Conference Report: The Second ACM Conference on Electronic Commerce

David M. Pennock
NEC Research Institute

The second Association for Computing Machinery (ACM) Conference on Electronic Commerce (EC'00)—held October 17–20, 2000, in Minneapolis, Minnesota, USA—featured twenty-nine talks chronicling the latest progress on research at the boundary between computers and commerce. Other conference highlights included four tutorials covering ecommerce foundations, recommender systems, scaling for ebusiness, and mobile ecommerce; a keynote address by Robert Weber; and a panel discussion on the results of a recent trading agent competition.

1. CONFERENCE OVERVIEW

On October 17–20, 2000, one hundred twelve researchers gathered at the Hilton hotel in downtown Minneapolis for the second Association for Computing Machinery (ACM) Conference on Electronic Commerce (EC'00). Experts in computer science, economics, management, policy, and law, from academic, industrial, and government institutions in Brazil, Germany, Israel, Japan, Sweden, the United States, and even New Jersey, convened to discuss scientific and technological progress in ecommerce research.

The EC conference is the centerpiece event of the ACM Special Interest Group on Electronic Commerce (SIGecom). SIGecom was founded on January 11, 1999; since its inception, membership has grown dramatically,[1] prompting ACM to grant

---

[1] Membership has grown from three in February 1999 to six hundred ninety three in November 2000. Month by month statistics are available at `http://www.acm.org/sigs/membership/`.

Address: 4 Independence Way, Princeton, NJ 08540; `dpennock@research.nj.nec.com`

the SIG permanent status is record time.

Submissions to EC'00 roughly doubled as compared to the first conference in 1999, and only 19% of submissions were accepted for publication. Submissions were examined through a peer review process under the direction of the conference committee. The EC'00 general chair was Anant Jhingran; the program co-chairs were Jeff MacKie Mason and Doug Tygar; the publicity chair was Bill Rubin; the tutorial chair was Pattie Maes; the treasurer was Mark Scott Johnson; the program committee included Mark Ackerman, Ross Anderson, Jack Breese, Karen Clay, Michael Froomkin, Marti Hearst, Manoj Kumar, Larry Lessig, Pattie Maes, Paul Resnick, John Riedl, Yoav Shoham, Gene Tsudik, Michael Wellman, and Peter Wurman. Accepted papers covered a wide spectrum of ecommerce topics, ranging from combinatorial auctions to computational complexity, from interfaces to information goods, from privacy to personalization, from reputation systems to recommender systems, and from trading agents to targeted advertising.[2]

The EC'00 program included four tutorials, a keynote address, a panel discussion, and the main technical program of contributed papers. Section 2 describes the tutorials, keynote, and panel, and Section 3 reviews the technical program in detail.

2. TUTORIALS, KEYNOTE, AND PANEL DISCUSSION

Four tutorials were offered at EC'00, representing a broad sample of theoretical and practical topics in ecommerce. In "Foundations of electronic marketplaces", Dr. Tuomas Sandholm, Professor of Computer Science at Washington University in St. Louis, reviewed the most relevant aspects of microeconomic theory, game theory, voting, auction theory, and coalition formation applicable to electronic marketplaces and multiagent systems. He discussed when designers of agents and mechanisms should consider a multiagent perspective, which theories and algorithms are likely to be most useful, and what effect real-world constraints (e.g., bounded rationality) have on practical implementations. Dr. Joseph Konstan and Dr. John Riedl, Professors of Computer Science at the University of Minnesota and founders of Net Perceptions,[3] the leading commercial vendor of recommender systems, led a tutorial on "Recommender systems in e-commerce: How collaborative filtering is helping businesses serve customers better". They outlined the history and use of recommender systems to automatically suggest products to buyers, reviewing today's state-of-the-art algorithms, and touching on promising future developments that integrate collaborative techniques with more traditional methods in machine

