Preface

The design of economic agents, mechanisms, and systems has received growing attention in the agents and multi-agent systems communities. Electronic commerce is rich with focused yet challenging problems, ripe for technical advances and practical application of agent technologies. As the domain is characterized by individual agent self-interest and private information, agent-mediated trade requires principled design, often incorporating novel combinations of theories from different disciplines. Thus, techniques from fields such as computer science, operations research, artificial intelligence, and distributed systems are integrated with principles from economics and game theory. Furthermore, there are challenges to eliciting human preferences and requirements and ensuring that they are represented in automated agent behavior.

The goal of this workshop was to explore research in the principled design of economic agents, mechanisms, and systems. To promote a forward-looking discussion, we encouraged the submission of extended abstracts, work-in-progress reports, and position papers, as well as more developed technical papers, expressing innovative and principled ideas and approaches.

Areas of particular interest included:

- mechanisms, negotiation protocols, and auctions (especially advanced designs such as multi-attribute auctions),
- bidding and negotiation strategies,
- integration of negotiation with broader decision making,
- economic-based theory and design methodologies,
- market-based problem solving,
- trading and pricing,
- eliciting human preferences and requirements and ensuring that they are represented in automated agent behavior,
- significant new problem domains.

The selected results of the call for papers are now presented here, in revised form, as a formal record of the meeting. The papers fell broadly into three categories: auctions, negotiation, and markets. The meeting took the conventional form of presentation, followed by questions, although there was a poster session mid-morning and mid-afternoon, which generated plenty of interest, and the day concluded with a panel session on "Moving E-Commerce Research to the Real World: Key Technologies and Practical Challenges." This was chaired by William Walsh (IBM, Hawthorne) and the participants were: Tuomas Sandholm (Carnegie Mellon), Norman Sadeh (Carnegie Mellon), Yoav Shoham (Stanford), and Onn Shehory (IBM, Haifa).

The Papers

Auctions formed the largest group of papers at this fourth running of the AMEC workshop, but several themes emerged within the topic. For example, aspects of combinatorial auctions were addressed in four papers.

In *Bidtree Ordering in IDA* Combinatorial Auction Winner-Determination with Side Constraints*, authors John Collins, Güleser Demir, and Maria Gini describe the incorporation of side constraints into Sandholm's bidtree algorithm– although these mostly have the consequence of preventing many of the optimizations described for CABOB by Sandholm. They present experimental evidence of the benefits of ordering the items in the bid tree according to the number of bids containing the item.

In *Differential-Revelation VCG Mechanisms for Combinatorial Auctions*, authors Wolfram Conen and Tuomas Sandholm put forward two auction mechanisms for the combinatorial setting in which bidders reveal differential rather than absolute valuations on a subset of the combinations offered. The mechanisms result in an efficient allocation of items and also determine the actual (Vickrey) payments.

In *Effectiveness of Preference Elicitation in Combinatorial Auctions*, authors Benoît Hudson and Tuomas Sandholm present empirical and theoretical evidence of the benefit of preference elicitation (as assumed in the previous paper), where bidders only put forward bids on a subset of the offered bundles. They also show that the information elicited is a vanishingly small percentage of that revealed in conventional mechanisms, and furthermore how this information can be acquired incrementally and refined as needed.

In *Minimal-Revelation Combinatorial Auctions*, author David Parkes also considers the undesirability of bidders having to reveal complete preference information, but from the perspective of determining how little information must be made available to achieve the same outcome. He establishes the existence of a class of certificates which are necessary and sufficient for a dominant strategy mechanism.

The interest in combinatorial auctions was complemented by three papers exploring the problems of multiple auctions, auctions for divisible items, and auctions for multi-attribute items.

In *A Comparison Among Bidding Algorithms for Multiple Auctions*, author Andrew Byde examines the results of simulations of different bidding strategies for English auctions. He analyzes the conditions leading to success and demonstrates that a dynamic programming approach to constructing models of opponents is as robust as and outperforms a conventional greedy approach.

In *Auctions for Divisible Resources: Price Functions, Nash Equilibrium and Decentralized Update Schemes*, authors Rajiv T. Maheswaran and Tamer Başar are driven by the application of agent-mediated auctions to the distribution of computational and network resources. They propose the use of a "divisible" auction as being more appropriate to these kinds of goods and, because optimal responses can be represented as price functions, this leads to an existence proof of a Nash equilibrium and the development of a distributed algorithm.

In *An English Auction Protocol for Multi-attribute Items*, authors Esther David, Rina Azoulay-Schwartz, and Sarit Kraus consider a common class of auction application (task allocation in its various forms), where the buyer takes the role of the auctioneer requesting bids and the sellers have the role of bidders offering particular configurations matching the request. They propose the use of the English auction protocol, such that each bid proposes an improvement on earlier bids and a better match with the original request, concluding with optimal bidding strategies and optimal auction designs.

