

Peyman Faratin

Juan A. Rodríguez-Aguilar (Eds.)

LNAI 3435

# Agent-Mediated Electronic Commerce VI

Theories for and Engineering  
of Distributed Mechanisms and Systems

AAMAS 2004 Workshop, AMEC 2004

New York, NY, USA, July 2004

Revised Papers



Springer

F713.36-53

A265

2004

Peyman Faratin

Juan A. Rodríguez-Aguilar (Eds.)

# Agent-Mediated Electronic Commerce VI

Theories for and Engineering  
of Distributed Mechanisms and Systems

AAMAS 2004 Workshop, AMEC 2004  
New York, NY, USA, July 19, 2004  
Revised Selected Papers



E200603436



Springer

## Series Editors

Jaime G. Carbonell, Carnegie Mellon University, Pittsburgh, PA, USA

Jörg Siekmann, University of Saarland, Saarbrücken, Germany

## Volume Editors

Peyman Faratin

Massachusetts Institute of Technology

Computer Science and Artificial Intelligence Laboratory

Massachusetts Avenue, Cambridge, 02139, USA

E-mail: peyman@mit.edu

Juan A. Rodríguez-Aguilar

Artificial Intelligence Research Institute (IIIA)

Spanish Council for Scientific Research (CSIC)

Campus de la Universitat Autònoma de Barcelona, 08193 Bellaterra, Spain

E-mail: jar@iiia.csic.es

Library of Congress Control Number: 2005938890

CR Subject Classification (1998): I.2.11, K.4.4, C.2, H.3.4-5, H.5.3, I.2, J.1

LNCS Sublibrary: SL 7 – Artificial Intelligence

ISSN 0302-9743

ISBN-10 3-540-29737-5 Springer Berlin Heidelberg New York

ISBN-13 978-3-540-29737-6 Springer Berlin Heidelberg New York

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, re-use of illustrations, recitation, broadcasting, reproduction on microfilms or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer. Violations are liable to prosecution under the German Copyright Law.

Springer is a part of Springer Science+Business Media

springer.com

© Springer-Verlag Berlin Heidelberg 2005

Printed in Germany

Typesetting: Camera-ready by author, data conversion by Scientific Publishing Services, Chennai, India

Printed on acid-free paper SPIN: 11575726 06/3142 5 4 3 2 1 0



## Preface

The design of intelligent trading agents, mechanisms, and systems has received growing attention in the agents and multiagent systems communities in an effort to address the increasing costs of search, transaction, and coordination which follows from the increasing number of Internet-enabled distributed electronic markets. Furthermore, new technologies and supporting business models are resulting in a growing volume of open and horizontally integrated markets for trading of an increasingly diverse set of goods and services. However, growth of technologies for such markets requires innovative solutions to a diverse set of existing and novel technical problems which we are only beginning to understand. Specifically, distributed markets present not only traditional economic problems but also introduce novel and challenging computational issues that are not represented in the classic economic solution concepts. Novel to agent-mediated electronic commerce are considerations involving the computation substrates of the agents and the electronic institutions that supports trading, and also the human-agent interface (involving issues of preference elicitation, representation, reasoning, and trust). In sum, agent-mediated electronic trade requires principled design (from economics and game theory) and incorporates novel combinations of theories from different disciplines such as computer science, operations research, artificial intelligence, and distributed systems.

The collection of above-mentioned issues and challenges has crystallized into a new, consolidated agent research field that has become a focus of attention in recent years: *agent-mediated electronic commerce*.

The papers in this volume originate from the 6th Workshop on Agent-Mediated Electronic Commerce (AMEC VI), held in conjunction with the 3rd International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS) in July 2004. The AMEC VI workshop continued with the tradition and built upon the success of the previous AMEC workshops.

Thus, the primary goal of this workshop was to continue to bring together novel work from diverse fields such as computer science, operations research, artificial intelligence and distributed systems that focus on modeling, implementation, and evaluation of computational trading institutions and/or agent strategies over a diverse set of goods. Along this direction, areas of particular interest included:

- Distributed (scalable) algorithmic mechanism design
- Mechanisms for unreliable, dynamic, and asynchronous environments
- Mechanisms for incomplete and/or imperfect information environments
- Mechanisms for information goods and services
- Mechanisms for security, privacy, accounting, verification, and auditing
- Distributed (agent and mechanism) learning models
- Agents strategies in multi-institutional environments

