

Conference Report: The Fifth ACM Conference on Electronic Commerce

David M. Pennock
Yahoo! Research Labs

This report summarizes events and activities at the Fifth Association for Computing Machinery (ACM) Conference on Electronic Commerce (EC'04), held May 17–20 2004 in New York City. The program featured twenty-four talks, twenty-one brief announcements, three tutorials, an award for Best Student Paper, a joint session with the Thirteenth International World Wide Web conference, and a joint invited keynote speech by Udi Manber.

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, Association for Computing Machinery (ACM) Conference on Electronic Commerce (EC'04) report and summary, Special Interest Group on Electronic Commerce (SIGecom)

1. CONFERENCE OVERVIEW

For the fifth time since its founding in 1999, the ACM Special Interest Group on Electronic Commerce (SIGecom) held its flagship event: the Conference on Electronic Commerce¹ (EC'04). The conference, held at the Sheraton hotel in New York City May 17–20 2004, brought together a variety of researchers exploring computational aspects of commerce and economics. EC'04 drew people from a number of fields, including artificial intelligence, auctions, cryptography, decision theory, distributed systems, e-business, experimental economics, game theory, human-computer interaction, finance, machine learning, psychology, and social network analysis. Participants represented both academia and industry from a number of countries around the world. EC'04 was co-located with the Thirteenth International World Wide Web conference² (WWW2004), and several joint activities were

¹<http://www.research.microsoft.com/acmec04/>

²<http://www.www2004.org/>

Address: 74 N. Pasadena Ave., 3rd Floor, Pasadena, CA 91103; pennockd@yahoo-inc.com

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works, requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept, ACM Inc., 1515 Broadway, New York, NY 10036 USA, fax +1 (212) 869-0481, or permissions@acm.org.

held. WWW suspended its usual ecommerce track for 2004, encouraging authors instead to submit ecommerce-related papers to EC'04. A joint technical session was held allowing participation from both conferences' registrants. Finally, there was a joint invited keynote speech by Udi Manber, CEO of A9.com, a search startup wholly owned by Amazon.com.

This report is a summary of my perspective on the conference: the report's contents (and any omissions) reflect the subjective and at-times arbitrary nature of my attendance, attentiveness, memory, and opinion, as well as the writing process itself; my apologies for anything missing or misrepresented in what follows.

EC'04 received 146 submissions, one of the largest numbers of submissions since the conference's inception. Submissions were examined through a peer review process under the direction of the program co-chairs, Joan Feigenbaum and Margo Seltzer. The EC'04 general chair was Jack Breese. The tutorial chair was David Parkes. The program committee responsible for peer reviews included Ross Anderson, Yair Bartal, Martin Bichler, Aaron Brown, Chris Dellarocas, Amos Fiat, Lance Fortnow, Nick Jennings, Michael Littman, Mark Manasse, Moni Naor, David Pennock, Mema Roussopoulos, Subhash Suri, Jim Waldo, and Tim van Zandt. A number of external reviewers assisted the program committee; the full list appears in the conference proceedings [Feigenbaum and Seltzer 2004]. The chair of SIGecom is Michael Wellman. Of the 146 submissions, 24 were accepted for full length presentation and 21 for short presentations. As a reviewer, I was struck by the quality of many submission, including many that ultimately were not accepted. Accepted papers covered a spectrum of ecommerce topics including auctions, automated learning, competitive analysis, cryptography, decision theory, experimental economics, game theory, interfaces, information markets, matching algorithms, mechanism design, mixed Bayesian worst-case analysis, negotiation, online algorithms, peer-to-peer (P2P) systems, preference elicitation, privacy, recommender systems, reputation systems, spam protection, and trading agents. Ryan Porter won the Best Student Paper award for his paper titled "Mechanism design for online real-time scheduling".

Section 2 summarizes the tutorial program, Section 3 reviews the technical program, and Section 4 describes Manber's invited talk. Section 5 ends with some concluding remarks.

2. TUTORIALS

Three tutorials were offered at EC'04: (1) *Market Clearing Algorithms* by Tuomas Sandholm, Computer Science Professor at Carnegie Mellon University; (2) *E-Commerce Applications of Semantic Web Services* by Benjamin Grosz, Information Technology Professor at MIT Sloan School of Management; and (3) *The Price of Anarchy* delivered by Tim Roughgarden, Computer Science Professor at Stanford, and prepared in conjunction with Christos Papadimitriou, Computer Science Professor at UC Berkeley. Tutorials (1) and (2) were offered in parallel.

