Conference Report: The Sixth ACM Conference on Electronic Commerce

Jeffrey Shneidman
Division of Engineering and Applied Sciences
Harvard University

This report summarizes events and activities at the Sixth Association for Computing Machinery (ACM) Conference on Electronic Commerce (EC'05) held June 5-8 2005 in Vancouver, British Columbia, Canada. The program featured thirty-one talks, four tutorials, and two invited talks by Ehud Kalai and Jennifer Rexford.

Categories and Subject Descriptors: A.1 [**General Literature**]: Introductory and Survey; J.4 [**Computer Applications**]: Social and Behavioral Sciences—*Economics*

General Terms: Algorithms, Design, Economics, Experimentation, Human Factors, Theory

Additional Key Words and Phrases: ecommerce, sigecom, economics and computer science

1. CONFERENCE OVERVIEW

This paper reports on the Sixth ACM Conference on Electronic Commerce (EC'05), held June 5-8, 2005, in Vancouver, British Columbia, Canada. Four tutorials and a simultaneously run full-day Workshop on Sponsored Search Auctions preceded two and a half days of full paper presentations. There were 31 full paper presentations, representing papers drawn from 114 submissions. There was no poster or short paper session. For the first time in several years, the conference was not co-located with any other conference. Perhaps as a result, the conference was attended by approximately 70 people, a slightly lower figure than in previous years. However, as in previous years, the talks and discussions were lively. The dominant topics presented and discussed at the conference concerned mechanism design and related algorithmic issues. Other main topics of presentation included privacy, reputation and collaborative filtering, and empirical measurements / systems work.

2. TUTORIALS AND WORKSHOP

EC'05 participants were able to attend either the tutorial program or the workshop on sponsored search results on Sunday, June 5th. The tutorial program was divided into two sets of parallel sessions. In the morning tutorial set, Jason Hartline presented a tutorial on *Optimal Mechanism Design without Priors.* Michael Wellman surveyed the state-of-the-art in developing automated trading strategies in his

Address: 33 Oxford Street, Cambridge, MA 02138; jeffsh@eecs.harvard.edu

tutorial, *Trading Agent Design and Analysis.* In the afternoon tutorial set, Kamal Jain and Vijay Vazirani presented *Polynomial Time Algorithms for Market Equilibria*, while Tuomas Sandholm spoke on *Algorithms for Combinatorial Auctions and Exchanges.*

Running concurrently with these tutorials was the *First Workshop on Sponsored Search Auctions.* This full-day workshop contained ten short papers on mechanisms and behaviors in sponsored search. Pre-proceedings of this workshop are available online, linked from the EC'05 conference website.

3. PLENARY TALKS

What follows are this author's impressions of the paper presentations at the EC'05 conference. The papers and talks are presented in the order of presentation at the conference.

3.1 Day 1: Morning Session

The first morning started off with an energetic talk on *Ranking Systems: The PageRank Axioms*, by Altman and Tennenholtz. The authors consider web page ranking as a social choice problem, and present basic axiomatic properties that should be satisfied by page ranking systems. They apply a descriptive axiomatic approach to a simplified PageRank algorithm. The authors present five axioms, and show that with respect to these axioms, PageRank is sound and complete. Completeness is proven by showing that there is only one ranking system that satisfies these axioms, and that this system is PageRank.

Weak Monotonicity Suffices for Truthfulness on Convex Domains, by Saks and Yu, proves the conjecture by Lavi, Mu'alem, and Nisan that weak monotonicity is sufficient for truthfulness for functions defined on any convex domain. They do this by showing that in problems with a convex domain and finite range, the non-negativity two-cycle property implies the non-negative cycle property, which has been shown in previous work to be sufficient for showing truthfulness.

Marginal Contribution Nets: A Compact Representation Scheme for Coalitional Games, by Ieong and Shoham, presents the idea of Marginal Contribution Networks, where games are given as rules that map *conjunction* and *not* logic patterns to actual values. The value of some game outcome is computed as the sum of the values of the triggered rules. The authors present marginal contribution nets as fully expressive, potentially concise, efficient for calculating Shapley value and tractable in tree-width for Core problems, and having good usage simplicity.

The last talk of the first session was *Cost Sharing in a Job Scheduling Problem Using the Shapley Value*, by Mishra and Rangarajan, and was presented on behalf of the authors by David C. Parkes. This work considers a job allocation domain where jobs have a processing time and waiting cost. Transfers from a center are used to achieve fairness. The paper presents a Shapley value solution to this problem and contributes an axiomatic characterization of the Shapley value rule explaining how the Shapley value can be used in this domain.