- Economic and game theoretic specification, design, and analysis
- Bargaining, voting, and auction mechanisms
- Distributed reputation and trusted mechanisms
- User-agent interface design
- Agents that support bidding and negotiation
- Empirical evaluation of human-agent trading
- Eliciting human preferences and requirements
- Simulation and evaluation of properties of novel and complex mechanisms
- Goods, services, and contract description languages
- Mechanism description, verification, and testing languages
- Machine learning for mechanism identification problem
- Agent-mediated electronic system architectures and design principles
- Implemented agent-mediated electronic-commerce systems
- Mechanisms for business (supply chains, coalitions, and virtual enterprises)
- Mechanisms for Internet (congestion, routing, overlay, peer-to-peer, ad hoc networks)
- Mechanisms for novel applications

The workshop received a total of 39 submissions, from which 14 were selected for full presentation during the workshop. After the workshop, the authors were asked to submit their revised versions for publication in this volume. The result is that the volume contains 15 high-quality papers that can be regarded as representative of the field.

We have arranged the papers in the book around three major topics:

- Mechanism design
- Trading agents
- Tools

The first section contains eight papers dealing with a variety of issues on mechanism design. Dash et al. design an auction mechanism for allocating multiple goods when the buyers have interdependent valuations that turns out to be a generalization of the Vickrey-Clarke-Grove (VCG) mechanism. Conitzer et al. study two related problems concerning the VCG payment scheme: the problem of revenue guarantees and that of collusion. Motivated also by the problems of the VCG payment scheme, Faltings introduces a new mechanism that sacrifices Pareto-efficiency to achieve budget balance while being both incentive compatible and individually rational. Also motivated by problems with side-payment schemes, Jurca et al. present a mechanism that discovers (in equilibrium) the true outcome of a transaction by analyzing the two reports coming from the agents involved in the exchange. Larson et al. lay out mechanism design principles for deliberative agents: agents whose actions are modelled as part of their strategies. Juda et al. devise an options-based market infrastructure that enables bidders to use a dominant, truthful strategy across multiple, sequential auctions. A different, more empirical approach is taken by Phelps et al., who report on an evolutionary game-theoretic comparison of two double-auction market designs.

Finally, Kelly focuses on the computational realization of mechanisms by analyzing the use of generalized knapsack solvers for multi-unit combinatorial auctions.

The second section brings together a collection of papers on trading agents in a wide range of trading scenarios. Debenham et al. propose an agent bidding strategy that is not based on traditional game theory models but rather information theoretic; maximum entropy inference to determine the agent's actions taking into account the uncertain data he handles in actual-world scenarios. Gerding et al. design and empirically compare different bargaining strategies for selling agents when negotiating with many buyers. They show that bilaterally exchanging multiple offers combined with a random offer generation mechanism suffices to closely approximating Pareto-efficiency. Furthermore, they also analyse the versatility of combined strategies. Sherstov et al. report on the development and analysis of three autonomous stock-trading agents within the framework of the Penn Exchange Simulator, a novel stock-trading simulator. Approaches based on reinforcement learning, trend following, and market making are presented, evaluated individually against a fixed opponent strategy, and analysed comparatively. Pardoe et al. research on strategies for a different type of scenario: the trading agent competition supply chain management. They study the selling strategy of a supply chain agent to generate the set of bids to customers in simultaneous reverse auctions that maximizes the agent's expected profit. Sarne et al. focus on the analysis of agents' strategies for the dual parallel search in partnership formation applications. As a framework application they choose the classic voice communication partnerships application in an electronic marketplace. The authors manage to provide efficient means for the agents to calculate their distributed equilibrium strategies so that they can improve their expected utilities.

Finally, the third section contains two papers dealing with tools aimed at supporting the enactment of digital markets. On the one hand, the work by Michael et al. focuses on a scripting language and a run-time system that allow for the specification and monitoring of market mechanisms using rights and obligations. On the other hand, Reyes-Moro et al. introduce a bundling procedure intended to assist buyers when deciding whether to auction a bundle of goods as a whole or as separate, smaller bundles.

We would like to conclude by thanking the members of the Program Committee. They were able to produce a large number of high-quality reviews in a very short time span. Furthermore, we would also like to thank the authors for submitting their papers to our workshop, as well as the attendees and panelists for their valuable insights and discussions. Needless to say that these helped authors to improve the revised papers published in this book.

