

Conference Report: The Fourth ACM Conference on Electronic Commerce

Kevin Leyton-Brown

and

Bob McGrew

and

Ryan Porter

Computer Science Department, Stanford University

{kevinlb;rwporter;bmcgrew}@stanford.edu

Permission to make digital/hard copy of all or part of this material without fee for personal or classroom use provided that the copies are not made or distributed for profit or commercial advantage, the ACM copyright/server notice, the title of the publication, and its date 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.

© 2003 ACM 0000-0000/2003/0000-0041 \$5.00

1. INTRODUCTION

The 4th ACM Conference on Electronic Commerce (EC'03) took place in San Diego, California, from June 9 to June 12, 2003. The conference consisted of 21 research papers, 50 posters and four tutorials, covering topics at the interface between computer science, economics, game theory and management. The EC'03 general chair was Danny Menascé and the program chair was Noam Nisan. The tutorial co-chairs were Joan Feigenbaum and Benjamin Grosz. The program committee included Virgilio A. F. Almeida, Issy Ben-Shaul, Sushil Bikhchandani, Andrei Broder, Jeremy Epstein, Anna Karlin, Rudolf Müller, David Parkes, Tal Rabin, Michael Reiter, Tuomas Sandholm, Ilya Segal, Oded Shmueli, Eva Tardos, Moshe Tennenholtz, Jeff Ullman, Vijay Vazirani, and Andrew Whinston.

For the first time, the conference was held in the early summer rather than the fall. This allowed EC'04 to be co-located with ACM's Federated Computing Research Conference (FCRC). This move allowed a great deal of cross-fertilization between EC and other ACM conferences. In particular, participants were able to attend several STOC sessions on topics related to mechanism design, and EC was attended by a larger contingent of theoretical computer scientists than in past years. Despite the twenty-month gap between EC'03 and EC'04, the conference once again attracted very high quality papers. New to the conference this year was the addition of a poster track, which allowed a large number of papers to be presented while still maintaining a single-track plenary session.

2. TUTORIALS

The first day of the conference consisted of four half-day tutorials in two parallel tracks. In the morning, David Parkes surveyed *Computational Issues in Auction Design*. This tutorial discussed four sources of complexity in electronic auctions: centralized computational complexity, communication and preference-elicitation complexity, the complexity of computing solution concepts and the complexity of designing economic mechanisms. In the second morning tutorial, *Mobile Commerce*, Norman Sadeh introduced participants to the technologies, applications, services and business models associated with mobile commerce.

In the afternoon, David Pennock and Michael Wellman led a tutorial on *Markets in Uncertainty: Risk, Gambling and Information Aggregation*, discussing financial instruments that pay off based on the realization of an uncertain variable. This tutorial covered the economic background on securities markets as well as more recent theoretical and empirical work on these markets' computational characteristics. Finally, Michael Kearns gave an introduction to *Computational Game Theory*, first describing the fundamentals of complete-information games and then discussing recent work on compact representations of games, the computation of different solution concepts, and connections to reinforcement learning.

3. PLENARY TALKS

The conference's plenary talks covered 21 papers on a wide variety of topics in both Computer Science and Economics, with methodologies running the full range from theoretical to experimental. This report organizes the papers into six categories: auction design; other mechanism design; computing solution concepts for complete-information games; other computational issues; agent behavior; securities markets.

3.1 Auction Design

A topic that has attracted many papers at ACM-EC since the first conference is design of novel auction mechanisms to solve specific e-commerce problems.

Efficiency through feedback-contingent fees and rewards in auction marketplaces with adverse selection and moral hazard, by Chrysanthos Dellarocas, proposed a mechanism for inducing cooperation in online auctions with noisy monitoring of quality and adverse selection. A listing fee contingent on announced quality was combined with feedback from buyers to induce sellers to produce high quality goods and to honestly declare these quality levels.

Envy-Free Auctions for Digital Goods, by Andrew Goldberg and Jason Hartline, performed a theoretical analysis of auctions for commodities in unlimited supply. It showed that no constant-competitive auction is both truthful and envy-free (i.e., makes use of non-discriminatory prices), but that constant-competitive auctions can be constructed when either truthfulness or envy-freedom is relaxed with arbitrarily small probability.

Highly Efficient, Budget-Balanced, Incentive-Compatible, Individually Rational Auctions for Supply Chain Formation, by Moshe Babaioff and William Walsh, considered the design of mechanisms for the supply chain setting. The Myerson-Satterthwaite impossibility theorem shows that it is impossible to construct a mechanism for this setting that is incentive-compatible, budget-balanced and efficient;

this paper showed an incentive-compatible and budget-balanced mechanism that gives a worst-case bound on efficiency.

A robust open ascending-price multi-unit auction protocol against false-name bids, by Atsushi Iwasaki, Makoto Yokoo and Kenji Terada, considered the problem of designing a mechanism for ascending multi-unit auctions when bidders can place bids under more than one identity. This work extended the Ausubel auction to the case where bidders' valuations may not be marginally decreasing, and showed experimental results.