3.2 Day 1: Midday Session

This session had a networking theme, and started with a talk on *Inefficiency in Provisioning Interconnected Communication Networks*, by Ferreira and Sirbu. In-

spired by previous work on Cost of Anarchy, this work focuses on interconnected networks and understanding the structure and bounds of routing inefficiency, considering how much the market can do for efficiency. The authors develop a model based on multi-commodity flows in which Internet Service Providers (ISP) play a non-cooperative game, and each ISP seeks to maximize profit. Unlike previous models, this work focuses on the presence and effect of Internet Service Providers (ISPs) and their role in routing traffic.

Hidden-Action in Multi-Hop Routing, by Feldman et al. looks at a traffic relay situation where *principals* send traffic to other principals via *agents*. Agents may choose to drop packets to conserve their own resources, and so these intermediate nodes must be incented to relay messages by receiving compensation. However, the relay/no-relay actions of these agents may be hidden from principals and links between nodes may be lossy so that relay actions are lost by the network. The work studied several problems in this domain, such as the effect of path length and the type of payment contracts between principals and agents. One interesting result of this work is that costly per-hop monitoring is actually not required to achieve good incentive properties, given some assumptions.

Content Availability, Pollution, and Poisoning in Peer-to-Peer File Sharing Networks, by Christin et al. seeks to measure the effect (in terms of availability vs. perceived availability) of junk files in the network, introduced either accidentally (pollution) or intentionally (poisoning). As possible motivation, this talk acknowledges that some organizations fight piracy by injecting garbage into file sharing networks. This work experiments with deployed peer-to-peer file sharing systems in an attempt to measure how different network types are resilient to file pollution.

The motivation for the work presented in *A Price-Anticipating Resource Allocation Mechanism for Distributed Shared Clusters*, by Feldman et al., is the perceived resource allocation problem in real distributed systems, such as on the Grid or Planetlab. The authors seek to build a distributed allocation mechanism that is economically efficient, envy-free, and has high agility in allocating resources as users change their demand. Their approach is a *fixed-budget resource allocation game*, where users place bids in a running series of simultaneous auctions, where each auction is responsible for a resource. This talk presented analytical and empirical simulation results, and pointed to ongoing work of a system deployment.

3.3 Day 1: Invited Talk by Ehud Kalai

Ehud Kalai gave an interesting and well-received invited talk[1] titled *Structural Robustness in Large Games*. Kalai is a Distinguished Professor of Decision and Game Sciences and Professor of Mathematics at the Kellogg School of Management. His talk presented the idea that in large games (with many semi-anonymous players), equilibria are structurally robust. In other words, with many players, equilibria survive major changes in the structure of the game. This work can be motivated by the issue that in reality, most games don't have formal rules that can be analyzed using game theory, and so it is nice if large games (unlike small games) have stability properties even in the absence of formal rules. One may wish equilibria to be robust to changes in structure, such as in order of play, ex-post stability, informa-

[1]Some of this talk was earlier published as *Large Robust Games* in Econometrica, 72(6), 2004.

tion leakage, or manipulations by insiders or outsiders. In these games, equilibria robustness is formalized by considering meta-games of a simultaneous move game and by requiring the simultaneous move equilibria to "survive" in all meta-games. The talk discussed these meta-games, and related problems, in more detail.

3.4 Day 1: Afternoon Session

Nearly Optimal Multi Attribute Auctions, by Ronen and Lehmann, presents a generic method (with many extensions) for the construction of nearly optimal multi-attribute auctions. This method always extracts at least half of the optimal expected utility and is almost optimal for independent cost functions. The computational time for the mechanisms equals the time required for computing the optimal mechanism on a small number of agents.

Optimal Design of English Auctions with Discrete Bid Levels, by David et al., muses on optimal auction design for an English auction when discrete bid levels, as opposed to continuous values, are used. The authors derive an expression that relates the expected revenue of an English auction to the discrete bid levels, number of participants, and distribution of bidders, and compare this result with previous continuous auction work.

Robust Solutions for Combinatorial Auctions, by Holland and O'Sullivan, is concerned with solution brittleness in combinatorial auctions. Specifically, the authors are worried about bid withdrawal, which can occur for both internal factors, such as winner's curse, strategic manipulation, or by external factors, such as a natural disaster. Recognizing that non-withdrawal rules may be unenforceable in practice, this work asks how one can repair a damaged combinatorial auction outcome. Their approach is to use constraint programming to find "close" robust solutions.