June 2005

Peyman Faratin  
Juan A. Rodríguez-Aguilar

# Workshop Organization

## Organizing Committee

Peyman Faratin	Massachusetts Institute of Technology, USA
Juan A. Rodríguez-Aguilar	IIIA-CSIC, Spain

## Program Committee

Martin Bichler	Technical University of Munich, Germany
Chris Brooks	University of San Francisco, USA
Monique Calisti	Whitestein Technologies, Switzerland
Jesús Cerquides	University of Barcelona, Spain
Vincent Conitzer	Carnegie Mellon University, USA
Frank Dignum	Universiteit Utrecht, The Netherlands
Marc Esteva	IIIA-CSIC, Spain
Boi Faltings	EPFL, Switzerland
Shaheen Fatima	Liverpool University, UK
Andrew Gilpin	Carnegie Mellon University, USA
Amy Greenwald	Brown University, USA
Nick R. Jennings	Southampton University, UK
Jeff Kephart	IBM Research, USA
Kate Larson	Carnegie Mellon University, USA
Kevin Leyton-Brown	University of Stanford, USA
Asuman Ozdaglar	Massachusetts Institute of Technology, USA
Julian Padget	University of Bath, UK
David Pennock	Overture Services, USA
Joerg Muller	Siemens, Germany
Simon Parsons	Brooklyn College, USA
Antonio Reyes-Moro	iSOCO, Spain
Jeff Rosenschein	Hebrew University, Israel
Rahul Sami	Massachusetts Institute of Technology, USA
Onn Shehory	Carnegie Mellon University, USA
Yoav Shoham	Stanford University, USA
Gerry Tesauro	IBM Research, USA
William Walsh	IBM Research, USA
Mike Wellman	University of Michigan, USA
Steven Willmott	UPC, Spain
Makoto Yokoo	Kyushu University, Japan



