Decision-theoretic bidding in online-auctions

Eberhard Schmid and Torsten O. Paulussen

Working Paper 4/2004
April 2004

Working Papers in Information Systems

University of Mannheim
Department of Information Systems 1
D-68131 Mannheim/Germany
Phone +49 621 1811691, Fax +49 621 1811692
E-Mail: wifo1@uni-mannheim.de
Internet: http://www.bwl.uni-mannheim.de/wifo1

Decision-theoretic bidding in online-auctions

Eberhard Schmid
Dept. of Information Systems 1
University of Mannheim
68131 Mannheim
Germany
eschmid@rumms.uni-mannheim.de

Torsten O. Paulussen
Dept. of Information Systems 1
University of Mannheim
68131 Mannheim
Germany
paulussen@bwl.uni-mannheim.de

Abstract

With the increasing role of electronic commerce in business applications, much attention is paid to online-auctions. As auctions become more and more popular in electronic commerce, agents face the problem of participating in multiple independent auctions simultaneously or in sequence. Decision making of agents becomes difficult when they have to buy bundles of goods. In this case the agents have to cope with substitutable or complementary effects between the single goods. In this paper we analyse existing approaches of tackling the problem of decision making in multiple, heterogeneous auctions and develop a flexible Dynamic Programming-based decision-making framework for agents, participating in multiple auctions. This work extends existing Dynamic Programming-approaches in this field.

1 Introduction

With the increasing role of electronic commerce in business applications much attention is paid to online-auctions. As auctions become more and more popular in electronic commerce, agents face the problem of participating in multiple independent auctions simultaneously or in sequence [7, 1]. From a decision-theoretic point of view one faces a dynamic decision-problem under uncertainty. The task of decision making becomes complex, if the agents are supposed to buy bundles of goods with components, which are sold in multiple independent auctions. In this case the agent cannot assign a value to a single good but only to the whole bundle. Thereby, complementary and substitutable effects between goods may arise. In such cases, strategies developed in auction-theory cannot be applied, because of their assumptions concerning the internal valuation of goods, the amount and types of players [15, 16]. Therefore bid distributions are usually learned from historical data and represent an external, given factor. In this paper we focus on the decision making part. For the learning aspect see for example [10, 4].

In this paper, first, we provide a brief review of approaches dealing with bidding strategies in multiple auctions. Our main focus will be on agents which want to buy bundles of goods. Second, we develop a conceptual decision-making framework for agents that participate in multiple auctions and whereas the agents value whole bundles of goods. Our framework is modeled with Markov Decision Processes (MDPs) and extends the work of [4] and [5].

The remainder of this paper is organised as follows. In section 2 we give a brief overview of related work in the field of bidding strategies for agents. In Section 3 we develop our MDP-Model. The conceptual framework incorporates several different auction protocols as well as sequential and overlapping auctions. In Section 4 we summarise our findings an give an outlook on future research in the field of MDP-based bidding models in trading-agent scenarios.

2 Related Work

In decision theory, one can distinguish between decision problems under certainty and risk. Furthermore, problems can be divided into static and dynamic problems. In static problems, time does not play any role. Dynamic problems explicitly deal with time, i.e. actions taken in the present may have consequences for decisions in the future [2].

Most of the existing approaches have in common, that they do not explicitly incorporate time in the decision models. Rather, the dynamic decision problem is modeled as a static problem, which means a simplification. We first focus on approaches which deal with the purchase of a single good as well as with the purchase of multiple goods, where the reservation-price of each good is independent of the amount of goods the agent holds. A static decision model for multiple overlapping English auctions is provided in [20, 19, 6, 5]. In this case, the agent determines whether to bid in an English auction or not, by taking into account the expected utility of bidding for the desired good in an auction which closes later. An approach without explicit decision-theoretic background is presented in [1]. In this approach the agent determines maximum bid values at a given time as a polynomial function, as proposed in [11]. The maximum bid value depends on characteristics of the current state. The agent combines several tactics, which express for example a bid value as a function of the remaining time, or number of remaining open auctions. The weights of the single tactics, as well as the parameters of the polynomial functions are optimised with genetic algorithms for different environments. Dumas et al. [10] present an approach where no explicit utility functions are given. In this work, the user has to provide the agent a desperateness-factor, which specifies the probability with which the agent has to acquire the good. The agent has to compute a bid, which it is willing to place in several, sequential auctions. The probability of winning at least one auction with that bid has to be equal to the desperateness-factor announced by the user. An approach which is beyond classical decision- and probability theory is developed by Garcia et al. [12]. The core of this approach is the possibility theory [9]. This theory orders the utility of consequences on an ordinal rather than on a cardinal scale (like in classical utility theory). Furthermore, no distribution functions are used. Consequences are ranked on an ordinal scale, which denotes the degree of plausibility of occurrence.

