Conference Report: The Third ACM Conference on Electronic Commerce

David C. Parkes
Division of Engineering and Applied Science,
Harvard University,
33 Oxford Street,
Cambridge, MA 02138
parkes@eecs.harvard.edu

The 3rd ACM Conference on Electronic Commerce (EC'01) took place in Tampa, Florida, from October 14th until October 17th, 2001. The conference contained 35 research papers in topics at the interface between computer science, economics, commerce and law. The 2001 Trading Agent Competition (TAC) was also associated with the conference. This report summarizes the program.

1. INTRODUCTION

The 3rd ACM Conference on Electronic Commerce (EC'01) took place in Tampa, Florida, from October 14th until October 17th, 2001. The conference contained 35 research papers in topics at the interface between computer science, economics, commerce and law [EC 2001]. The EC'01 general chair was Michael Wellman and the program chair was Yoav Shoham.[1]

The EC conference continues to attract researchers from across the traditional computer science disciplines– theoretical computer science, AI, systems, and networking –in addition to researchers from operations research, economics, and schools of information science. Following EC'99 and EC'00, the conference once again attracted high quality papers, with the presentation of both theoretical and experimental results. The ACM E'commerce conference is organized around a single paper track, which allows for a healthy exchange of viewpoints.

[1]The tutorial co-chairs were Chris Dellarocas and Benjamin Grosof; the program committee included Martin Abadi, Dan Ariely, Erik Brynjolfsson, Lorrie Cranor, Matthew Franklin, Amy Greenwald, Andrew Goldberg, David Heckerman, Jeff Kephart, Ronny Kohavi, Piero La Mura, John Ledyard, Daniel Lehmann, Daniel Menasce, Haim Mendelson, Christos Papadimitriou, Amir Ronen, Scott Shenker, Yoav Shoham, Dan Suciu, and Moshe Vardi.

Many papers in this year's conference were in the general area of computational mechanism design, with new theoretical and experimental results presented that relate to the principled design of systems that consider the incentives of individual agents and users together with computational and informational efficiency. Other papers covered topics as diverse as privacy, dynamic information bundling, bidding agents, user modeling, web server QoS optimization, network resource allocation, and preference elicitation. Conference highlights included invited talks by Al Roth (Harvard) and Jeremy Bulow (Stanford, FTC).

2. TRADING AGENT COMPETITION, INVITED LECTURES, AND PANEL DISCUSSION

The 2001 Trading Agent Competition (TAC), organized by Michael Wellman and a team at the University of Michigan, with significant contributions from Amy Greenwald (Brown), Peter Stone (AT&T) and Peter Wurman (NCSU), was associated with the main conference. By the day of the competition an initial field of 28 agents had been reduced to 16, that was reduced to 8 after the semifinals, which were during the morning. The finals consisted of 24 games among the 8 surviving agents, with `livingagents` (Living Systems AG) the eventual winner and `ATTac` (AT&T Research) a close runner-up. Although the design of the competition was changed considerably from the 2000 competition [Stone and Greenwald 2001] there was evidence that teams had learned lessons from the previous year, for example in the method used to compute an agent's utility over different outcomes.

Approaches remained more decision-theoretic than game-theoretic, although some agents learned different strategies from historical data about the games played by the opponents in a particular round. Technical methodologies included *Bayesian modeling* (e.g. `whitebear`, Cornell), *mathematical programming* (e.g. `TacsMan`, Stanford), *statistical methods* (e.g. `polimi_bot`, Polit. di Milano), *classifier learning* (e.g. `ATTac`), *monte-carlo methods* (e.g. `RoxyBot`, Brown), *Markov-chain modeling* (e.g. `Retsina`, CMU), *Q-learning* (e.g. `jbroadw`, McGill), and *fuzzy decision rules* (e.g. `SouthamptonTAC`, U. Southampton).

In an invited lecture, Al Roth, an experimental economist and game-theorist at Harvard University, presented his recent work with Axel Ockenfels [Roth and Ockenfels 2001] on last-minute bidding in on-line auctions, in which the effect of the closing rule on bidding behavior is examined. The work is interesting in that it shows that a even a very subtle change in the rules of an economic mechanism can have dramatic effects on the strategies of participants. Jeremy Bulow, Stanford University and Director of the Bureau of Economics at the FTC, discussed technical innovation and regulatory issues, in particular with respect to the software industry.