# Lecture Notes in Artificial Intelligence (LNAI)

- Vol. 3863: M. Kohlhase (Ed.), *Mathematical Knowledge Management*. XI, 405 pages. 2006.
- Vol. 3848: J.-F. Boulicaut, L. De Raedt, H. Mannila (Eds.), *Constraint-Based Mining and Inductive Databases*. X, 401 pages. 2006.
- Vol. 3847: K.P. Jantke, A. Lunzer, N. Spyros, Y. Tanaka (Eds.), *Federation over the Web*. X, 215 pages. 2006.
- Vol. 3835: G. Sutcliffe, A. Voronkov (Eds.), *Logic for Programming, Artificial Intelligence, and Reasoning*. XIV, 744 pages. 2005.
- Vol. 3817: M. Faundez-Zanuy, L. Janer, A. Esposito, A. Satue-Villar, J. Roure, V. Espinosa-Duro (Eds.), *Nonlinear Analyses and Algorithms for Speech Processing*. XII, 380 pages. 2006.
- Vol. 3814: M. Maybury, O. Stock, W. Wahlster (Eds.), *Intelligent Technologies for Interactive Entertainment*. XV, 342 pages. 2005.
- Vol. 3809: S. Zhang, R. Jarvis (Eds.), *AI 2005: Advances in Artificial Intelligence*. XXVII, 1344 pages. 2005.
- Vol. 3808: C. Bento, A. Cardoso, G. Dias (Eds.), *Progress in Artificial Intelligence*. XVIII, 704 pages. 2005.
- Vol. 3802: Y. Hao, J. Liu, Y.-P. Wang, Y.-m. Cheung, H. Yin, L. Jiao, J. Ma, Y.-C. Jiao (Eds.), *Computational Intelligence and Security*, Part II. XLII, 1166 pages. 2005.
- Vol. 3801: Y. Hao, J. Liu, Y.-P. Wang, Y.-m. Cheung, H. Yin, L. Jiao, J. Ma, Y.-C. Jiao (Eds.), *Computational Intelligence and Security*, Part I. XLI, 1122 pages. 2005.
- Vol. 3789: A. Gelbukh, Á. de Albornoz, H. Terashima-Marín (Eds.), *MICA 2005: Advances in Artificial Intelligence*. XXVI, 1198 pages. 2005.
- Vol. 3782: K.-D. Althoff, A. Dengel, R. Bergmann, M. Nick, T.R. Roth-Berghofer (Eds.), *Professional Knowledge Management*. XXIII, 739 pages. 2005.
- Vol. 3763: H. Hong, D. Wang (Eds.), *Automated Deduction in Geometry*. X, 213 pages. 2006.
- Vol. 3735: A. Hoffmann, H. Motoda, T. Scheffer (Eds.), *Discovery Science*. XVI, 400 pages. 2005.
- Vol. 3734: S. Jain, H.U. Simon, E. Tomita (Eds.), *Algorithmic Learning Theory*. XII, 490 pages. 2005.
- Vol. 3721: A.M. Jorge, L. Torgo, P.B. Brazdil, R. Camacho, J. Gama (Eds.), *Knowledge Discovery in Databases: PKDD 2005*. XXIII, 719 pages. 2005.
- Vol. 3720: J. Gama, R. Camacho, P.B. Brazdil, A.M. Jorge, L. Torgo (Eds.), *Machine Learning: ECML 2005*. XXIII, 769 pages. 2005.
- Vol. 3717: B. Gramlich (Ed.), *Frontiers of Combining Systems*. X, 321 pages. 2005.
- Vol. 3702: B. Beckert (Ed.), *Automated Reasoning with Analytic Tableaux and Related Methods*. XIII, 343 pages. 2005.
- Vol. 3698: U. Furbach (Ed.), *KI 2005: Advances in Artificial Intelligence*. XIII, 409 pages. 2005.
- Vol. 3690: M. Pěchouček, P. Petta, L.Z. Varga (Eds.), *Multi-Agent Systems and Applications IV*. XVII, 667 pages. 2005.
- Vol. 3684: R. Khosla, R.J. Howlett, L.C. Jain (Eds.), *Knowledge-Based Intelligent Information and Engineering Systems, Part IV*. LXXIX, 933 pages. 2005.
- Vol. 3683: R. Khosla, R.J. Howlett, L.C. Jain (Eds.), *Knowledge-Based Intelligent Information and Engineering Systems, Part III*. LXXX, 1397 pages. 2005.
- Vol. 3682: R. Khosla, R.J. Howlett, L.C. Jain (Eds.), *Knowledge-Based Intelligent Information and Engineering Systems, Part II*. LXXXIX, 1371 pages. 2005.
- Vol. 3681: R. Khosla, R.J. Howlett, L.C. Jain (Eds.), *Knowledge-Based Intelligent Information and Engineering Systems, Part I*. LXXX, 1319 pages. 2005.
- Vol. 3673: S. Bandini, S. Manzoni (Eds.), *AI\*IA 2005: Advances in Artificial Intelligence*. XIV, 614 pages. 2005.
- Vol. 3662: C. Baral, G. Greco, N. Leone, G. Terracina (Eds.), *Logic Programming and Nonmonotonic Reasoning*. XIII, 454 pages. 2005.
- Vol. 3661: T. Panayiotopoulos, J. Gratch, R.S. Aylett, D. Ballin, P. Olivier, T. Rist (Eds.), *Intelligent Virtual Agents*. XIII, 506 pages. 2005.
- Vol. 3658: V. Matoušek, P. Mautner, T. Pavelka (Eds.), *Text, Speech and Dialogue*. XV, 460 pages. 2005.
- Vol. 3651: R. Dale, K.-F. Wong, J. Su, O.Y. Kwong (Eds.), *Natural Language Processing – IJCNLP 2005*. XXI, 1031 pages. 2005.
- Vol. 3642: D. Ślęzak, J. Yao, J.F. Peters, W. Ziarko, X. Hu (Eds.), *Rough Sets, Fuzzy Sets, Data Mining, and Granular Computing, Part II*. XXIII, 738 pages. 2005.
- Vol. 3641: D. Ślęzak, G. Wang, M. Szczuka, I. Düntsch, Y. Yao (Eds.), *Rough Sets, Fuzzy Sets, Data Mining, and Granular Computing, Part I*. XXIV, 742 pages. 2005.
- Vol. 3635: J.R. Winkler, M. Niranjan, N.D. Lawrence (Eds.), *Deterministic and Statistical Methods in Machine Learning*. VIII, 341 pages. 2005.
- Vol. 3632: R. Nieuwenhuis (Ed.), *Automated Deduction – CADE-20*. XIII, 459 pages. 2005.
- Vol. 3630: M.S. Capcarrère, A.A. Freitas, P.J. Bentley, C.G. Johnson, J. Timmis (Eds.), *Advances in Artificial Life*. XIX, 949 pages. 2005.
- Vol. 3626: B. Ganter, G. Stumme, R. Wille (Eds.), *Formal Concept Analysis*. X, 349 pages. 2005.
- Vol. 3625: S. Kramer, B. Pfahringer (Eds.), *Inductive Logic Programming*. XIII, 427 pages. 2005.