Having briefly described approaches which assign values to individual goods, we now turn attention to the more complex case, where goods are valued in bundles. When an agent values goods in bundles, no reservation price can be assigned to individual goods. In fact, the agent has to distribute the value of the bundle over the individual goods [4]. Byde et al. [7] provide a generic framework for agents that operate in multiple heterogeneous auctions. Different auction protocols are modeled within this framework. The agent can purchase multiple items of one good, where the agents' utility is expressed as a function of the amount of goods the agent owns. Heuristics, based on thresholds for maximum bids in each auction, are presented. The purchase of bundles of different goods is addressed in [21, 18]. In this work, only English-auctions are addressed. Beside these general generic decision frameworks, domain-specific approaches emerged from the Trading Agent Competition (TAC) [13]. TAC describes a scenario where agents, which represent travel agencies, have to purchase trips for eight clients. The single components of the trips - i.e. flights, hotel rooms, and entertainment tickets - are traded in separate auctions, causing complementary and substitution effects between the goods. Most approaches in the TAC-field model the stochastic dynamic problem as a static problem under certainty, which means a rough simplification of the problem. Probability functions, as for the expected closing price, are reduced to deterministic expected values, which enter in an optimisation model. It is therefore a challenging task to accurately predict closing prices of auctions [24]. Different approaches to predict closing prices of hotel-room auctions have been undertaken. Besides simple historical averaging [24] there exist more sophisticated approaches based on logistic regression [23], fuzzy-logic [14] and competitive equilibrium analysis [8].

3 A Dynamic Decision-Framework for Multiple Auctions

In this section, we develop a dynamic decision-framework that is based on Markov Decision Processes (MDPs) and can be solved via Dynamic Programming (DP) [22]. So far, only a few MDP-approaches for the determination of bidding strategies in online auctions exist. Byde [5] develops an Dynamic Programming approach for agents that bid in multiple overlapping English auctions, in which the agent wants to buy a single item of a good. Boutilier et al. [4] present an approach where an agent wants to buy a bundle of goods, which components are sold in strictly sequential first-price-sealed-bid auctions. Based on these approaches, we develop a generalised decision-framework, which incorporates different auction protocols. Furthermore different start and closing times are possible. We will describe our model in the remainder of this section.

3.1 Assumptions

We assume that the agent may bid in a heterogeneous auction scenario. Therefore, we want to incorporate English, Dutch, first-price-sealed-bid and Vickrey auctions. In the case of English and Dutch auctions we consider auctions, in which the ask-price increases respectively decreases by a fixed increment at every time step. Further the agents have to decide whether they want to continue to bid. Only one agent can hold the active bid in English auctions. Further, we assume, that the agent has beliefs about the distributions of closing prices of the auctions. Finally, for computational reasons, we assume that the agent has quasi-linear preferences concerning the amount of money it holds, and that the agent has no budget constraint. Based on these assumptions we describe the decision model in the next section.

3.2 The Model

In this section we develop a MDP-based model for the heterogeneous auction scenario. To this end, we model decision-stages, the state- and action-space, the transition probabilities and the rewards and costs. Finally, we present the derived functional equations.

Stages, states and actions

We introduce the decision-stages t, which are synchronisation time points. At these discrete time points ask-prices in English and Dutch auctions change, and sealed-bid auctions may close. Let E, D, F and V denote the sets of English, Dutch, sealed-bid-first-price and Vickrey auctions. Let Z denote the space of states. Every state $z \in Z$ consists of the actual holdings h in every auction and the status s, which denotes the status of every auction. The set of possible statuses S contains the statuses open (op), closed (cl), active (ac), and inactive (in). Let s^z denote the vector of statuses of auctions, and h^z the vector of holdings in every auction in state z. $s^z(i)$ denotes the status of auction i in state z, whereas $s^z(i) \in \{ac, in, cl\} \forall i \in E$ and $s^z(i) \in \{op, cl\} \forall i \in D, F, V$. Therefore

$$z = (h, s)^t.$$

The action space A is given by the possible actions, which are represented by bids the agent can submit at every decision stage. In English and Dutch auctions the agent only can decide to bid or not to bid. Therefore $Bid(i) \in \{bid, nobid\} \forall i \in E, D$, where $Bid(i)$ denotes the bid in auction i. In sealed-bid auctions the agent can bid every positive number, therefore $Bid(i) \in R^{+} \forall i \in F, V$. The agent can only bid in English auctions when its status is inactive.