An intriguingly different perspective on auction mechanism design was presented by Steve Phelps, Phil McBurney, and Mike Wooldridge in *Co-evolutionary Auction Mechanism Design: A Preliminary Report*. Instead of traditional analytic methods, mechanism design is seen instead as an evolutionary process which automatically generates strategies for traders and auctioneers. The idea is presented in the context of double auctions for a wholesale electricity marketplace.

In *ATTac-2001: A Learning, Autonomous Bidding Agent* authors Peter Stone, Robert Schapire, János Csirik, Michael Littman, and David McAllester discuss their experience in building the ATTac-2001 agent to participate in the second Trading Agent Competition. Their focus is on the development of a model for the learning of price dynamics based on past data in order to permit the calculation of optimal bids. They include experimental evidence on the effectiveness of the approach against some alternatives.

Negotiation formed the second major theme at AMEC-IV, comprising one theory paper, one practical paper, and several that built bridges between the two.

In *The Influence of Information on Negotiation Equilibrium*, authors Shaheen Fatima, Michael Wooldridge, and Nicholas Jennings undertake a theoretical analysis of how the information that an agent has about its bargaining partners affects the equilibrium behaviour. Such information might be deadlines, discounting factors, or reservation prices, for example. By systematic variation of the amount of information available, the authors show the relative impacts of the opponent's parameters on the negotiation outcome, and hence on which of its opponent's parameters an agent should learn in order to maximize its utility. As a consequence, an agent could select its conegotiator on the basis of its information state.

In *A Software Infrastructure for Negotiation Within Inter-organisational Alliances*, authors Mihnea Bratu, Jean-Marc Andreoli, Olivier Boissier, and Stefania Castellani aim to construct a middleware framework for collaboration between organizations, while preserving participants' autonomy. The negotiation process is turned into a distributed constraint satisfaction problem operating at three levels (middleware, multi-agent, and human) and the approach is illustrated using a scenario in which printshops trade printing jobs.

In contrast to the few orthogonal issues with which computational models of negotiation are mostly concerned, real-world contracts are usually more complex. In *Using an Annealing Mediator to Solve the Prisoner's Dilemma in the Negotiation of Complex Contracts*, authors Mark Klein, Peyman Faratin, and Yaneer

Bar-Yam propose a mediator-based negotiation protocol for this situation of many overlapping issues. The mediator puts forward a contract that is either accepted by the parties to the negotiation, in which case the process terminates, or rejected, leading to a revised (better) proposal from the mediator. However, individual agent strategies can lead to adverse outcomes, and so the contribution of this paper is to modify the mediator to demonstrate a time-decreasing willingness to pursue contracts that were rejected by both parties.

In *Automated Negotiation in Many-to-Many Markets for Imperfectly Substitutable Goods*, authors Chris Preist and Carlos Mérida-Campos describe an agent designed to carry out the negotiation of buying and selling in a double auction style market. In the market, the goods are almost the same (substitutable), but one may be preferred over another. The given for-instance of such a good is memory chips which have different fault rates. The agent itself is a generalization of earlier work on ZIP agents. The main results are the extension to imperfectly substitutable goods, experimental results showing that the new agents converge to the predicted general equilibrium.

Negotiation protocols can vary widely in practice, even if they are similar in principle, making it difficult for an agent to be effective in many environments. Valentina Tamma, Michael Wooldridge, Ian Blacoe, and Ian Dickinson outline in *An Ontology Based Approach to Automated Negotiation* an alternative approach in which agent interactions are in terms of a shared ontology of negotiation terms. They demonstrate the idea by application to the trading agent competition scenario.

Markets defined the third theme and brought together a diverse collection of material from frameworks for constructing agent economies to the focussed application of agents in advertising space trading, from modeling individual reputation to guaranteeing properties of systems, and from steps towards the generation of trading platforms from specifications to the analysis of the effects of agent trading behavior on markets.

The possibility, or even probability, of widespread application of agents and the emergence of agent-based economies inspired Steven Willmott, Monique Calisti, and Emma Rollon to propose, in *Challenges in Large-Scale Open Agent Mediated Economies*, such scenarios as potentially important tools for agent-based research now, which will in turn aid future adoption and integration of agents into real markets.

