

# Combinatorial Exchanges for Coordinating Grid Services

BJÖRN SCHNIZLER

and

DIRK NEUMANN

Institute of Information Systems and Management, University Of Karlsruhe, Germany

---

In this paper, two combinatorial exchange mechanisms MACE and GreedEx are introduced that are suitable for resource allocation in service-oriented environments such as Grids. MACE provides users with a fairly complex bidding language offering flexibility in the bidding process. This flexibility comes at the expense of computational tractability. GreedEx tremendously restricts the bidding language and thereby reduces the problem complexity considerably. As with any practical mechanism design effort, the designed artifact does not implement desirable allocations in dominant strategies. In this short paper we also introduced jCase as tool for evaluating the market outcome when agents are acting strategically on the market for Grids.

Categories and Subject Descriptors: K.6 MANAGEMENT OF COMPUTING AND INFORMATION SYSTEMS

General Terms: Integer programming; Combinatorial exchange; Grid

Additional Key Words and Phrases: Auctions and bidding; Market engineering

---

This research was supported by the Funding Agency. This work has partially been funded by EU IST programme under grant 034286 "SORMA".

Authors' addresses: [bjoern.schnizler@googlemail.com](mailto:bjoern.schnizler@googlemail.com) (Björn Schnizler), [dirk.neumann@kit.edu](mailto:dirk.neumann@kit.edu) (Dirk Neumann)

Permission to make digital/hard copy of part of this work for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication, and its date of appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee.

© 2001 ACM 1073-0516/01/0300-0034 \$5.00

ACM SIGecom Exchanges, Vol. 7, No. 1, December 2007.

## 1. INTRODUCTION

Since the early 1990s, the integration of computers and other computational resources into parallel and distributed systems has become common practice. In this context, the term Grid denotes a computing model, where computational resources (e.g. CPUs, applications) are managed as cohesive distributed systems. State-of-the-art Grid middleware provides the technical infrastructure to allow the sharing of resources over multiple geographic and administrative domains. Local resource managers are implemented for allocating and scheduling computational jobs. These managers typically employ simple objective functions (e.g. maximize utilization) or first-come first-serve strategies. Generally, these schedulers work well as long as information about supply and demand is truthfully reported. Since Grids tackle resource sharing across organizations, those scheduling mechanisms suffer under insincere revelation of job information (e.g. priority level) [Neumann et al. 2006; Schnizler et al. 2006]. Consequently, these mechanisms are inapplicable in most commercial Grid settings.

Market mechanisms are known to attain fairly efficient allocations in situations where the participating agents may conceal their private information about costs and valuations [Shneidman et al. 2005]. Nevertheless, most available market mechanisms are not applicable to Grids. Reasons for this stem from the fact that demand for Grid resources refers to bundles, i.e. combinations of multiple resources. This demand arises from the fact that Grid resources display complementarities [Subramoniam et al. 2002]. Participants have super-additive valuations for the resources, as the sum of the valuations for the single resources is less than the valuation for the whole bundle. This is quite logical as jobs may need both computation and storage to be executable. Additionally, computer resources can have different quality attributes (e.g. 10 MB storage) and are time dependent (e.g. from 7 – 9 o'clock).

In the following sections, we present two combinatorial exchanges that attempt to solve the resource allocation problem in Grids [Neumann et al. 2008]. The first mechanism, called MACE, allows a rich bidding language that addresses most general requirements upon a Grid scheduler. As MACE is hard to solve in large-scaled settings we present a second mechanism called GreedEx. GreedEx implements a heuristic that attains approximately truthful revelation of private information. For evaluating the mechanisms, we propose using the open-source toolkit Java Combinatorial Auction Simulation Environment (jCase).