---

[2]In addition to a diverse collection of topics, ACM SIGecom's EC'00 featured WAABNO (a Wide Array of Acronyms, Both New and Old), including LP, VCG, MUCAS, WSP, CAMUS, AkBA, PCA, GVA, SAA, AUSM, RAD, DARPA, IBM, LDS, IDA*, CASS, FCC, VSA, HTTP, HTML, NCSA, URL, IP, IT, ISBN, SKU, WWW, PC, VTP, CA, VCR, CATS, US, CDF, U, LPM, MARI, MIT, UIM, BVM, SVM, B/SVM, VFG, VFT, MCM, OHM, MM, USM, UNM, DM, AUL, TPL, SQL, DTD, XML, AP, UPI, AOL, VHS, API, B2C, CD, MSNBC, ACM, KDD, CF, SVD, ML, EC, IR, P3P, GCE, CGI, PAC, DNS, PDB, DAG, B2B, WES, eGS, HR, CM, PG, CC, CSL, PC, PM, FP, CCL, CMG, QH, LFD, MLD, CQ, JDK, RAM, ART, EPC, HI, IZ, T, CSUQ, GMD, GMD-FIT, CORBA, EJB, RMI, GUI, OBI, JETS, CNET, RFB, MAS, RCV, EH, MTBF, SAM, CMAP, SAMP-SB, and SAMP-SB-D. Two papers tied for the most acronyms (19) [Tewari and Maes, p. 86; VanderMeer, Dutta, and Datta, p. 185]. Interestingly, the former's institution is itself an acronym.

[3]http://www.netperceptions.com/

learning and information retrieval. Dr. Daniel A. Menasce, Professor of Computer Science at George Mason University, and Dr. Virgilio A. F. Almeida, Professor of Computer Science at the Federal University of Minas Gerais in Brazil, presented "Scaling for e-business: Technologies, models, performance and capacity planning". With low overhead and worldwide accessibility, ebusinesses have the potential for sudden demand and rapid growth. Menasce and Almeida described models of typical ebusinesses and their customers, and discussed how ebusinesses can prepare for uneven and accelerated demand. Dr. Anupam Joshi, Professor of Computer Science at the University of Maryland, Baltimore County, spoke on "Mobile e-commerce: Challenges and solutions". The talk covered the basics of mobile computing, from hardware to applications, focusing on the promise and pitfalls of ecommerce applications. In order to support robust search, payment, and content delivery services, developers must address many challenging problems, including recovery from disconnection, authentication, multimedia synchronization, and heterogeneous system interoperability. All four tutorials were well attended, and feedback was overwhelmingly positive.

Dr. Robert Weber, Professor of Decision Sciences at the Kellogg Graduate School of Management at Northwestern University, opened the EC'00 technical program with an invited keynote address on the theory and practice of auctions. In the first half of his talk, Weber reviewed some of the key results in auction theory. For example, under simplifying assumptions, auction sellers can expect to obtain equivalent revenue regardless of which of the major auction types they utilize. When buyers compete for a good of uncertain value, they must contend with the so-called *winner's curse*: the very act of winning suggests that the good is worth less than the winner expected, since everyone else apparently values it less. According to the *linkage principle*, full revelation of information helps alleviate the winner's curse by reducing uncertainty. Moreover, in repeated-game situations, disclosing all relevant information is in the best interest of the seller as well. The *generalized Vickrey auction* (GVA) is a mechanism for allocating bundles of goods in arbitrary combinations. Although the GVA is Pareto optimal and incentive compatible, it does suffer from some counterintuitive properties: for example, it is possible that winning bidders actually pay less even when all bidders offer more.