Approximately-Strategyproof and Tractable Multi-Unit Auctions, by Anshul Kothari, David Parkes and Subhash Suri, also considered single-good multi-unit auctions, but focused on tractability and efficiency in the case of marginally-decreasing valuations. A polynomial-time mechanism was presented which gives tight bounds on both efficiency and the amount an agent can gain by lying.

3.2 Other Mechanism Design

This section describes mechanism design papers that did not focus on auctions.

k-Implementation, by Dov Monderer and Moshe Tennenholtz, discussed a setting in which a third party wishes to influence the behavior of self-interested agents playing a game by offering non-negative financial transfers whenever agents play certain strategy profiles. This paper introduced the notion of k -implementation, where k is the cost to the third party of causing a particular strategy profile to become a dominant strategy for all agents, and considered both complete and incomplete-information games.

How Much Can Taxes Help Selfish Routing?, by Richard Cole, Yevgeniy Dodis and Tim Roughgarden, studied economic incentives for influencing selfish behavior in networks, considering a model of selfish routing in which latency is a function of congestion and users route selfishly on minimal-latency paths. They showed that marginal cost taxes do not improve the cost of the Nash equilibrium in a large class of networks, and that the best-possible benefit from arbitrary taxes does not exceed that of edge removal.

Strategyproof Cost-sharing Mechanisms for Set Cover and Facility Location Games, by Nikhil R. Devanur, Vijay V. Vazirani and Milena Mihail, took a cooperative game theory perspective. The paper described polynomial-time strategyproof cost allocations for the set cover game and the facility location game.

3.3 Agent Behavior

A fundamental problem in e-commerce is determining how agents ought to behave in complex economic environments.

Agent-Based Service Composition Through Simultaneous Negotiation, by Chris Preist, Claudio Bartolini and Andrew Bye, considered the problem of automatically bundling together several component products or services to meet the needs of a given customer. The paper described a service composition agent that manages the risk of simultaneously buying components in English auctions and selling services in RFQ reverse auctions.

On Cheating in Sealed-Bid Auctions, by Ryan Porter and Yoav Shoham, analyzes two forms of cheating in sealed-bid auctions and discusses ways in which agents should respond. First, it considered a seller who spies on the bids of a second-

price auction and inserts a fake bid to increase the payment of the winning bidder; second, it considered first-price auctions in which bidders can cheat by examining other competing bids before submitting their own bids.

Exploring Bidding Strategies for Market-Based Scheduling, by Michael Wellman, Jeffrey K. MacKie-Mason, Daniel Reeves and Sowmya Swaminathan, considered a scheduling mechanism that allocates a single resource over time based on the bids of participating agents, using separate, simultaneous markets for the different time slots. The paper explored strategies that agents might employ in this setting, identifying these strategies primarily with evolutionary search methods.

Selling Online versus Offline: Theory and Evidences at Sotheby's, by Eiichiro Kazumori and John McMillan, asked the question, "How and why does a firm use online markets versus traditional offline markets?" using a dataset of 3000 auctions held by Sotheby's both online and offline in summer 2002. They found that offline auctions tended to make it easier for the seller to alleviate buyers' uncertainty about goods' values, but also have much higher transaction costs for buyers. The authors modeled this tradeoff between information revelation and participation costs using a simple microstructure model.

3.4 Computing Solution Concepts for Complete-Information Games

Another major topic at EC-03 was the computation of solution concepts for games.

Playing Large Games Using Simple Strategies, by Richard Lipton, Vangelis Markakis and Aranyak Mehta, proved the existence of ϵ -Nash equilibrium strategies with support logarithmic in the number of pure strategies, and where the players' payoffs are almost the same as in the original Nash equilibrium. They used this result to give a quasi-polynomial time algorithm for computing an ϵ -Nash equilibrium.

Correlated Equilibria in Graphical Games, by Sham Kakade, Michael Kearns, John Langford and Luis Ortiz, introduced a Markov network which represents all correlated equilibria of a graphical game. Using this network, they provided a general algorithm for computing correlated equilibria in graphical games.

A polynomial-time Nash equilibrium algorithm for repeated games, by Michael Littman and Peter Stone, presented an algorithm for finding finite-state equilibrium strategies in average-payoff infinitely-repeated bimatrix games. This work drew on a finite-state automaton representation and on the folk theorem.

On the Core of the Multicommodity Flow Game, by Vangelis Markakis and Amin Saberi, considered a formalization of the notion of routing stability in BGP as a coalitional game theory problem. They showed that that the core is nonempty (i.e., that no node in the network has incentive to secede) both when nodes are allowed to transfer payoff within a coalition and when they are not.

3.5 Other Computational Issues

In this section we describe papers that addressed other computational issues in e-commerce.

Convergent Algorithms for Collaborative Filtering, by Jon Kleinberg and Mark Sandler, sought to provide a theoretical underpinning to the intuition that collaborative filtering algorithms should offer better recommendations when they have more data. They described a recommendation algorithm that converges in the sense that as the amount of data increases, its recommendations approach those of

an omniscient algorithm that knows users' hidden preferences.