Sandholm's tutorial covered a huge array of economic matching algorithms for allocating resources, from the canonical English "open outcry" auction for selling a single good, to advanced combinatorial auctions allowing expressive bids over bundles of multiple goods. After reviewing results from auction theory in economics, he moved on to problems where computation complexity and algorithm design

become of paramount concern, covering a number of fairly recent results, including some from his own group, one of the pioneering groups in this burgeoning field.

Grosov highlighted the synergistic rise of two potential “game changers” on the web: semantic markup languages and web services. Web authors have so far been slow in adopting extensive or complicated markup languages. But with demand high and rising for web services (e.g., RSS, Google API, etc.)—and with some of the most creative services being developed by third parties—the value of expressive markups is becoming clearer to many content providers. Automated or semi-automated tagging mechanisms are also driving growth. Grosov surveyed the latest advances in this area, covering a number of ecommerce applications including negotiation, finance, and security/privacy. He highlighted business-to-business ecommerce as an area of enormous potential and a likely playground for early adopters.

Papadimitriou and Roughgarden’s tutorial centered around a natural measure of the (always nonnegative) difference between the value of a centralized and decentralized solution to a noncooperative game. They call the measure the “price of anarchy”, and define it as the worst-case difference in social welfare between the absolute global optimum and the equilibrium solution in the game. Of course the measure varies depending how social welfare and equilibrium are defined. The tutorial described new techniques for bounding the price of anarchy in various mechanisms associated with a number of application areas.

3. TECHNICAL SESSIONS

The EC’04 technical program comprised twenty-four presentations and twenty-nine brief communications reporting on recent research results across a range of topics in ecommerce. Corresponding written reports appear in the archival proceedings [Feigenbaum and Seltzer 2004]. I attempt to summarize and highlight some of the talks here, organized by session.

Gonen et al. [p. 1] provide an ingenious mechanism that, in a sense, circumvents the famous Myerson-Satterthwaite (M-S) impossibility theorem. The mechanism employs a bi-partite graph matching algorithm and does achieve all four properties that M-S proved were mutually incompatible: budget balance, efficiency, individual rationality, and incentive compatibility. The catch is that the mechanism returns upper and lower ranges for bilateral trade, rather than point prices, leaving it up to the two parties to engage in final negotiation however they see fit.

Babaioff et al. [p. 9] examines the problem of auctioning a good in independent spatially-distributed markets with transportation costs, where buyers and sellers are strategic. They show that prices can be seen as distances in a “residual flow” graph representation, and provide two mechanisms: one that is efficient, incentive compatible, and individually rational, and one that is budget balanced, incentive compatible, and individually rational (the M-S theorem ensures that no mechanism can have all four properties).

Acquisti [p. 21] takes a behavioral economics approach to analyzing privacy, presenting a model to explain extensive empirical data showing that people do not act to protect their privacy in proportion to the amount they express worry about privacy violations. He explains that part of the problem is that the benefits of revealing private information are immediate while the drawbacks, though potentially

severe, are generally long term and of low probability.

Goodman and Rounthwaite [p. 30] explore mechanisms to control spam originating from a free email service like Hotmail.com. They show that, surprisingly, imposing sender costs only during an initial period works almost as well as perpetual fees. They note that costs need not take the form of actual currency: they could involve barriers like solving *human interactive proofs* (HIPs, a.k.a. CAPTCHAs) or computational costs, and costs can be imposed in proportion to a user’s spam filter rating. The authors analyze and explain the economics of spam, and the costs of for example outsourcing the solution of HIPs to low wage countries, and calculate an approximate threshold where untargeted spamming becomes economically unsustainable. They show that complaint feedback is key, and discuss ways to encourage legitimate complaints.

Loder et al. [p. 40] also analyze the economics of spam, and argue that even perfect filters are socially nonoptimal. The reason is that some unsolicited communication is mutually desirable if it’s combined with a sufficient payment from the sender to the receiver. The authors analyze an *attention bond mechanism* where the sender sets aside an amount of money per message (a signal of the sender’s value of the communication) and the receiver has the option of seizing the bond. They solve for equilibrium solutions in several cases, showing how the mechanism outperforms a perfect filter in terms of total social welfare.

