Structure and Adequate Mechanisms

4.1 Centralized Service and Resource Markets

Formerly, auctions have been successfully applied to trade a variety of different commodities such as financial shares, electricity, or logistic scenarios. Auctions are institutions with an explicit set of rules determining resource allocation and prices on the basis of bids from the market participants [13]. As auctions can achieve economically efficient outcomes, their application for the CATNETS scenario is considered as a promising approach. In this subsection, an auction schema for the service market and another schema for the resource market are introduced.

Service Markets Designing a market mechanism for the service market firstly requires to analyse the underlying environment. Basically, buyers and sellers are services, which require other auxiliary services. That is, we distinguish basic services as sellers (e.g. a PDF creator service) and complex services as buyers (e.g. agents requesting a specific service). The basic services offer one or more specific auxiliary services. Hence, they are responsible for providing the auxiliary services to the buyers as well as for acquiring the required resources for the services on the resource market. Obviously, the products traded on the service market are completely standardized. For example, an instance of a PDF creator traded once does not differ from a PDF creator instance traded at a later time. For trading standardized resources, standard double auctions are successfully applied in theory and practice. Double auctions can achieve (approximate) efficient allocations and are computationally manageable in large-scaled markets.

In a double auction market, a large number of participants trade a common object and can submit bids (buy orders) and asks (sell orders). Trading in double auctions is organized by means of order books, each for a set of homogenous goods. An order book is responsible for storing non executed orders of the agents. For instance, in the service market there will be n different order books, each for one of the n different services. Buyers and sellers submit their bids in a sealed envelope to the auctioneer. The auctioneer aggregates the bids to form supply and demand curves. Once these curves are aggregated, they are used to set a specific price for trading – the price at which supply equals demand[1].

A key consideration in double auctions is the timing of the clearing process, i.e. the timing of determining the auction winners and thereby the allocation of the resources. Double auctions can be either cleared continuously (Continuous Double Auction) or periodically (Periodic Double Auction, Call Market): A Continuous Double Auction (CDA) is a double auction where buyers and sellers simultaneously and asynchronously announce bids and offers. Whenever a new order enters the market, the auctioneer tries to clear the market immediately. Thus, the CDA is advantageous especially in terms of immediacy. A Call Market is a double auction with periodic uniform clearing, e.g. the auctioneer clears the market every fives minutes. All orders in a period are collected in an order book and will be cleared periodically. Assuming none time-critical resources, the call market is advantageous in terms of enhancing the overall welfare in a market. A short time period may increase the overall welfare of a market; considering the immediate service allocation, a continuous clearing would be superior. The effects of both concepts have to be evaluated for the service market scenario by means of simulations.

Resource Markets In the resource market, participants are the basic services as resource consumers (buyers) and resource services (sellers) offering computational services having specific capacities, e.g. processing power. The same resources (e.g. CPUs) can differ in their quality attributes, e.g. a hard disk can have 30GB or 200GB of space. An adequate market mechanism for the resource market has to support simultaneous trading of multiple buyers and sellers, as well as an immediate resource allocation. Furthermore, the mechanism has to support bundle orders – i.e. all-or-nothing orders on multiple resources – as basic services usually demand a combination of computer resources. For comprising the different capacities of the resources (i.e. resources can differ in their quality), the mechanism has to support bids on multi-attribute resources.

Reviewing the requirements and surveying the literature, no classical auction mechanism is directly applicable to the resource market. Instead, a multi-attribute combinatorial exchange (MACE) is applied that satisfies the described requirements [25].

MACE allows multiple buyers and sellers simultaneously the submission of bids on heterogonous services expressing substitutabilities (realized by XOR bids) and complementarities (realized by bundle bids). Furthermore, the mechanism is capable of handling cardinal attributes as well as an immediate execution of given orders as the clearing can be done continuously. For instance, a resource consumer can bid on a bundle consisting of a computation service and a storage service. The computing service should have two processors where each processors should have at least 700MHz. Furthermore, the storage service should have 200MB of free space. The bids can be formulated as WS-Agreement offers [12] and thereby comply with standard resource negotiation mechanisms applied in current Grid systems. After the participants submitted

[1] Price tunnels (i.e. ranges where any price will be acceptable because the supply and demand curves overlap) are resolved using the k-pricing schema as presented in [9].

their bids to the auctioneer, the allocation (winner determination) and the corresponding prices are determined.