The second half of Weber's talk concerned real-world auctions, emphasizing when practical experience differs from theoretical predictions. For example, although Dutch auctions are theoretically equivalent to first-price sealed-bid auctions, the former tends to yield higher revenues empirically, when bidders "panic" and jump in too quickly. In 1981, Sotheby's auctioned off seven identical leases for RCA transponders. Because the bidders did not fully think through their strategies, each successive lease sold for less than the previous, in direct opposition to what theory predicts. Buyers were outraged at paying vastly different prices for identical leases, and the FCC actually invalidated the results of the auction. Weber relayed a fascinating account of his involvement in the 1994 FCC auction of broadband PCS licenses. Although prior agreements and direct communication among the bidders were prohibited, many bidders engaged in complex negotiations by signaling their intent through their bids. Some even encoded messages in the lower order digits of their bids. Consulting for American Portable Telecommunications (APT), Weber's team was able to orchestrate strategic demand reduction (e.g.,

"I'll back off license X if you back off Y") between APT and WirelessCo through its bidding behavior. Weber also discussed bid shielding, bid sniping, and fraud on the popular auction website eBay.com. A guide to Weber's talk is available at http://www.kellogg.nwu.edu/faculty/weber/ibm-acm/index.htm.

EC'00 also featured a panel discussion on the results of the Trading Agent Competition (TAC) held at the Fourth International Conference on Multiagent Systems in Boston in July 2000. TAC was organized by Dr. Michael P. Wellman, Professor at the University of Michigan, and Dr. Peter R. Wurman, Professor at North Carolina State University; both were present to field questions. Panelists included Wurman, along with contest finalists Peter Stone at AT&T Research, Justin Boyan at MIT, Amy Greenwald at Brown University, Andrew Goldberg at InterTrust Technologies, and Youyong Zou at the University of Maryland, Baltimore County. The goal of the contest was to design and implement trading agents to automatically configure travel packages on behalf of multiple customers, while competing with other agents for limited resources. The problem was complex enough to encourage sophisticated solutions, some including adaptation, learning, and combinatorial optimization; indeed, all finalists employed quite elaborate strategies. The number of competitors was large enough to make fully rational game-theoretic reasoning intractable, yet not too large as to justify efficient-market assumptions. The meta-goal of the contest was to generate ideas and stimulate progress in trading-agent research. The meta-assessment of the panelists was that this meta-goal was met in an enjoyable way, and all looked forward to the second competition, TAC-01, scheduled to coincide with EC'01. More information about TAC is available at http://tac.eecs.umich.edu/.

3. TECHNICAL SESSIONS

The EC'00 technical program served as the foundation of the conference. Over the course of two and a half days, twenty-nine presentations spanned a range of research topics in ecommerce. Corresponding written reports appear in the archival proceedings [MacKie Mason and Tygar 2000]. Papers can be roughly divided into six categories (with considerable overlap): auctions, empirical studies, simulation studies, analytical models, systems and algorithms, and interfaces.

The program began and ended with talks on auctions, reflecting an increasing awareness of the role of computation in auctions, and mirroring the recent commercial growth of electronic auction houses, whether consumer-to-consumer (e.g., eBay.com), business-to-consumer (e.g., priceline.com), or business-to-business (e.g., ariba.com). In particular, the cacophony of conclusions concerning combinatorial auctions (CAs) clearly cemented CA as the three-satisfiability problem of ecommerce. Following two full sessions on CAs at the Seventeenth National Conference on Artificial Intelligence in Austin in July 2000,[4] seven talks at EC'00 dealt with CAs. CAs allow bids on arbitrary combinations of goods, so solving for the optimal allocation is NP-hard. Two speakers presented improved winner-determination algorithms, one optimal [Gonen and Lehmann, p. 13] and one approximate [Sakurai, Yokoo, and Kamei, p. 30]. Other aspects of CA research discussed included the tradeoff between bid language expressivity and computational complexity [Nisan,

---

[4]http://www.aaai.org/

p. 1]; progressive and non-discriminatory pricing mechanisms [Wurman and Wellman, p. 21]; the generation of realistic bids for testing and comparing algorithms [Leyton-Brown, Pearson, and Shoham, p. 66]; the design of incentive compatible and computationally tractable CAs [Nisan and Ronen, p. 242]; and the use of CAs for market-driven supply chain formation [Walsh, Wellman, and Ygge, p. 260]. Other (non-CA) auction-related talks addressed the incentive compatibility of online auctions [Lavi and Nisan, p. 233] and methods for setting up stable coalitions among bidders [Leyton-Brown, Shoham, and Tennenholtz, p. 253].