Conitzer and Sandholm [p. 51] explore novel mechanisms for allocating charitable donations that in a sense generalize the familiar “matching pledge”, allowing more expressive coordination among donors. The problem is to produce socially optimal allocations given a number of pledges that each designate an amount willing to be contributed as a function of the vector of total donations to all charities. The authors show the problem to be NP-hard and present heuristic solutions and restricted settings with polynomial solutions.

Ryan Porter’s work on decentralized online scheduling earned him the conference’s Best Student Paper award. Porter [p. 61] proves the first competitive bounds for online scheduling where incentives are taken seriously: that is, the agents performing tasks are self-interested and may misrepresent their capabilities or delay their responses if it helps increase their own expected utility. Surprisingly, Porter provides an incentive compatible mechanism that increases the competitive ratio by only one, as compared to the optimal non-strategic bound. The work should serve as a valuable template as more and more problems of distributed algorithm design are recognized naturally as problems of mechanism design where understanding incentives is of paramount importance.

Hajiaghayi et al. [p. 71] generalize a classic result known as the *secretary problem*³ to the strategic setting in the form of an online limited-supply auctioneer’s problem. They prove several competitive ratio bounds for algorithms that are strategy proof in terms of both value and time, in the process demonstrating a fairly generic technique for proving temporal strategyproofness. Their adversary model is worst-case under an independence assumption, a natural form of mixed Bayesian worst-

³The secretary problem is the online problem of finding the maximum in a sequence (e.g., the best secretary via a series of interviews); Gilbert and Mosteller prove the startling result that a simple selection algorithm gives the maximum with constant $(1/e)$ probability.

case analysis.

Singh et al. [p. 81] generalize some promising work on graphical games—which in certain circumstances provide huge gains in representational conciseness and associated solution complexity—to the case of incomplete information.

Marti and Garcia-Molina [p. 91] examine how limited reputation systems can help address problems like freeriding and spamming in peer-to-peer (P2P) systems.

Feldman et al. [p. 102] also tackle the problem of freeriding in P2P networks, which they show is exacerbated by a number of factors endemic to the P2P setting. The authors model the problem as a generalized prisoner dilemma game and show in simulation how certain policies can help.

Awerbuch et al. [p. 112] formulate a general framework for studying reputation systems, applicable in a number of domains included P2P systems, auctions (e.g., Ebay), product reviews (e.g., Amazon), merchant reviews (e.g., Shopping.com), and even email. They propose a method to balance exploration (random selection) and exploitation (reputation-biased selection) and show its benefits in theory and in simulation.

Mehta and Vazirani [p. 120] prove a surprising result in the context of auctioning a digital good (unlimited supply). The authors show that, for every mechanism that is (ex ante) incentive compatible in expectation, there exists an equivalent auction that a randomization over incentive compatible deterministic auctions, where truth-telling is weakly dominant even after all randomization is revealed. The proof is constructive and the authors prove conditions under which the constructed truthful auction is ϵ -equivalent to the original and runs in polynomial time.

Lipton et al. [p. 125] prove a number of complexity results, both positive and negative, for the problem of fair (envy free) division of discrete indivisible goods. The authors give (1) an absolute bound with polynomial solution that substantially improves a previous bound, (2) show that minimizing envy and approximating minimum envy are computationally hard, (3) show that approximating the envy *ratio* is possible in polynomial time, and (4) perhaps most surprisingly show that no minimum envy algorithm can be truthful.

Conitzer and Sandholm [p. 132] extend their recent work on automated mechanism design—where the designer solves for the optimal mechanism at “run time” specifically to address a particular scenario—to the case of a *self-interested* designer wishing to maximize its own utility rather than some social welfare objective.

Kim and Karp [p. 142] describe the workings of HP’s e-speak platform for web services in detail, where the goal is to build an expressive ontology to support searching and browsing a catalog of products, resources, or services. Among a number of nice features, the platform supports configurable security settings so that searchers can only see the availability and description of services they are authorized to see.