The objective of the winner determination problem in MACE is the maximization of social welfare, i.e. the difference between the buyers' valuations and the sellers' reservation prices. The problem is formulated as a linear Mixed Integer Programm (MIP) and thus can be solved by optimization solvers such as CPLEX[2]. The winner determination is, however, a generalization of the combinatorial allocation problem (CAP) and thus $\mathcal{NP}$ complete. For large-scaled scenarios, the use of approximations have to be evaluated [15]. Nevertheless, the application of such a complex problem seems to be promising, as the number of different bundles in the resource market is restricted.

The outcome of the winner determination model is allocative efficient, as long as buyers and sellers reveal their valuations truthfully. The incentive to set bids according to the valuation is induced by an efficient pricing mechanism. With respect to the economic objective of achieving an efficient allocation, a pricing scheme based on a Vickrey-Clarke-Groves (VCG) mechanism would attain this objective. Moreover, VCG mechanisms are the only allocative-efficient and incentive compatible mechanisms [10].

The basic idea of a VCG mechanism is to grant a participant a discount on its bids. This discount reflects the impact of that bid on the social welfare. A VCG mechanism is efficient, incentive-compatible, and individual rational for participants with quasi linear utility functions [20]. However, [16] proved that it is impossible to design an exchange, which is incentive compatible, (interim) individually rational, and budget balanced that achieves efficiency in equilibrium. In MACE, a VCG mechanism is efficient and individual rational, however, not budget balanced. In this case, the auctioneer has to endow the exchange, which is practical not realizable.

Relaxing the efficiency property, a possible implementation of a budget-balanced pricing rule for MACE is the k-pricing scheme. The underlying idea of the k-pricing scheme is to determine prices for a buyer and a seller on the basis of the difference between their bids [24]. For instance, suppose a buyer wants to buy a computation service for 5 and a seller wants to sell a computation service for at least 4. The difference between these bids is $\pi = 1$, i.e. π is the surplus of this transaction and can be distributed among the participants. This schema can be applied to MACE and results in an approximately efficient outcome [26].

4.2 Decentralized Service and Resource Markets

The computational complexity and the inability to obtain global knowledge in dynamic environments at one single instance are the main obstacles when being exposed to centralized control. In contrast, a decentralized market economy allows economic tradeoffs to be made by local decision makers, guided by price signals and constrained by general rules [14]. The absence of global knowledge leads to a self-organizing system with decentralized market-based control.

The market in the decentralized approach is mainly a communication bus where self-interested agents signal prices and barter for services and resources. First, a search process is processed, that sends the demand of a seller to possible buyer entities. In

[2] CPLEX is a commercial solver for optimization problems (`http://www.ilog.com/`).

a second step, the bargaining between one buyer and seller follows which determines the price for the good. Thereby, the local decision makers follow a strategy which decides about the price signals. A set of communicative acts in the negotiation protocol constraint rules for price signalling.

Search and selection of a possible negotiation partner The search methods at the markets are quite simple, because the search mechanism has to search only for the specified service or resource type. No additional constraints have to be taken into account. Therefore, the search is performed using widely accepted P2P search methods like simple flooding, CAN or CHORD [29,21].

After the termination of the search process, a list of suppliers is generated at the client's side which is expected to be incomplete but is assumed to be good enough for the following selection process. A complete list of all possible suppliers is impossible to achieve because of the system's dynamic behavior. Compared to the centralized market, this selection process fundamentally differs. The global resource broker of the centralized matchmaker now appears to be a local resource broker of each buyer entities (complex service on the service market and basic service on the resource market) which covers all offers of one demand. The local resource broker puts these offers into an order, using a utility function. The negotiation partner with the expected highest utility is chosen for the following bargaining process.

Bargaining and Pricing The definition of a strategy for bartering is essential in the decentralized CATNETS market. The initial situation is described like depicted in figure 2. A (human) principal defines an indifference price that equals his estimation about the value of the good. For a buyer, this is a maximum price, for the seller a minimum price. So, the utility gain equals the amount between price of the purchase and the indifference price. The start price represents the price where the strategy begins to negotiate. By agreeing concessions, the opponents come closer to the middle and a possible contract. A transaction is unlikely, if the closure zone is empty, which might result when indifference prices do not build an overlapping zone.

The goal of the software agent strategies is maximization of their own utility. For a buyer, this means an maximization of the distance between the indifference and purchase price and for the seller the enlargement of the interval between the seller's purchase price and his indifference price. The proposed realization for the CATNETS market is the usage of a heuristic factor, that decides on the percentage change of the negotiation's starting price in the following, successful completed transactions. The higher this parameter is, the lager is the aspiration towards the enlargement of price distances. The agent is not aware of his reached level and there is no optimal goal. The attempt to maximize his own utility is not limited by any restriction regarding the absolute value nor time constraints and is a never ending process.