An interesting panel, moderated by Daniel Menascé, discussed teaching electronic commerce in computer science. Participants included Erik Brynjolfsson, Stuart Feldman, Michael Wellman and Yoav Shoham. There was some consensus on an approach oriented on the "automation of commerce", and on computing with incentives as a useful unifying theme. Participants expressed concerns about the need to identify the underling core concepts and disciplinary foundations of e-commerce, and the challenges of positioning an e-commerce course within a traditional undergraduate computer science curriculum.

3. TECHNICAL SESSIONS

The following section presents a brief and personal review of some of the papers presented at the conference. Many of the papers fit into one of the following major themes: (a) computational mechanism design; (b) computational complexity of clearing in markets; (c) preference elicitation; (d) optimization for e-commerce servers; (e) information and price dynamics.

Computational Mechanism Design

Computational mechanism design continues to emerge as an important research area, attracting considerable interest in AI and more recently within theoretical CS. Successful mechanisms to coordinate agents, and optimize outcomes in open decentralized systems should be computationally tractable, while also handling the incentive issues that follow from agent rationality.

Moshe Babaioff and Noam Nisan proposed a protocol to chain double auctions together across a supply chain, such that participants' need only bid in local double auctions, and such that truth-revelation is a dominant strategy. The protocol carefully coordinates clearing across markets via information aggregation of supply and demand curves before clearing, to achieve strategy-proofness, budget feasibility, and as much efficiency as possible in any single double auction.

Amir Ronen proposed a method to approximate optimal, or revenue maximizing, auction design in the case of general distributions over agent preferences. Ronen proposes a polynomial time computable and dominant strategy auction, that is a subtle variation on Vickrey's [61] second-price auction. The auction guarantees at least half of the optimal revenue. In a later paper, Ronen proposed a method to adapt Vickrey-Clarke-Groves (VCG) schemes to handle incomplete bidding languages, while retaining truthfulness. Agents can submit "appeal functions" in addition to bids. The truthfulness of the resulting mechanism is proved within a bounded-rational equilibrium solution concept.

Joan Feigenbaum, Arvind Krishnamurthy, Rahul Sami and Scott Shenker presented an interesting extension of their earlier work on the network implementation of cost-sharing mechanisms for multicast trees. An approximation to the budget-balanced Shapley mechanism was presented, with exponentially less worst-case communication, that nevertheless remains group-strategyproof and provides bounds on the effect of the approximation on budget loss and efficiency loss. In addition, the authors have characterized the successful collusive strategies available to receivers in the individual strategy-proof and efficient VCG mechanism.

Kevin Leyton-Brown, Ryan Porter, Shobha Venkataraman and Balaji Prabhakar, introduced a set of mechanisms to provide incentives for temporal demand shifting in a network resource allocation problem. Network slots are associated with fixed prices, but this price is waived probabilistically to smooth out network demand and achieve useful incentive properties.

Computational Complexity of Clearing in Markets

Benny Lehmann, Daniel Lehmann and Noam Nisan presented a hierarchical characterization of the class of combinatorial allocation problems (CAPs) for agents with decreasing marginal utilities. A transition from polynomial time solvable to NP-hard is demonstrated between the class of *gross-substitutes* preferences and the

class of *submodular* preferences, with NP-hardness with submodular preferences established by reduction from the knapsack problem.

Edo Zurel and Noam Nisan proposed an approximation algorithm to compute solutions to the winner determination problem in combinatorial auctions. The algorithm, ALPH, is a novel combination of an approximate solution to an linear program relaxation of the CAP followed by a sequence of greedy hill-climbing methods to improve the final solution quality. The algorithm is 2-3 orders of magnitude faster than the fastest known *optimal* algorithms, computing solutions in around the same time it takes to send the problem instance over a T1 line, and with less than 1% error on average.

Preference Elicitation

While the cost structure in electronic commerce has dramatically expanded the applicability of dynamic mechanisms, bringing down the costs of dynamic pricing and negotiation, the cost of preference elicitation, from individuals and/or businesses, continues to provide an important bottleneck.