Benisch et al. [p. 152] present methods for handling uncertainty in a trading scenario, specifically exploring the scheduling component of the Trading Agent Competition in Supply Chain Management (TAC SCM). In extensive experiments using TAC SCM data, they show that the expected value method (converting the problem to a deterministic one by taking expectations) is inferior to a greedy sample average approximation (SAA), which in turn is inferior to a lookahead SAA. The benefits of lookahead are shown to increase quickly over time.

Kirovski [p. 160] examines algorithms to strengthen the security of a fascinating physical authentication method called a certificate of authenticity (COA). A COA is a cheap physical object whose random structure is difficult and expensive to replicate; Kirovski explores a COA constructed by randomly placing fiber optics in a square. The COA's signature consists of the pattern of lights that results from placing a source light in certain fixed positions. Authentication involves matching the COA's pattern of lights to an encrypted encoding of the pattern placed in an associated bar code or RFID tag. Kirovski shows that improving the compression ratio of the encoding increases the cost of forgery exponentially. Though he proves that optimal compression is NP-hard, he shows that his heuristic compression algorithm tailored to the setting outperforms standard techniques.

Pennock [p. 170] describes a new mechanism supporting speculation (wagering) on an event, and hedging or risk sharing. The mechanism can be thought of as a cross between a continuous double auction and a pari-mutuel market, combining desirable properties of both. The mechanism has infinite (one sided) liquidity, supports dynamic response to information, and guarantees no risk to the auctioneer.

Lahaie and Parkes [p. 180] make precise an analogy between value and demand queries in preference elicitation and membership and equivalence queries in function learning theory. Since preference elicitation need not learn all agents' preferences exactly (as long as an equilibrium is found), the two problems are not equivalent, and some previous negative results can be sidestepped. The authors prove that polynomial queries and communications are sufficient to obtain the optimal allocation in some fairly general settings.

Kakade et al. [p. 189] derive online algorithms for executing financial transactions where the goal is to obtain prices as close as possible (in the worst case) to the volume-weighted average price or (alternatively) the maximum price. The authors derive competitive bounds within several natural models, each increasingly realistic, including models with access to the full order book and that account for distributional assumptions.

Spiekermann [p. 200] reports on empirical research backing the hypothesis that consumers prefer different contexts for different products. For example, products with higher purchase uncertainty benefit from more detailed context, trend- or fashion-related products benefit from more emotional context, and technological products benefit from more feature-oriented context.

Pu and Kumar [p. 208] describe and evaluate an example-based search tool called SmartClient, which goes beyond the standard one dimensional ranked list to allow multi-dimensional tradeoff exploration and relevance feedback. They show that SmartClient performs as well as ranked lists in simple settings, while significantly reducing search time and error rate in more complex and realistic scenarios.

Wang et al. [p. 218] find that clustering ecommerce customers according to three metrics—services requested, navigation patterns, and server resource usage, corresponding to three common site optimizations—all yield roughly equivalent groupings, suggesting that similar customers remain similar across all three dimensions. The authors then argue for clustering by services requested, since this yields the cleanest results, is simplest to implement, and seems sufficient for all three purposes.

Many of the brief announcements were interesting and impressive, at times on par with or even exceeding the quality of full papers. I briefly mention only a

few. Tennenholtz extends the classic setting of static voting to a dynamic setting that captures many aspects of ranking algorithms used on the web (e.g., in web search). Garcia et al. describe a promising reputation mechanism that takes incentives seriously. Chen et al. describe a laboratory economics experiment showing that reputation information makes a difference in market efficiency, and that self-reported ratings are nearly as good as complete historical information. Hogg and Adamic argue for the use of social networks to validate or augment reputation information. Smith et al. describe a system where consumers can scan a product's bar code using a mobile device and obtain related information from the web. Shoham and Tennenholtz coin the phrase *behavioral mechanism design* to incorporate many intangible incentives that affect people's interactions with online auctions and other mechanisms, like entertainment value, competitive spirit, and interactivity, as well as the designer's objectives, like buzz, marketing value, and stickiness. Lam describes a test of the hypothesis that social network information should improve a recommender system. O'Mahony describes a promising new notion of "neighborhood" for use in recommender systems. Conitzer and Sandholm provide compelling arguments that, once computational complexity is taken into consideration, care must be taken in invoking the classic *revelation principle*. Larson and Sandholm prove a strong impossibility result for auctions where agent preferences require some computation to fully resolve.