The continuous aspiration, to be better than other agents, is also influenced by the selected learning algorithm. The search for a good parameter configuration which assures good utility gain, leads to comparison with and re-combination of known configurations of other agents. This results in the decentralized, evolutionary learning algorithm STDEA [27].

Figure 2. Bilateral negotiation process.

The local decision makers follow a strict negotiation protocol for exchanging their price signals in a one-to-one communication. Both agents converge to a tradeoff point in an iterative way using the exchange of offers and counter-offers. Rosenschein and Zlotkin [22] call this a monotonic concession protocol. In the CATNETS project, the negotiation protocol has basically three different states, where decision about the next communicative act has to be made: accept the offer, propose a counter-offer or reject the offer.

The strategy chooses at every negotiation step one of these three states of the negotiation protocol until an agreement is reached or the negotiation is aborted. This specifies the actions, an agent uses for attainability of his goals. He chooses the best alternative action for his goal on the basis of information about the environment. The negotiation strategy, which is described here, is based on the AVALANCHE strategy [7]. The strategy defines 5 basic parameters, which define the individual behavior of each agent:

- priceNext:
  The parameter priceNext modifies the particular starting value of a new negotiation. The result of the last negotiation is saved and it is esteemed as quite sure, to achieve the almost the same value as in the last negotiation. An achievement of the last would not increase the utility of the agent. Therefore, he tries to buy cheaper the next time and decreases his initial price with a winner discount.
- weightMemory:
  The weighting ratio of current price information and historic price information influences the strategy of the agents largely. The higher the weighting factor is, the faster the market price adapts to the current market situation. But, the agents can be influenced faster through short-time price fluctuation on the market. This will emphasize the characteristic behavior of the agent. A "correct" value of the weighting factor cannot be defined a priori. The value arises from the cooperation of the agents. The parameter is predefined at the initialization of the agent.
- satisfaction:

If the offer price is higher than the market price, but less than double market price, an agent has to decide whether to continue. It is possible that the market price of the negotiated good increased his price. To simulate a more complex deterministic "probe", the stochastic factor satisfaction is checked which has values between 0 and 1. A satisfaction value of 0.75 means, that the agent is satisfied to 75% with the negotiation process and continues. An agent with satisfaction 0 will abort a negotiation at once and an agent with satisfaction 1 will never abort a negotiation.

– acquisitiveness:

  The parameter acquisitiveness defines the probability to make an uniliteral concession in the following negotiation. The real decision is determined with a stochastic "probe". The value of the parameter does not determine the action, but sets a specific probability. The value interval is between 0 and 1. A value of 0.7 means a probability of 70%, that an agent follows a competitive strategy not making concessions. A seller agent with a true boolean value will not adapt his price and will signal the same price to his opponent. If the boolean value is false, the agent computes an new price, adapting the price using his concession parameter. A buyer agent will rise his offer, a seller agent will lower his price. An agent with acquisitiveness value 1 will never change his price and an agent with acquisitiveness value 0 will make an unilateral concession, adapting his price towards the opponents price.

– priceStep:

  The concession's level (parameter priceStep) is defined at the beginning of the negotiation. The definition of an absolute concession level does not take into account the level of the demanded price which leads to an implementation with a percentage computation. The direct indication of the concession as a percentage of the bid discriminates the seller, because he assumes an absolute value during the negotiation. With such a process, the market price decreases over time. That would lead to a wrong picture of the market-based coordination mechanism. Thus, a percentage of the difference between the initial starting prices of both parties is introduced in the strategy which does not change during the negotiation. A value of priceStep = 0.25 means a computation of the concession level as 1/4 of the first stated difference. If both opponents negotiate in the same way and make concession to each other, they meet each other on the half way in the third negotiation round under the assumption of no negotiation abortion.

This heuristic strategy is used to realize market-based control on the decentralized CATNETS market. The 5 parameters influence the price signalling of each agent and influence the outcome of the negotiations. The combination of historic data of market prices on the one hand and stochastic probes on the other hand ensure the adaptability to the dynamic environment. Additionally the strategy is supported using the Smith Taylor Distributed Evolutionary learning Algorithm (STDEA) for faster adoption of the local decision makers to new market circumstances.