# Table of Contents

## Mechanism Design

Revenue Failures and Collusion in Combinatorial Auctions and Exchanges with VCG Payments <i>Vincent Conitzer, Tuomas Sandholm</i> .....	1
A Mechanism for Multiple Goods and Interdependent Valuations <i>Rajdeep K. Dash, Alex Rogers, Nicholas R. Jennings</i> .....	15
A Budget-Balanced, Incentive-Compatible Scheme for Social Choice <i>Boi Faltings</i> .....	30
An Options-Based Method to Solve the Composability Problem in Sequential Auctions <i>Adam I. Juda, David C. Parkes</i> .....	44
“CONFESS”. Eliciting Honest Feedback Without Independent Verification Authorities <i>Radu Jurca, Boi Faltings</i> .....	59
Generalized Knapsack Solvers for Multi-unit Combinatorial Auctions: Analysis and Application to Computational Resource Allocation <i>Terence Kelly</i> .....	73
Designing Auctions for Deliberative Agents <i>Kate Larson, Tuomas Sandholm</i> .....	87
An Evolutionary Game-Theoretic Comparison of Two Double-Auction Market Designs <i>Steve Phelps, Simon Parsons, Peter McBurney</i> .....	101

## Trading Agents

Auctions and Bidding with Information <i>John Debenham</i> .....	115
Multi-attribute Bilateral Bargaining in a One-to-Many Setting <i>E.H. Gerding, D.J.A. Somefun, J.A. Han La Poutré</i> .....	129

Bidding for Customer Orders in TAC SCM  
*David Pardoe, Peter Stone* ..... 143

Agents' Strategies for the Dual Parallel Search in Partnership  
Formation Applications  
*David Sarne, Sarit Kraus* ..... 158

Three Automated Stock-Trading Agents: A Comparative Study  
*Alexander A. Sherstov, Peter Stone* ..... 173

**Tools**

Specifying and Monitoring Market Mechanisms Using Rights  
and Obligations  
*Loizos Michael, David C. Parkes, Avi Pfeffer* ..... 188

iAuctionMaker: A Decision Support Tool for Mixed Bundling  
*Antonio Reyes-Moro, Juan A. Rodríguez-Aguilar* ..... 202

**Author Index** ..... 215

# Revenue Failures and Collusion in Combinatorial Auctions and Exchanges with VCG Payments<sup>\*</sup>

Vincent Conitzer and Tuomas Sandholm

Computer Science Department, Carnegie Mellon University,  
5000 Forbes Avenue, Pittsburgh, PA 15213  
{conitzer, sandholm}@cs.cmu.edu

**Abstract.** In a *combinatorial auction*, there are multiple items for sale, and bidders are allowed to place a bid on a *bundle* of these items rather than just on the individual items. A key problem in this and similar settings is that of *strategic bidding*, where bidders misreport their true preferences in order to effect a better outcome for themselves. The *VCG payment scheme* is the canonical method for motivating the bidders to bid truthfully. We study two related problems concerning the VCG payment scheme: the problem of revenue guarantees, and that of collusion. The existence of such problems is known by many; in this paper, we lay out their full extent.

We study four settings: combinatorial forward auctions with free disposal, combinatorial reverse auctions with free disposal, combinatorial forward (or reverse) auctions without free disposal, and combinatorial exchanges. In each setting, we give an example of how additional bidders (colluders) can make the outcome much worse (less revenue or higher cost) under the VCG payment scheme (but not under a first price scheme); derive necessary and sufficient conditions for such an effective collusion to be possible under the VCG payment scheme; and (when nontrivial) study the computational complexity of deciding whether these conditions hold.

## 1 Introduction

In a *combinatorial auction*, there are multiple items for sale, and bidders are allowed to place a bid on a *bundle* of these items rather than just on the individual items. A rapidly growing body of computer science literature is devoted to the study of combinatorial auctions, and, to a lesser extent, variations of it, such as combinatorial reverse auctions (where the auctioneer seeks to procure certain items) and combinatorial exchanges (where bidders can offer goods for sale as well as express demand for goods—even within the same bid). One of the main reasons for the computer science community's interest in combinatorial auctions and exchanges is the hardness of the *clearing problem*. The clearing problem is to label bids as accepted or rejected to maximize the total value of the bids accepted (or, in the case of a reverse auction, to minimize their total value), under the natural constraint that the corresponding allocation of items does

---

<sup>\*</sup> This work is supported in part by NSF under CAREER Award IRI-9703122, Grant IIS-9800994, ITR IIS-0081246, and ITR IIS-0121678.

not require more items than are available (or, in the case of a reverse auction, under the constraint that all the desired items are procured). For example, the combinatorial auction clearing problem is NP-complete [15] (even to approximate [16]). Much research has focused on developing worst-case exponential time algorithms as well as approximation algorithms for the clearing problem [12,16,17,6].