Vijay Iyengar, Jon Lee and Murray Campbell introduced a principled approach to minimal preference elicitation in the special case of ranking a set of multiattribute items. The method sequences ordinal comparisons across items to try to maximize the reduction in uncertainty about the user's preferences. In related work, Wolfram Conen and Tuomas Sandholm addressed the problem of preference elicitation in application to a combinatorial auction, in which a *preference elicitator agent* asks users for partial information about their preferences across different bundles of items to determine the efficient allocation of items.

Optimization for E-commerce Servers

A number of papers considered applications of optimization, queueing models, and statistics to configuration and load balancing problems in web servers for e-commerce applications. Zhen Liu, Mark Squillante and Joel Wolf introduced a queueing-theory and optimization-based model to formalize and solve the problem of assigning server loads to maximize expected profit, given a set of service level agreements that provide customers with Quality of Service (QoS) guarantees. In related work, Daniel Menascé, Daniel Barbará and Ronald Dodge presented an approach to automatically monitor and configure an e-commerce site to attain desired levels of QoS, using hill climbing to search for the optimal combination of configuration parameters. Moises Goldszmidt, Derek Palma, and Bikash Sabata introduced a metric, *effective eBusiness capacity*, for web servers, and presented an automatic method using techniques from machine learning and statistical pattern recognition to measure operational capacity.

Information and Price Dynamics

A number of participants presented research relating to information and price dynamics within agent-mediated marketplaces. Panos Markopoulos and Lyle Ungar presented a study of the value of *price information* for commodities in electronic marketplaces, with a model in which sellers charge for price information and rational buyers make decisions based on the expected value of additional price information and the cost of that information.

Chris Dellarocas presented an interesting equilibrium model of a reputation mechanism for use in a market in which the quality of goods is unknown to buyers until after a transaction, but sellers are rated by buyers after each transaction. Dellarocas considers a binary reputation mechanism in which ratings are either *positive*, *negative*, or *no-rating*, and presented in aggregation to buyers. The model considers the *leniency* of a buyer in rating a seller, and the *strictness* of a buyer in assessing the aggregate rating of a seller. A *well-functioning* reputation mechanism is defined as one in which there is a steady-state equilibrium in which sellers choose to truthfully reveal information about the quality of goods. Extended analysis, both theoretical and experimental, suggests that high rating leniency together with high assessment strictness provides a well-functioning mechanism with good stability properties.

Jeffrey Kephart, Chris Brooks and Rajarshi Das considered the dynamic optimization problem faced by a monopolist in an information goods market with shifting user demand. The study considers the effectiveness of dynamic price optimization for different levels of price complexity, ranging from *pure bundling* to *two-part prices* to *general non-linear prices*. The analysis establishes relationships between the complexity, learnability, and profitability of the various price schedules.

In related work, Robert Gazzale and Jeffrey MacKie-Mason presented a model of dynamic competition between two firms that compete in price and product configuration in an information market in which there is uncertainty about consumer preferences. Firms must compete for mass and niche markets and balance current profits with learning about the profitability of alternative strategies. Learning and uncertainty about consumer preferences leads to a higher product diversity and lower competition than in a model in which the firms are well-informed about customer preferences.

Finally, Philippe Golle, Kevin Leyton-Brown and Ilya Mironov proposed a number of simple mechanisms to provide incentives for sharing of information goods in peer-to-peer networks. The goal is to address the free-rider problem that can arise in systems such as Napster in which individual users are provided with no incentive for adding value to the network. Golle et al. analyze the game-theoretic equilibria in a system with heterogeneous agents, including some altruistic agents.

4. CONCLUSION

The ACM Electronic Commerce conference, sponsored by the ACM Special Interest Group on E-commerce (SIGecom), continues to attract high quality academic papers and serve as an important forum for the publication and dissemination of important new results in this emerging and exciting area of computer science that interfaces with economics, commerce, and law.

REFERENCES

EC 2001. *Proc. 3rd ACM Conf. on Electronic Commerce* (2001). ACM Press, New York.

Roth, A. E. and Ockenfels, A. 2001. Last-Minute bidding and the rules for ending second-price auctions: Evidence from eBay and Amazon auctions on the internet. *American Economic Review*. forthcoming.

Stone, P. and Greenwald, A. 2001. The first international trading agent competition: Autonomous bidding agents. *Journal of Electronic Commerce Research*. To appear.