5 Conclusions & Outlook

This contribution provides a selection of market mechanisms for application and computational services for both, a centralized and a decentralized approach. These findings are considered to be a first step towards the evaluation of a comparison between these two fundamentally different paradigms.

In both cases, an end-to-end – complex service to resource service – perspective is considered. Under real market circumstances, a complex service which represents a customer will not be able to judge the computational demand of his job. However, the intermediate basic services or service providers offer the service to the complex services and purchase at the same time the computational resources at the resource services.

In the approach at hand, in the centralized case, the service market is designed as a double auction market. The resource market is carried out as a combinatorial exchange with a k-pricing approach. In the decentralized case, both markets are carried out with an approach of bilateral bargaining using heterogeneous strategies.

The next steps in our work comprise the evaluation of both approaches. This issue is approached twofold: In one track, a large scaled simulator based on OptorSim [3]. In this, both, the centralized and the decentralized market mechanisms will be implemented. The outcomes will be compared using different economic metrics.

In the second track, a prototype implementation based on WS-agreement protocol is carried out, which fosters the decentralized approach for bilateral bargaining on computational resources.

Acknowledgements

This work was partly funded by the European Commission under contract FP6-003769.

References

1. ACHARYA, N., CHOKKAREDD, C., DESAI, P., DEVRAJAN, R., FRANK, M. P., CHAKRAVARTHUL, S. K. N., NAGENDRANATH, M., PADALA, P., PARK, H., SUR, G., TOBIAS, M., WONG, C.-K., AND YU, B. The Open Computation Exchange & Auctioning Network (OCEAN). Tech. rep., Department of Computer & Information Science & Engineering, University of Florida, 2001.
2. AUYOUNG, A., CHUN, B. N., SNOEREN, A. C., AND VAHDAT, A. Resource Allocation in Federated Distributed Computing Infrastructures. In *Proceedings of the 1st Workshop on Operating System and Architectural Support for the On-demand IT InfraStructure* (2004).
3. BELL, W. H., CAMERON, D. G., MILLAR, A. P., CAPOZZA, L., STOCKINGER, K., AND ZINI, F. OptorSim: A Grid Simulator for Studying Dynamic Data Replication Strategies. *International Journal of High Performance Computing Applications 17* (2003), 403–416.
4. BHARGAVA, H., AND SUNDARESAN, S. Computing as Utility: Managing Availability, Commitment, and Pricing Through Contingent Bid Auctions. *Journal of Management Information Systems 21*, 2 (2004), 201–227.
5. BULOW, J., AND KLEMPERER, P. Auctions versus negotiations. *American Economic Review 86* (1996), 80–194.