Another key problem in auctions and exchanges (combinatorial or not) is that in general, the bidders may not bid their true valuations for the goods. For example, under a first-price payment rule, where bidders pay the value of their accepted bids, a bidder that bids her true valuation is entirely indifferent whether her bid is accepted or not. Thus, in order to benefit from the auction or exchange at all, bidders necessarily need to “shave” their bids, that is, report a lower value than their true value. The problem with untruthful bidding is that the clearing algorithm can only base the final allocation of the goods on the reported valuations, and thus the final allocation may not be optimal relative to the bidders’ true valuations. Thus, economic efficiency may be lost. Additionally, by a result known as the *revelation principle*, for any nontruthful mechanism, there is a truthful mechanism that performs just as well (under some assumptions on the strategic behavior of the bidders) [11]. It turns out, however, that by changing the payment rule, it is possible to motivate bidders to report their true valuations. The best-known such payment rule is the Vickrey-Clarke-Groves (VCG) scheme [18,3,8]. Here, a bidder must pay the total value of the bids that would have been accepted if she had not participated, minus the total value of bids that did get accepted (excluding her own bids). Because the bidder pays the externality she imposes on the other bidders (based on their reported valuations), she will bid to maximize the value of the final allocation—measured by her true valuation and the others’ reported valuations. If the clearing algorithm always finds the optimal allocation, bidding her true value will always effect this.<sup>1</sup> Not only is the VCG payment scheme the best-known payment scheme for motivating truthfulness, if the setting is general enough, given certain requirements, it (or its generalization to *Groves mechanisms*) is also the only one [7,9,20].

Unfortunately, there are also many problems with VCG mechanisms. They introduce the problem of lying auctioneers; they are bad from a privacy perspective; they are vulnerable to collusion; and they may lead to low revenue for the auctioneer. (In this paper, we will focus on the last two problems, which are closely related.) While most researchers in combinatorial auctions and exchanges acknowledge these problems, we believe that their severity may not be fully appreciated. For one, the Vickrey auction for a single item (the second-price sealed-bid auction) has some nice properties that unfortunately do not generalize to multi-item settings. For example, in a single-item Vickrey auction, it is not possible for colluders to obtain the item at a price less than the bid of any other bidder. Additionally, for the single-item Vickrey auction, various types of

---

<sup>1</sup> Of course, in general, the clearing problem may be too hard to always be solved optimally; and in general, the VCG scheme will not motivate bidders to bid truthfully when the final allocation may be suboptimal. A growing body of research is dedicated to finding approximation schemes for the (forward auction) clearing problem that still motivate bidders to bid truthfully, or attempting to prove that this is impossible in general [13,10,2,9]. Throughout this paper we will assume that the auctions and exchanges are cleared optimally.

revenue equivalence with (for example) first-price sealed-bid auctions hold. As we will show, in the multi-item setting these properties do not hold at all (and can be violated to an arbitrary extent). We hope that a greater awareness of these issues will help bridge the gap between theory and practice in mechanism design for combinatorial auctions and exchanges.

To that end, in this paper we give some detailed worst-case results about collusion and revenue. For the various variants of combinatorial auctions and exchanges, we study the following single problem that relates both issues under consideration:

*Given some of the bids, how “bad” can the remaining bidders make the outcome?*

“Bad” here means that the bidders are paid an inordinately large amount, or pay an inordinately small amount, relative to the goods they receive and/or provide. This is closely related to the problem of making any revenue guarantees to the auctioneer. But it is also the collusion problem, if we conceive of the remaining bidders as colluders.

As it will turn out, our fundamental problem is often computationally hard. Computational hardness here is a double-edged sword. On the one hand, if the problem is hard, collusion may not occur (or to a lesser extent) because the colluders cannot find the (most) beneficial collusion. On the other hand, if the problem is hard, it is difficult to make strong revenue guarantees to the auctioneer.

## 2 The VCG Mechanism

All the results in this paper hold even when all bidders are *single-minded*: that is, they are interested in only one bundle of items (or, in the case of a reverse auction, can provide only one bundle of items). In this case, every bid corresponds to a unique *utility function*, namely the one for which this bid would have been a truthful revelation of the bidder’s valuation for the bundle.

The VCG payment scheme proceeds as follows: accept bids so that the resulting allocation maximizes the sum of the bidders’ utilities as implied by the bids, not taking payments into account. (This is simply maximizing the sum of the values of the accepted bids in a forward auction, or minimizing their sum in a reverse auction.) Call this sum of utilities  $a$ . Then, to determine winning bidder  $i$ ’s payment, remove that bidder’s bid, and see what the maximum sum of the utilities (disregarding payments) would have been with only the remaining bids. Call this sum of utilities  $b_i$ . Winning bidder  $i$  then must pay the second sum of utilities, minus the original sum of utilities of the other bidders—that is, the externality she imposed on the other bidders. (Thus, if the value of winning bidder  $i$ ’s bid is  $v_i$ , the payment is  $b_i - a + v_i$ .) We observe that this payment can be negative, if the bidder’s presence actually makes the other bidders better off (disregarding the payments).<sup>2</sup> For instance, in a reverse auction where goods are disposable, each winning bidder’s payment will be nonpositive (the other bidders need to supply fewer items when this bidder is present, so (disregarding payments) this bidder makes the others better off).