Searching Dynamically Bundled Goods with Pairwise Relations, by Yuan-Chi Chang, Chung-Sheng Li and John Smith, presented a dynamic programming based algorithm for finding the best bundles of goods to satisfy user preferences. They demonstrated this algorithm using an example from packaged travel, where discounts arising from business partnerships between airline, hotel and car rental companies could be exploited.

On Polynomial-Time Preference Elicitation, by Martin Zinkevich, Tuomas Sandholm and Avrim Blum, studied special cases of preferences for which elicitation can be done in polynomial time using value queries. They also considered the problem of allocating items given these preferences, showing that in some of the cases this problem is \mathcal{NP} -hard, and in other cases that it can be solved in polynomial time.

3.6 Securities Markets

Finally, there were two papers that examined securities markets.

Betting Boolean-style: A Framework for Trading in Securities Based on Logical Formulas, by Lance Fortnow, Joe Kilian, David Pennock and Michael Wellman, considered compound securities, financial instruments that pay off contingent on the outcomes of arbitrary statements in propositional logic. They addressed the auctioneer's matching problem, showing that it is $\text{co-}\mathcal{NP}$ -complete in the divisible case and Σ_2^P -complete in the indivisible case.

Computation in a Distributed Information Market, by Joan Feigenbaum, Lance Fortnow, David Pennock and Rahul Sami, investigated the claim that the equilibrium price of a financial security reflects all of the available information regarding the security's value. They showed that securities whose payoffs cannot be expressed as weighted threshold functions of distributed input bits are not guaranteed to converge to the proper equilibrium; when the security *can* be expressed in this way, convergence always occurs and is linear in the number of bits of distributed information.

4. POSTERS

The plenary talks were complimented by a very active poster session containing 50 posters. This section gives a broad—though incomplete—overview of the wide range of topics that were covered, briefly touching on several illustrative posters.

A large number of the posters fell under the general heading of mechanism design. Multicast networks remained a popular computer science application for mechanism design, with two posters (by Chawla, Kitchin, Rajan, Ravi and Sinha; Mehta, Shenker and Vazirani) focusing on profit maximization and another by Bläser and Manthey focusing on budget balance. An interesting variant of the mechanism design setting was introduced by Friedman and Parkes, who considered the online problem of allocating and pricing bandwidth to users who arrive over time. A novel challenge in this setting is ensuring that agents have incentive to truthfully declare their arrival. A similar challenge was tackled by Ng, Parkes, and Seltzer, who looked at strategyproof mechanisms for distributed computation through series of auctions.

Another variant of mechanism design looked at implementing protocols in an incentive-aware fashion. For example, one poster by Brandt looked at what sorts

of mechanisms can be implemented privately using cryptography, showing completeness results for cryptographic implementations of mechanisms. Considering a related topic, Schniedman and Parkes looked at using distributed systems techniques to make the BGP protocol incentive-compatible, so that a machine executing the BGP protocol would never have an incentive to disobey the protocol.

Another well-represented category was combinatorial auctions. Motivated by the fact that an exponential number of XOR bids may be needed to represent a bidder's valuation, three posters (by Sunderam and Parkes; Hudson and Sandholm; Ghosh and Kalagnanam) offered ways to incrementally elicit bidders' preferences. Another poster by Wurman, Cai, and Sureka presented a method for efficiently determining the outcome of a myopic, open-outcry combinatorial auction in which all agents use proxy bidders. This algorithm breaks the problem up into periods of time in which the bid increments for each bundle can be represented by a linear function, then uses linear programming to calculate the state of the auction at the end of each such stage.

Various papers discussed reputation, trust, and authentication. Fujimura and Nishihara looked at reputation systems in e-commerce with multiple sellers and multiple buyers, while Xiong and Liu considered trust models in a peer-to-peer setting. Nuckolls, Martel, and Stubblebine considered how to authenticate data for third-party publishing.

Finally, several posters addressed the issue of pricing. Pricing of information goods was the topic of two posters (by Liu, Wynter and Xia; Polk, Hanson, Ledyard and Ishikida), while a poster by Bichler and Kalagnanam offered a technique for setting the reserve price in a procurement auction. A fourth poster, by Wellman, Reeves, Lochner and Vorobeychik, surveyed the price prediction strategies in the last Trading Agent Competition, and offered a new metric for comparison, called the value of perfect prediction.

5. CONCLUSIONS

The ACM Electronic Commerce conference, sponsored by the ACM Special Interest Group on E-commerce (SIGecom), had a very successful fourth session in San Diego, California. The addition of the poster session meant that dramatically more authors were involved this year, as compared to any of the three previous sessions. The new tutorial tracks gave attendees the chance to learn about areas of e-commerce outside their specialties. The conference continues to serve as a top-quality venue for the presentation and publication of cutting-edge research in the emerging and exciting research area at the interface between computer science, economics, game theory and management.