4. KEYNOTE SPEECH

Udi Manber delivered the conference's keynote address to a combined EC-WWW audience. Manber moved from academia as a U. Arizona professor into industrial research as the Chief Scientist at Yahoo!, then the Chief Algorithms Officer at Amazon.com, and now the CEO of Amazon.com subsidiary A9.com. Manber and A9.com received something of a blank (or at least large) check and an open mandate from Amazon to build the infrastructure and expertise necessary to pursue their goal of, simply stated, "inventing new things in search". Manber expressed a desire not to simply "build another Google", but to expand into new areas, algorithms, and applications. A9's shorter term strategy is to focus on ecommerce search. One of the first of A9's products is the heavily publicized (and widely praised) "search inside the book" service. In what Manber described as "the most exciting project" of his career, he orchestrated cooperation among hundreds of publishers to allow Amazon to scan and OCR hundreds of thousands of books, allow users to search through the full text of the books, and allow limited viewing of the scanned pages. Designing a system that disables large-scale copyright infringement—and perhaps more importantly convincing publishers of its efficacy—was no easy task.⁴ Manber emphasized the importance of being "customer centric", and illustrated A9's efforts along these lines in terms of interface design and personalization. For example, A9 users can ask the site to remember their page layout customizations, remember past searches, and remember pages previously clicked on. A9 toolbar users can record their browsing and searching history and store annotations for any

⁴Manber illustrated one of the "bugs" that can arise in the midst of a massive and semi-automated effort to scan in millions of pages from thousands of books: as identified by an Amazon.com user, one page's digital image contained the silhouette of a squashed fly.

web page, all on the A9 server for access from any computer. More paranoid users can search anonymously at `generic.a9.com`. Beyond a peek inside A9 and a bit of salesmanship, Manber offered both a vision and some advice for fellow visionaries. He believes that in the next ten years, search technology is in for a revolution comparable to the size and scale of the Internet/Web revolution of the past ten years. He believes current solutions have barely scratched the surface of what is possible for algorithms, interfaces, and relevancy. In particular, today's focus on short one-shot queries will evolve into longer descriptive queries with interactive relevance feedback, only possible once algorithms are good enough to provide users sufficient benefit for their additional effort. Manber ended with a number of "what if" scenarios: thought experiments where he encouraged the audience to "look at problems backward", imagining the solution first, then figuring a way to get there. What if you had unlimited computation, storage, and bandwidth? What services could you provide? What if computers could judge relevance with machine speed and human accuracy? By all accounts, Manber's talk was well delivered, well motivated, entertaining, and thought provoking, capping a very productive and interesting EC'04 conference.⁵

5. CONCLUSION

EC'04 served as an excellent venue for reporting research results in electronic commerce, fostering communication among like-minded researchers in a way that (still today) only a physical event can. Cooperation with the World Wide Web conference encouraged cross fertilization, with many WWW registrants attending the joint session and nearly all attending Manber's talk showcasing A9.com as a successful merger between electronic commerce and search.

EC'04 papers and participants spanned a number of disciplines, with a bias toward researchers in computer science theory and artificial intelligence, many focusing on mechanism design. Although this core in many ways represents a healthy broadening from an initial emphasis in 1999 and 2000 on auctions, and in particular on combinatorial auctions, a continuing goal of conference organizers is to broaden the scope of the conference further while retaining its core areas of strength. At the conference business meeting, many agreed that striking the right balance between breadth and depth is difficult, especially given the self-reinforcing nature of expectations among contributors. Still, it seemed that the general sentiment among conference organizers was to strive for greater breadth in the coming years.

At the close of EC'04, the organizing committee for the 2005 conference was announced. John Riedl will serve as general chair, with Michael Kearns and Michael Reiter serving as program co-chairs. The 2005 organizing committee, as well as SIGecom Chair Michael Wellman, welcome any suggestions for the 2005 conference from all interested readers.

ACKNOWLEDGMENTS

Thanks to Varsha Dani for helpful comments.

⁵A popular press account of Manber's talk can be found at <http://www.internetnews.com/ent-news/article.php/3356831>

REFERENCES

- FEIGENBAUM, J. AND SELTZER, M. Eds. 2004. *Proceedings of the Fifth ACM Conference on Electronic Commerce (EC'04)*. ACM Press, New York.