The last talk of the day was *Online Auctions with Re-usable Goods*, by Hajiaghayi et al. This talk presents a "what's new" in online auctions, and follows up on the work presented at EC'04. The general online mechanism problem has agents arriving and departing over some time interval. The mechanism makes allocation and payment decisions, without knowledge of future agent behaviors. This work adds re-usable goods to the online allocation problem, and provides new mechanisms that are constant competitive with respect to efficiency, and have new lower bounds with respect to revenue.

3.5 Day 2: Morning Session

The first two talks on Tuesday considered a routing problem where there are multiple competing paths from a source to a sink. Edges on these paths have costs for routing traffic and are asked by the mechanism to reveal their costs.

The first presentation on this subject was *First-Price Path Auctions*, by Immorlica et al. Unlike previous work that has focused on cost declaration truthfulness, this work focuses on reducing the payment that must be made when routing traffic on these graphs, using a first-price auction. The authors focus on strong epsilon-Nash and epsilon-Nash equilibrium concepts, and show that in strong epsilon-Nash, payments are often much less than Vickrey-Clarke-Groves (VCG) payments, while requiring strong player coordination. In epsilon-Nash, which doesn't require player coordination, there is a randomized mechanism that has provably low payments.

The second talk using this network model, *True Costs of Cheap Labor Are Hard*

To Measure: Edge Deletion and VCG Payments in Graphs, by Elkind, considers reducing routing payments by simply deleting edges in a communication graph to reduce the Vickrey-Clarke-Groves (VCG) payment. But which edges should one delete? The paper shows that finding the best set of edges to delete is NP-hard and not able to be approximated. The author explores series-parallel graphs as a special case of restricted graphs, and shows that even in this special case where edge costs are 0/1, finding the set of edges to delete is hard.

The third talk, *From Optimal Limited Supply Auctions to Unlimited Supply Actions* by McGrew and Hartline, seeks a profit maximizing auction (using competitive ratio as a metric) in the worst case for limited/unlimited supply auctions. The authors show a non-trivial limited supply optimal auction, and then show a general way to turn limited supply auctions into unlimited supply auctions. When this transformation is applied to their 3-bidder limited supply auction, they find an unlimited supply auction which bests the previously known competitive ratio.

The last talk of the morning was *Multi-unit auctions with budget-constrained bidders* by Borgs et al. This work characterized truthful mechanisms for multi-unit auctions with private valuations and bidder budgets. The authors prove that there is no non-trivial truthful auction that can allocate all goods, but introduce a revenue-maximizing truthful mechanism that allocates some units. The talk argued that two properties, *consumer sovereignty* and *independence of irrelevant alternatives*, are natural when thinking about multi-unit auctions with budget constraints.

3.6 Day 2: Midday Session

The midday session on the second day had a more applied feeling. The talks in this session were very different in spirit from most of the other sessions in that they focused on "real" and perceived e-commerce problems and possibilities.

The first talk was *Graceful Service Degradation (or, How to Know your Payment is Late)* by Andoni and Staddon. This work takes a cryptographic approach to graceful degradation of quality of service for users that are a "little late" on payment. Inspired by hard functions for spam-fighting, the authors propose using variably hard functions to increase the amount of computation of a problem that a "pay-late" user must experience. Their idea is to encrypt content (for instance, a movie from an online video rental service) and then disseminate varying levels of helpful hints to users appropriate to their various levels of arrears with the help of degradation protocols.

Privacy-Preserving Credit Checking by Frikken et al. presents a scheme for solving the problem of keeping information (belonging to both a consumer and a bank) private from each other and from a third-party credit reporting agency (CRA), when attempting to verify whether or not a consumer can qualify for a bank loan. Private information in this case are details of the consumer, and the loan approval algorithm used by the bank. The authors' approach is based on Scrambled Circuit Evaluation. The talk presented the protocol and a number of extensions that were mainly geared toward making this type of protocol more practical.

Dynamic and Secure B2B E-Contract Update Management, by Angelov et al., focuses on automated electronic contracts, where the contracts are established and fulfilled automatically based on precise contract descriptions. This type of contract is not currently used in practice, but this type of contract could be interesting

in areas like online advertising. These contracts may need to be changed, and the authors discuss two protocols that can be used by companies to update these electronic contracts.

Secure Distributed Human Computation, by Gentry et al., is an idea-piece (read: interesting open questions and few answers) that argues that there are many cases where humans and computers could and should work together. For instance, computers are poor at doing image analysis, but humans can label images with probable search terms very efficiently. As a result, some online image search databases actually use humans to label images. This paper is a call for researchers to find ways to combine human and computer labor to solve interesting tasks that would be difficult for either a computer or human working alone. The talk (and the paper) presents a number of examples and open problems related to this idea.