<sup>2</sup> This can only happen if the bidder’s presence actually makes more allocations to the other bidders possible. For instance, in a forward auction with free disposal, when we remove a bidder we can always throw away the items allocated to her and keep the allocation to all other bidders the same. Thus, payments cannot be negative in this setting.

### 3 Combinatorial (Forward) Auctions

#### 3.1 Review

In a combinatorial auction, there is a set of items  $I = \{A_1, A_2, \dots, A_m\}$  for sale. A bid takes the form  $b = (B, v)$ , where  $B \subseteq I$  and  $v \in \mathbb{R}$ . The clearing problem is to label bids as accepted or rejected, to maximize the sum of the values of the accepted bids, under the constraint that no item occurs in more than one accepted bid. (This is assuming *free disposal*, that is, items do not have to be allocated to anyone.) We say a bid is *truthful* if the value attached to the bundle is the bidder's utility for that bundle.

#### 3.2 Motivating Example

Consider an auction with two items,  $A$  and  $B$ . Suppose we have collected two bids (from different bidders), both  $(\{A, B\}, N)$ . If these are the only two bids, one of the bidders will be awarded both the items and, under the VCG payment scheme, have to pay  $N$ . However, suppose two more bids (by different bidders) come in:  $(\{A\}, N + 1)$  and  $(\{B\}, N + 1)$ . Then these bids will win. Moreover, neither bidder will have to pay anything!

This example demonstrates a number of issues. First, the addition of additional bidders may actually decrease the auctioneer's revenue from an arbitrary amount to 0. Second, the VCG mechanism is not revenue equivalent to the sealed-bid first-price mechanism in combinatorial auctions, even when all bidders' true valuations are common knowledge—unlike in the single-item case. (The first-price mechanism will generate positive expected revenue for these valuations; we omit the proof because of space constraint.) Third, even when the other bidders by themselves would generate nonnegative revenue for the auctioneer under the VCG payment scheme, it is possible that two colluders can bid so as to receive all the items without paying anything.

The following proposition sums up the properties of this example.

**Proposition 1.** *In a forward auction (even with only 2 items), there exists a family of instances (sets of bids) such that: 1. The winning bidders pay nothing under the VCG payment scheme; 2. If the winning bids are removed, the remaining bids actually generate revenue  $N$  under the VCG payment scheme; 3. If these bids were truthful (as we would expect under VCG), then if we had run a first-price sealed-bid auction instead (and the bidders knew each other's true valuations), any equilibrium would have generated revenue  $\Theta(N)$ .*

#### 3.3 Characterization

We now proceed to characterize the settings where the colluders can receive all the items for free.

**Lemma 1.** *If the colluders receive all the items at cost 0, then for any positive bid on a bundle of items by a noncolluder, at least two of the colluders receive an item from this bundle.*

*Proof.* Suppose that for some positive bid  $b$  on a bundle  $B$  by a noncolluder  $i$ , one of the colluders  $c$  receives all the items in  $B$  (and possibly others). Then, in the auction

where we remove that colluder's bids, one possible allocation gives every remaining bidder all the goods that bidder received in the original auction; additionally, it gives  $i$  all the items in bundle  $B$ ; and it disposes of all the other items  $c$  received in the original auction. With this allocation, the total value of the accepted bids by bidders other than  $c$  is at least  $v(b)$  more than in the original auction. Because the total value obtained in the new auction is at least the value of this particular allocation, it follows that  $c$  imposes a negative externality of at least  $v(b)$  on the other bidders, and will pay at least  $v(b)$ . ■

**Lemma 2.** *Suppose all the items in the auction can be divided among the colluders in such a way that for any positive bid on a bundle of items by a noncolluder, at least two of the colluders receive an item from this bundle. Then the colluders can receive all the items at cost 0.*