Empirical studies investigated the use (and potential misuse) of web server logs to deduce customer activity [Rosenstein, p. 38]; the divergence of online book prices from theoretical predictions [Clay, Krishnan, and Wolff, p. 44]; regularities and self-similarities in ebusiness workloads [Menascé *et al.*, p. 56]; and comparisons among product recommendation algorithms [Sarwar *et al.*, p. 158]. Deck and Wilson [p. 77] conducted laboratory experiments to test subjects' preferences among several automated and manual price-setting strategies.

Several authors employed simulations to explore new types of markets and new agent strategies. Two papers studied markets of information goods, where marginal costs and bundling costs are almost negligible. Brooks, Durfee, and Das [p. 95] showed how producers can discover and exploit product niches, thus avoiding destructive competition and price wars. Kephart and Fay [p. 117] proposed a model of information-good bundling, and examined its behavior through game-theoretic analysis and simulation. Morris, Ree, and Maes [p. 128] used a market simulator to compare adaptive pricing strategies in a hypothetical airline-seat auction. Dellarocas, Klein, and Rodriguez-Aguilar [p. 225] built a Swarm-based simulation of a distributed network to bolster their argument that institutional-level exception handling mechanisms are preferable to individual-level mechanisms.

Two papers developed analytic models that account for new economic phenomena arising from the sale and use of information technology. One model showed how producers can benefit by offering free information products that are either strategic complements (e.g., Adobe Acrobat Reader as a complement to Acrobat) or strategic substitutes (e.g., Microsoft Internet Explorer as a substitute for Netscape Navigator) [Parker and Van Alstyne p. 107]. The other model described how information technology facilitates product customization and differentiation, thus explaining why larger-than-expected price variability still persists on the Internet [Farag and Van Alstyne, p. 135].

Five speakers presented new systems, algorithms, or mechanisms tailored for ecommerce environments. Three systems offered expanded options for online buyers and sellers: the Multi-Attribute Resource Intermediary, or MARI, supports agent-mediated negotiations on a variety of product attributes beyond simply price, including characteristics of the transaction partner [Tewari and Maes, p. 86]; MultE-commerce allows multiple users to shop together on the web and even share payments [Puglia, Carter, and Jain, p. 215]; and Micro-Options give consumers the right—but the the obligation—to purchase resources, thereby reducing the risks involved in trying to acquire multiple complementary resources [Ginis and Chandy, p. 207]. Chickering and Heckerman [p. 145] developed and tested an algorithm for targeting online ads that maximizes click-through rates while satisfying contractual obligations. Dellarocas [p. 150] proposed and evaluated several mechanisms

for combating fraud in reputation reporting systems (e.g., eBay.com's user ratings system).

Four papers addressed interface and usability issues in ecommerce. Lee and Podlaseck [p. 168] illustrated how starfield visualization techniques can be used to display online retailer data and guide business decisions. Arlein *et al.* [p. 176] described a protocol allowing website customization to depend on user activity across multiple sites, while preserving user and merchant privacy constraints. VanderMeer, Dutta, and Datta [p. 185] developed a scalable system for website customization that tracks changing user preferences and operates much more efficiently than off-the-shelf database software. Callahan and Koenemann [p. 197] conducted usability experiments that compared a standard hierarchical electronic product catalog interface to the InfoZoom interface, which features flexible navigation (including the ability to "zoom" in), and compression techniques to display many items on a single page.

4. CONCLUSION

SIGecom intends for the EC conference to serve as the core academic venue for research in computational aspects of commerce. Given the growth in SIGecom's membership, the rapid increase in submissions to EC, and, more importantly, the strong technical content of EC'99 and EC'00, SIGecom seems well on its way to achieving this goal.

REFERENCES

MacKie Mason, J. K. and Tygar, D. Eds. 2000. *Proceedings of the Second ACM Conference on Electronic Commerce (EC-00)*. ACM Press, New York.