In *An Extensible Agent Architecture for a Competitive Market-Based Allocation of Consumer Attention Space*, authors Pieter Jan 't Hoen, Sander Bohte, Enrico Gerding, and Han La Poutré describe a distributed recommendation mechanism based on adaptive software agents, whose task is to allocate advertising banners in an electronic shopping mall. The task of evaluating and classifying customers' bids on banners is distributed amongst individual shop agents, rather than performed centrally. Hence each agent may apply its own private strategy, learning mechanism, and specific domain knowledge without revealing potentially commercially sensitive information. The extension over earlier work is the

distributed aspect, leading to scalability and extensibility, combined with the means to operate concurrently in multiple markets.

Modeling reputation via statistical analysis of past ratings works well for some contexts, but in *Goodwill Hunting: An Economically Efficient Online Feedback Mechanism for Environments with Variable Product Quality* author Chris Dellarocas considers how to deal with environments in which goods intrinsically vary in quality. He puts forward a new feedback mechanism that uses the threat of biased reporting in the future to induce sellers to be accurate in declaring the quality of the goods offered.

The converse of individual reputation is institution or system reputation, which is an aspect of *Guaranteeing Properties for E-Commerce Systems* by Frank Guerin and Jeremy Pitt. The focus of the paper is on how solutions from game theory and from computing can be used to specify rules and prove properties of agent systems. The objective is to enable the creation of a wider range of systems to which agent owners may be more willing to let their delegated agents go and enter into legally binding contracts. Although, as the authors point out, neither the use of game theory nor the verification techniques are new in the context of MAS, what the paper does provide is a formal framework which allows both to be used to engineer agent systems.

The well-known FishMarket auction house is the subject of *Skeletal Jade Components for the Construction of Institutions*, in which authors Oliver Vickers and Julian Padget examine the process of developing another implementation of that auction house, but this time working from the ISLANDER institution specification language towards stylized Jade components with the long-term objective of "round-trip" engineering of trading platforms. A novel aspect of this implementation included the generation of the ontology from a specification in RDF, which was written using Protege by means of the Jade Beans generator. Also novel is the implementation of the distributed auction protocol, described in an earlier work, that does without the need for an auctioneer.

In *Self-enforcing Strategic Demand Reduction* authors Paul S.A. Reitsma, Peter Stone, János A. Csirik, and Michael L. Littman consider the effect of agent bidding strategies on both the overall market and the other agents participating in the market, taking as a simulation scenario the FCC spectrum auctions. Having examined several well-known strategies in this context, they also put forward a new strategy *punishing randomized strategic demand reduction* and demonstrate that if all bidders use this then it is mutually beneficial to all the participants. Furthermore, the mechanism automatically detects and sanctions non-cooperating bidders, making it robust against defection.

The Panel

The overarching issue presented to the panel was to consider some of the key challenges involved in translating agent-mediated electronic commerce research into products, both within an R&D organization and in the marketplace. As a starting point, the chair asked the panel to provide examples of technology

transfer of agent-based e-commerce research and ideas, identify past successes and failures, and offer lessons for the future.

Sandholm reported on the success of the technology deployed by CombineNet Inc. for combinatorial market winner determination. In the first quarter of 2002, some 30 combinatorial auctions had been run by CombineNet. The value of goods cleared totaled over USD 1 billion while showing a saving of USD 100 million over the previous year's pricing. In the medium term, he foresees the increased utilization of preference elicitation and indirect mechanisms. For the longer term, he identifies the importance of *automated* mechanism design, in which the objective is to automatically design the rules of the game so that a good outcome happens in spite of the agents' strategies. Sandholm identified some impediments to technology transfer: customers do not always know what they need and are often committed to a particular ad hoc mechanism. He also noted that it is computationally intractable to calculate equilibria for undesigned mechanisms. In response, David Parkes commented that mechanisms should be designed with actual agent capabilities and behaviors in mind, rather than relying on the impractical theoretical models of rationality and equilibrium generally assumed now.

Shehory began by reiterating what he believed to be widely accepted: that eCommerce adoption has been slower than initially expected, that agents are not yet a widespread technology, and that, while industry sees eCommerce agents as interesting, it does not yet see the economic value of the present technology. However, he also asserted that growth in both electronic transaction volume and the adoption of eCommerce solutions continues. Additionally, there is growing acceptance that eCommerce reduces transaction costs. These trends provide openings that could lead to wider adoption of agents and applications to a greater range of business functions. Technical challenges for greater business use of eCommerce technologies include enabling technologies such as security and privacy, and demonstrations of economic benefits. Scientific challenges include better human–computer interfaces, development tools, environments, and infrastructures.