*Proof.* For the given partition of items among the noncolluders, let each colluder place a bid with an extremely large value on the bundle consisting of the items assigned to him in the partition. (For instance, twice the sum of the values of all noncolluders' bids.) Then, the auction will clear awarding each colluder the items assigned to him by the partition. Moreover, if we remove the bids of one of the colluders, all the remaining colluders' bids will still win—and thus none of the noncolluders' bids will win, because each such bid requires items assigned to at least two colluders by the partition (and at least one of them is still in the auction and wins these items). Thus, each colluder (individually) imposes no externality on the other bidders. ■

Combining these two lemmas, we get:

**Theorem 1.** *The colluders can receive all the items for free if and only if it is possible to divide the items among the colluders in such a way that for any (nonzero) bid by a noncolluder, the items in that bid are spread across at least two colluders.*

### 3.4 Complexity

**Definition 1 (DIVIDE-SUBSETS).** *Suppose we are given a set  $S$ , as well as a collection of subsets of it,  $R = \{S_1, \dots, S_q\}$ . We are asked whether  $S$  can be partitioned into  $n$  parts  $T_1, T_2, \dots, T_n$  so that no subset  $S_i \in R$  is contained in one of these parts.*

**Theorem 2.** *DIVIDE-SUBSETS is NP-complete, even when  $n = 2$ .*

Complexity proofs are omitted because of space constraint.

## 4 Combinatorial Reverse Auctions

### 4.1 Review

In a combinatorial reverse auction, there is a set of items  $I = \{A_1, A_2, \dots, A_m\}$  to be procured. A bid takes the form  $b = (B, v)$ , where  $B \subseteq I$  and  $v \in \mathbb{R}$ . The clearing problem is to label bids as accepted or rejected, to minimize the sum of the values of the accepted bids, under the constraint that each item occurs in at least one accepted bid. (This is assuming *free disposal*, that is, items do not have to be allocated to anyone.) We say a bid is *truthful* if the value attached to the bundle is the bidder's cost for providing that bundle.



## 4.2 Motivating Example

Consider a reverse auction with  $m$  items,  $A_1, A_2, \dots, A_m$ . Suppose we have collected two bids (from different bidders), both  $(\{A_1, A_2, \dots, A_m\}, N)$ . If these are the only two bids, one of the bidders will be chosen to provide all the goods, and, under the VCG payment scheme, be paid  $N$ . However, suppose  $m$  more bids (by different bidders) come in:  $(\{A_1\}, 0), (\{A_2\}, 0), \dots, (\{A_m\}, 0)$ . Then these bids will win. Moreover, each bidder will be paid  $N$  under the VCG payment scheme. (Without this bidder, we would have had to accept one of the original bids.) Thus, the total payment that needs to be made is  $mN$ .

Again, this example demonstrates a number of issues. First, the addition of additional bidders may actually increase the total amount that the auctioneer needs to pay. Second, the VCG mechanism requires much larger payments than a first-price auction in the case where all bidders know each others' valuations (and the equilibrium is in pure strategies<sup>3</sup>). (The first-price mechanism will not require a total payment of more than  $N$  for these valuations in any pure-strategy equilibrium; we omit the proof because of space constraint.)

Third, even when the other bidders by themselves would allow the auctioneer to procure the items at a low cost under the VCG payment scheme, it is possible for  $m$  colluders to get paid  $m$  times as much for all the items.

The following proposition sums up the properties of this example.

**Proposition 2.** *In a reverse auction, there exists a family of instances (sets of bids) such that: 1. The winning bidders are paid  $mN$  under the VCG payment scheme; 2. If the winning bids are removed, the remaining bids allow the auctioneer to procure everything at a cost of only  $N$ ; 3. If these bids were truthful (as we would expect under VCG), then if we had run a first-price sealed-bid reverse auction instead (and the bidders knew each other's true valuations), any equilibrium in pure strategies would have required total payment of at most  $N$ . (However, there are also mixed-strategy equilibria with arbitrarily large expected total payment.)*

## 4.3 Characterization

Letting  $N$  be the sum of the values of the accepted bids when all the colluders' bids are taken out, it is clear that no colluder can be paid more than  $N$ . (With the colluder's bid, the sum of the values of others' accepted bids is still at least 0; without it, it can be at most  $N$ , because in the worst case the auctioneer can accept the bids that would be accepted if none of the colluders are present.) In this subsection, we will identify a necessary and sufficient condition for the colluders to be able to each receive  $N$ .

**Lemma 3.** *If a colluder receives  $N$ , then the items that she has to provide cannot be covered by a set of noncolluders' bids with cost less than  $N$ .*

<sup>3</sup> Perhaps surprisingly, the first-price combinatorial reverse auction for this example (with commonly known true valuations corresponding to the given bids) actually has mixed-strategy equilibria with arbitrarily high expected payments. We omit the proof because of space constraint.