3.7 Day 2: Invited Talk by Jennifer Rexford

The second day's invited talk was given by Jennifer Rexford. She has recently joined the faculty at Princeton, and was asked to speak to EC'05 not only because of her strong theoretical grounding, but also because of her practical exposure in her past experience with AT&T networking. Her talk was titled *Economic Incentives in Internet Routing*. After providing a brief architectural overview of how Internet routing works, she focused on convergence problems within the Border Gateway Protocol (BGP). Her talk argued that some of the convergence problems that have been worried about in the past are not realistic. Her talk was full of interesting lessons about how autonomous-system-level inter-domain networking is really done. She presented a set of BGP policies that can be proven to avoid oscillation. Her idea is to use black boxes that enforce these rules and relay the decisions of the mechanism to the actual BGP routers. Rexford closed off the talk by discussing open problems in economic incentives in inter-domain routing. Examples included: Can one model how nodes select a peer or a provider? How can one study pricing games where agents manipulate their declared ability to route traffic based on billing policies? How can one create mechanisms to avoid "Hot Potato" routing?

3.8 Day 2: Afternoon Sessions

The first paper of the afternoon session, *Communication Complexity of Common Voting Rules*, by Conitzer and Sandholm, presented eleven voting rules and the communication complexity of each rule. For each voting rule, an upper bound is given on a deterministic protocol for that rule, and a lower bound is given for both deterministic and nondeterministic versions of the protocol.

On the Computational Power of Iterative Auctions: Demand Queries and Ascending Auctions, by Blumrosen and Nisan, explored the power (and the limitations) of iterative auctions. This work contains several interesting results for auctions that repeatedly interrogate bidders for their demand on a changing set of posted prices.

The third talk, *Fairness and Optimality in Congestion Games*, by Chakrabarty et al., considers two problems in congestion games: finding a fair allocation, and computing a social optimum. The authors give simplification algorithms for an otherwise inapproximable problem. The talk, and the paper, focus on hardness results and approximate solutions.

The last talk of the day, *Congestion Games with Failures - CGFs*, by Penn et al., introduces a new class of congestion game that allows for failures by the service provider support. In this environment, an agent may rely on a provider that fails with some known probability. The authors show that basic CGFs always have a pure-strategy Nash equilibrium, and explore various properties of the CGFs, both in a basic form, and an extended model that allows for a cost when submitting jobs to service providers.

3.9 Day 3: Morning Sessions

ICE: An Iterative Combinatorial Exchange, by Parkes et al., presents the first design of a fully expressive iterative combinatorial exchange. ICE can allocate resources among mixed buyers and sellers (i.e., users that simultaneously buy or sell bundles of items). The iterative nature allows a user to be inexact in their initial valuation (often desirable in the real world) and reduces the communication and computation complexity of the mechanism. All features in this paper have been fully implemented in software.

Self-Selection, Salvaging, Slacking, and Stoning: the Impacts of Negative Feedback at eBay, by Khopkar et al., set out to measure the various effects of negative feedback left by users on eBay using a wealth of usage data from the online auction-house. The authors performed a statistical analysis in an attempt to measure the effects of three forms of user behavior with respect to negative feedback.

Information Markets vs. Opinion Pools: An Empirical Comparison, by Chen et al., makes the claim that predictive information markets predict events as accurately as pooled expert assessments. The authors use historical data from the 2003 US National Football League to support this claim using straightforward statistical analysis.

Integrating Tradeoff Support in Product Search Tools for E-Commerce Sites, by Pu and Chen, presents a user study in the context of decision support. Their talk and paper describes a series of experiments to determine how users go about making and remaining confident about decisions.

3.10 Day 3: Midday Sessions

The last session of the conference contained two talks. The third scheduled paper, *Towards Truthful Mechanisms for Binary Demand Games: A General Framework* by Kao et al., was not presented.

The first talk in the session, on the *Complexity of (Iterated) Dominance*, by Conitzer and Sandholm, considers several computational aspects of solving games using dominance and iterated dominance. The paper considers this problem in both normal form and Bayesian games. The talk and paper focus on the complexity of showing dominance.

The last talk of the conference was *On Decentralized Incentive Compatible Mechanisms for Partially Informed Environments* by Mu'alem. This work suggests a notion of *partially informed* agents. Using the context of peer-to-peer file sharing with no external failures, Mu'alem considers how being partially informed drives agent behavior on whether or not to provide files to others. In the partially informed model, cooperation is achieved independent of agents' beliefs.