Sadeh, like Shehory, took a similar broad view of the issues involved, expressing concern over the willingness of users to delegate to software and arguing that the benefits of using agents have to outweigh the efforts the user has to invest to program the agent, for example in specifying preferences/valuations, constraints, etc. He pointed out that this is more likely to be the case when looking at tasks that are time-consuming, computationally challenging and/or repetitive in nature, for example you can enter your preferences/valuations once and reuse them across a number of instances. He felt that today examples of such tasks are easier to find in the context of B2B than B2C scenarios. For B2C, and in particular mobile commerce, Sadeh suggested we need to develop techniques that are capable of *automatically* capturing part of the user's context, and even to plan dialogues with the user to elicit missing information. Unfortunately, contextual information tends to be distributed across a number of systems, such as user calendar information, location tracking functionality, etc., making

it difficult to build a comprehensive and consistent picture. In respect of this, he described the *Semantic eWallet* being developed by the mobile commerce group at CMU, which is a Semantic Web directory of personal resources such as those listed above, and through which agents can automatically discover and access relevant contextual information. In addressing the question of eCommerce failures, Sadeh asserted that, by and large, emarketplaces have failed because of their narrow focus on cost and suggested that it was important to compete on more than one factor, citing work with Mike Wellman on multi-attribute negotiation as one way of achieving this.

Shoham argued that researchers should focus on doing good research, while product developers should focus on the best product. Sometimes research and commercial objectives will coincide, but the best design for a product is not generally a significant research contribution. Conversely, good research does not always lead directly to a good product. Nevertheless, good research can still have an impact, even if there is no immediate commercial product, because the ideas get added to the corpora of knowledge on eCommerce. The topic of mechanisms and their design was regarded as very significant by all the panel, but in contrast to Sandholm's optimism for dominant-strategy mechanism design, Shoham did not believe that it would be tractable for all situations of importance. He identified agent strategic reasoning and learning as key areas for future research, but noted that our current theoretic models are inadequate or flawed.

To summarize the discussion of the panel: it is an uphill struggle to get industry acceptance of eCommerce; good research and good product design are generally separate endeavors; there were differing views on the feasibility of dominant-strategy mechanism design, but consensus that there are many computational and methodological challenges in developing robust agent strategic reasoning.

September, 2002

Julian Padget, David Parkes, Norman Sadeh,
Onn Shehory, William Walsh

Program Committee

Boi Faltings	Ecole Polytechnique Fédérale de Lausanne, Switzerland
Peyman Faratin	MIT Sloan School of Management, USA
Amy Greenwald	Brown University, USA
Martin Griss	Hewlett-Packard, Palo Alto Labs, USA
Robert Guttman	Frictionless Commerce Inc., USA
Nick R. Jennings	University of Southampton, UK
Sarit Kraus	Bar-Ilan University, Israel
Yannis Labrou	University of Maryland, USA
Joerg Muller	Siemens GmbH, Germany
David Pennock	NEC Research Institute, USA
Jeremy Pitt	Imperial College, London UK
Chris Preist	Hewlett-Packard, Bristol Labs, UK
Jeff Rosenschein	Hebrew University, Israel
Tuomas Sandholm	Carnegie Mellon University, USA
Katia Sycara	Carnegie Mellon University, USA
Michael Wellman	University of Michigan, USA
Mike Wooldridge	University of Liverpool, UK
Peter Wurman	North Carolina State University, USA

Additional Reviewers

Rina Azoulay-Schwartz	Bar-Ilan University, Israel
Andrew Byde	Hewlett-Packard, Bristol Labs, UK
Gangshu Cai	North Carolina State University, USA
Vincent Conitzer	Carnegie Mellon University, USA
Minghua He	University of Southampton, UK
Benoît Hudson	Carnegie Mellon University, USA
Kate Larson	Carnegie Mellon University, USA
Kevin Leyton-Brown	Stanford University, USA
Xudong Luo	University of Southampton, UK
Panos Markopoulos	University of Pennsylvania, USA
Pritti Phadke	North Carolina State University, USA

Workshop Organizers

Julian Padget Department. of Computer Science,
University of Bath, UK
jap@cs.bath.ac.uk

David C. Parkes Department of Computer Science,
Harvard University, USA
parkes@eecs.harvard.edu

Norman M. Sadeh eCommerce Institute, ISRI,
Carnegie Mellon University, USA
sadeh@cs.cmu.edu

Onn Shehory IBM – Haifa Research Lab, Israel
onn@il.ibm.com

William E. Walsh IBM T.J. Watson Research Center, USA
wwalsh1@us.ibm.com

Acknowledgements

The workshop organizers would like to thank the AAMAS 2002 conference and the University of Bologna for hosting the workshop. We are also very grateful to all the program committee, the reviewers, the authors and the participants in the panel session for their respective contributions to the process of creating a high quality, stimulating workshop. Finally, we are pleased to acknowledge the encouragement and support from Springer-Verlag, in the person of Alfred Hofmann, for helping to bring the workshop to this concrete conclusion.