## 2. MACE: A MULTI-ATTRIBUTE COMBINATORIAL EXCHANGE

MACE is a multi-attribute combinatorial exchange for allocating and scheduling resources in the Grid [Schnizler et al. 2006; Schnizler 2007]. In contrast to other approaches, the proposed mechanism accounts for the variety of Grid resources by incorporating time and quality as well as co-allocation constraints. The mechanism provides buyers and sellers with a rich bidding language, allowing them to compose bundles that express either substitutabilities or complementarities. The winner determination problem maximizes social welfare for the submitted bids. The winner determination scheme alone, however, is insufficient to guarantee an efficient allocation of the services. The pricing scheme must be constructed in a way that motivates buyers and sellers to reveal their true valuations and reservation prices. This is problematic in the case of combinatorial exchanges, since the only pricing schema that assures an efficient allocation of resources, the VCG mechanism, is not budget-balanced and must be subsidized from outside the mechanism. Although an approximation of the VCG mechanism results in budget-balanced results [Parkes et al. 2001], the computational effort that is required to compute the payments is high. MACE implements a k-pricing rule. In essence, the k-pricing rule determines prices such that the resulting surpluses to the buyers and sellers divide the entire surplus being accrued by the trade according to a predefined ratio k. The pricing rule is budget-balanced but cannot retain the efficiency property. However, simulations showed that the pricing schema results come very close to efficient outcomes.

The contributions of MACE are the following: It is the first auction mechanism that addresses simultaneously several Grid specific requirements such as quality characteristics, time attributes and co-allocation restrictions. MACE allows agents to bid on bundles rather than on single items. This leads to more efficient outcomes compared to traditional auctions, as Grid resources are complementary. The mechanism applies a new pricing schema for combinatorial exchanges. In contrast to VCG payments, the proposed k-pricing rule is feasible and computationally more efficient.

## 3. GREEDEX: A SCALABLE CLEARING MECHANISM FOR UTILITY COMPUTING

MACE allows an enormous degree of freedom in calculating efficient allocations, which comes at the expense of tractability. The participation of 500 participants already requires computation time of more than 5 minutes. This rather long computation time gave rise to the design of GreedEx. In essence, GreedEx retains the property of an exchange, where multiple requesters and providers can trade for resources for a sequence of time slots. Different

from MACE, GreedEx allows only trading on two goods, computation power and memory. Stoesser et al. design a scheduling heuristic that achieves truthful revelation of valuations in dominant strategies on the demand side by relying on the Vickrey principle [Stoesser and Neumann 2007; Stoesser et al. 2007]. Furthermore, the mechanism attains an approximately truthful revelation on the supply side by distributing the generated surplus among the resource providers according to their contribution to the allocation. The heuristic, however, is unbounded as the allocation can be made arbitrarily bad. This shortcoming can be alleviated by introducing a randomization of the mechanism. With this extension, GreedEx constitutes a very fast scheduling mechanism, which forces an approximately truthful revelation of valuations by the requesters and providers. In addition to that, GreedEx is split- as well as merge-proof, as requesters cannot benefit by dividing large jobs into smaller ones or by combining several jobs into a larger one. The mechanism has been successfully implemented in the Grid Operating System MOSIX and will be tested in a field study at the Campus of the Hebrew University at Jerusalem [Amar et al. 2007].

#### 4. JCASE: JAVA COMBINATORIAL AUCTION SIMULATION ENVIRONMENT

The evaluation of both mechanisms is performed by means of the Java Combinatorial Auction Simulation Environment (jCase) [Schnizler 2007], an open-source toolkit for simulating combinatorial mechanisms<sup>1</sup>. jCase is capable of generating different bidding streams with a varying number of participants, resources and bundles. The simulation tool implements, among others, the different bidding techniques for generating stochastically influenced bids. It integrates the CATS 2.0 framework [Leyton-Brown and Shoham 2006] and provides a set of different bundle distributions. The bidding strategy of the agent population can be varied from truth-telling to manipulation. Depending on the users' choice, the allocation and the system performance (e.g. revenue of the sellers, allocative efficiency) is computed.

Simulation settings are described using an XML based description language. Once a scenario is encoded by means of this specification, it can be executed using the graphical user interface or through a batch process.

The evaluation of MACE and GreedEx discovered that both mechanisms impose a strong incentive to truthfully reveal the drawn valuation. Both mechanisms lead to allocations not far from the efficient allocation. This results holds for GreedEx even more than for MACE, as the demand side implements truth-telling in dominant strategies.