6. CHELIOTIS, G., KENYON, C., AND BUYYA, R. Grid Economics: 10 Lessons from Finance. Joint Technical Report, GRIDS-TR-2003-3, IBM Research Zurich and Grid Computing and Distributed Systems Laboratory and University of Melbourne, 2003.
7. EYMANN, T. *AVALANCHE - An Agent-Based Decentralized Coordination Mechanism For Electronic Marketplaces*. PhD thesis, Wirtschaftswissenschaftliche Fakultaet Albert-Ludwigs-Universität Freiburg, 2000.
8. EYMANN, T., REINICKE, M., STREITBERGER, W., RANA, O., LIVIU, J., NEUMANN, D., SCHNIZLER, B., VEIT, D., ARDAIZ, O., CHACIN, P., CHAO, I., FREITAG, F., NAVARRO, LEANDRO CATALANO, M., GALLEGATI, M., GIULIONI, G., CARVAJAL SCHIAFFINO, R., AND ZINI, F. Catallaxy-based grid markets. *Multiagent and Grid Systems, Special Issue on Smart Grid Technologies and Market Mechanisms* (2005). forthcoming.
9. FRIEDMAN, D. The Double Auction Market Institution: A Survey. In *The Double Auction Market - Institutions, Theories, and Evidence*, D. Friedman and J. Rust, Ed. Cambridge MA, Perseus Publishing, 1991, pp. 3–26.
10. GREEN, J., AND LAFFONT, J.-J. Characterization of Satisfactory Mechanisms for the Revelation of Preferences for Public Goods. *Econometrica 45*, 2 (1977), 427–438.
11. HEVNER, A. R., MARCH, S. T., PARK, J., AND RAM, S. Design Science in Information Systems. *MIS Quarterly 28*, 1 (March 2004), 75–105.
12. LUDWIG, H., DAN, A., AND KEARNEY, B. Cremona: An Architecture and Library for Creation and Monitoring of WS-Agreements. In *Proceedings of the 2nd International Conference on Service Oriented Computing (ICSOC 2004)* (2004), pp. 65–74.
13. MCAFEE, R., AND MCMILLAN, J. Auctions and Bidding. *Journal of Economic Literatur 25*, 2 (1987), 699–738.
14. MILLER, M. S., AND DREXLER, K. E. *The ecology of computation*. Elsevier Science Inc., New York, NY, USA, 1988.
15. MITO, M., AND FUJITA, S. On Heuristics for Solving Winner Determination Problem in Combinatorial Auctions. *Journal of Heuristics 10*, 5 (2004), 507–523.
16. MYERSON, R. B., AND SATTERTHWAITE, M. A. Efficient Mechanisms for bilateral Trading. *Journal of Economic Theory 28* (1983), 265–281.
17. NAKAI, J., AND WIJNGAART, R. F. V. D. Applicability of Markets to Global Scheduling in Grids. NAS Report 03-004, NASA Advanced Supercomputing (NAS) Division, 2003.
18. NEUMANN, D. *Market Engineering - A Structured Design Process for Electronic Markets*. PhD thesis, Economics and Business Engineering. Karlsruhe, Germany, University of Karlsruhe (TH), 2004.
19. PADALA, P., HARRISON, C., PELFORT, N., JANSEN, E., FRANK, M., AND CHOKKAREDDY, C. Ocean: The open computation exchange and arbitration network, a market approach to meta computing, 2003.
20. PARKES, D. C., KALAGNANAM, J., AND ESO, M. Achieving Budget-Balance with Vickrey-Based Payment Schemes in Exchanges. In *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence* (2001), pp. 1161–1168.
21. RATNASAMY, S., FRANCIS, P., HANDLEY, M., KARP, R., AND SHENKER, S. A Scalable Content Addressable Network. Tech. Rep. TR-00-010, Berkeley, CA, 2000.
22. ROSENSCHEIN, J. S., AND ZLOTKIN, G. *Rules of Encounter: Designing Conventions for Automated Negotiation Among Computers*. MIT Press, Cambridge, Massachusetts, 1994.
23. SATTERTHWAITE, M., AND SHNEYEROV, A. Convergence of a dynamic matching and bargaining market with two-sided incomplete information to perfect competition. Tech. rep., 2003.
24. SATTHERTHWAITE, M. A., AND WILLIAMS, S. R. The Bayesian Theory of the k-Double Auction. In *The Double Auction Market - Institutions, Theories, and Evidence*, D. Friedman and J.Rust, Ed. Addison-Wesley, 1993, ch. 4, pp. 99–123.

25. Schnizler, B., Neumann, D., Veit, D., and Weinhardt, C. A Multiattribute Combinatorial Exchange for Trading Grid Resources. In *Proceedings of the Research Symposium on Emerging Electronic Markets (RSEEM), Amsterdam, Netherlands* (2005).
26. Schnizler, B., Neumann, D., Veit, D., and Weinhardt, C. Trading Grid Services – A Multi-attribute Combinatorial Approach. Working Paper, University of Karlsruhe (TH), Germany, 2005.
27. Smith, R. E., and N. Taylor, N. A Framework for Evolutionary Computation in Agent-Based Systems. In *Proceedings of the 1998 International Conference on Intelligent Systems* (1998), ISCA Press.
28. Smith, W., I.Foster, and Taylor, V. Scheduling with Advanced Reservations. In *Proceedings of the International Parallel and Distributed Processing Symposium (IPDPS) Conference* (2000).
29. Stoica, I., Morris, R., Karger, D., Kaashoek, F., and Balakrishnan, H. Chord: A Scalable Peer-To-Peer Lookup Service for Internet Applications. In *Proceedings of the 2001 ACM SIGCOMM Conference* (2001), pp. 149–160.
30. Thomas, C. J., and Wilson, B. J. A comparison of auctions and multilateral negotiations. *RAND Journal of Economics 33* (2002), 140–155.
31. Veit, D. *Matchmaking in Electronic Markets*. Springer Verlag, 2003.
32. Waldspurger, C. A., Hogg, T., Huberman, B., Kephart, J. O., and Stornetta, W. Spawn: A Distributed Computational Economy. *IEEE Transactions on Software Engineering 18*, 2 (1992), 103–117.
33. Wolski, R., Plank, J., Brevik, J., and Bryan, T. G-Commerce: Market Formulations Controlling Resource Allocation on the Computational Grid. In *Proceedings of the International Parallel and Distributed Processing Symposium (IPDPS)* (2001).