#### 5. CONCLUSION AND FUTURE TRENDS

In this paper, we briefly introduced the two combinatorial exchange mechanisms MACE and GreedEx that could be used for resource allocation in service-oriented environments such as Grids. MACE provides the user with a very rich bidding language, which allows trading multiple services over the Grid. This flexibility comes at the expense of tractability. GreedEx tremendously restricts the bidding language and thereby reduces the problem complexity considerably. As with any practical mechanism design effort, the designed artifact does not implement desirable allocations in dominant strategies. In this short paper we also introduced jCase as tool for evaluating the market outcome when agents are acting strategically on the market for Grids.

So far, both mechanisms MACE as well as GreedEx have been tested by using synthetically produced data. Whether those mechanisms will succeed in practice remains an open research question. Empirical data on the job characteristics, the number of participants and their valuations are needed to answer this. Both mechanisms will be applied in different field studies. The deployment of MACE will take place in the project Biz2Grid2, where the mechanism will be integrated into a commercial scenario. GreedEx will be launched within a Campus Grid project as reward mechanism within the EU project SORMA3.

#### REFERENCES

- AMAR, L., STOESSER, J., BARAK, A., AND NEUMANN, D. 2007. Economically Enhanced MOSIX for Market-based Scheduling in Grid OS. Workshop on Economic Models and Algorithms for Grid Systems (EMAGS) 2007, Austin, USA.
- LEYTON-BROWN, K. AND SHOHAM, Y. 2006. A test suite for combinatorial auctions. *Combinatorial Auctions*. P. Cramton, Y. Shoham and R. Steinberg, MIT Press: 451-478.
- NEUMANN, D., HOLTSMANN, C., AND ORWAT, C. 2006. Grid Economics. *Wirtschaftsinformatik* 48(3): 206-209.
- NEUMANN, D., STOESSER, J., AND WEINHARDT, C. 2008. Bridging the Adoption Gap - Developing a Roadmap for Trading Grids. *International Journal of Electronic Markets* 18(1): forthcoming.

<sup>1</sup> See <http://www.iw.uni-karlsruhe.de/jcase/> for details.

<sup>2</sup> See <http://www.biz2gro.de/> for details.

<sup>3</sup> See <http://www.sorma-project.eu/> for details.

- PARKES, D. C., KALAGNANAM, J., AND ESO, M. 2001. Achieving budget-balance with vickrey-based payment schemes in exchanges. International Joint Conference on Artificial Intelligence.
- SCHNIZLER, B. 2007. Resource Allocation in the Grid - A Market Engineering Approach. Karlsruhe, Universitätsverlag Karlsruhe.
- SCHNIZLER, B., NEUMANN, D., VEIT D., AND WEINHARDT, C. 2006. Trading Grid Services - A Multi-attribute Combinatorial Approach. *European Journal of Operation Research*: forthcoming.
- SHNEIDMAN, J., NG, C., PARKES, D.C., AU YOUNG, A., SNOEREN, A.C., VAHDAT, A., AND CHUN, B. 2005. Why Markets Could (But Don't Currently) Solve Resource Allocation Problems in Systems. In *Proceedings of the 10th conference on Hot Topics in Operating Systems*.
- STOESSER, J. AND NEUMANN D. 2007. GREEDEX – A Scalable Clearing Mechanism for Utility Computing. Networking and Electronic Commerce Research Conference (NAEC), Riva Del Garda, Italy.
- STOESSER, J., NEUMANN D., AND ANANDASIVAM A. 2007. A Truthful Heuristic for Efficient Scheduling in Network-Centric Grid OS. European Conference on Information Systems (ECIS 2007), St. Gallen.
- SUBRAMONIAM, K., MAHESWARAN M., AND TOULOUSE M. 2002. Towards a micro-economic model for resource allocation in grid computing systems. In *Proceedings of the 2002 IEEE Canadian Conference on Electrical & Computer Engineering*.