Transition probabilities

We assume, that the agent has beliefs about the distribution $P_i^t(x)$ of the highest bid of the opposing players at any time t for every auction i, where x denotes the bid. In English and Dutch auctions the actual bid price is a function of time, thus $x = x(t)$. For the case of English and Dutch auctions, we assume that the agent has beliefs about probabilities that it does not get the active bid, respectively does not win (see also [5]). We will denote these as the blocking-probability $B_i^t(x)$ which denotes the chance that the agent at time t does not get the active bid, or - in dutch auctions - does not win the good, if he bids x. For reasons of flexibility, we divide the transition-probabilities into *closing-probabilities*, which denote the probability that an auction closes at the next time step, and the *winning-probability* which denotes the probability of winning any specific goods, given the referring auction has closed or not. Transition probabilities depend on the type of the auction. We will derive them in the following paragraphs.

English auctions The *closing probability* in English auctions is the probability that no agent wants to place a bid higher than the actual bid. If the agent bids, the closing probability is zero. Thus

$$Prs_i^{t+1}(Bid, z) = \begin{cases} \frac{P_i^{t+1}(x(t+1)) - P_i^{t+1}(x(t))}{1 - P_i^{t+1}(x(t))} & \text{if } Bid(i) = nobid \\ 0 & \text{else.} \end{cases}$$

If the auction closes at $t+1$ one can build the *winning probabilities* for winning the auction:

$$Prg_i^{t+1}(1, cl|Bid, z) = \begin{cases} 1 & \text{if } s^z(i) = ac \\ 0 & \text{else,} \end{cases}$$

Respectively, the probability of loosing the auction is given by:

$$Prg_i^{t+1}(0, cl|Bid, z) = \begin{cases} 1 & \text{if } s^z(i) = in \text{ und } Bid(i) = nobid \\ 0 & \text{else.} \end{cases}$$

If the referring auction does not close, the *winning probabilities* are related to resulting states.

$$\overline{Prg_i^{t+1}}(0, ac|Bid, z) = \begin{cases} 1 - B_i^{t+1}(x(t+1)) & \text{if } s^z(i) = in \text{ und } Bid(i) = bid \\ 0 & \text{else,} \end{cases}$$

$$\overline{Prg_i^{t+1}}(0, in|Bid, z) = \begin{cases} B_i^{t+1}(x(t+1)) & \text{if } s^z(i) = in \text{ und } Bid(i) = bid \\ 1 & \text{else.} \end{cases}$$

Dutch auctions For the case of Dutch auctions, transition probabilities may be build very similar. If the agent bids, the *closing probability* must be one. Else, the probability is given by the probability, that the highest opposing bid is higher than the bid price at the next stage. Thus

$$Prs_i^{t+1}(Bid, z) = \begin{cases} \frac{P_i^{t+1}(x(t)) - P_i^{t+1}(x(t+1))}{P_i^{t+1}(x(t))} & \text{if } Bid(i) = nobid \\ 1 & \text{else.} \end{cases}$$

If the auction closes the agent only may win if he did bid in the previous stage. The *winning probabilities are*:

$$Prg_i^{t+1}(1, cl|Bid, z) = \begin{cases} 1 - B_i^{t+1}(t+1) & \text{if } Bid(i) = bid \\ 0 & \text{else,} \end{cases}$$

$$Prg_i^{t+1}(0, cl|Bid, z) = \begin{cases} B_i^{t+1}(t+1) & \text{if } Bid(i) = bid \\ 1 & \text{if } Bid(i) = nobid. \end{cases}$$

If auction i does not close at the next stage the only possible resulting status of the auction is open:

$$\overline{Prg_i^{t+1}}(0, op|Bid, z) = 1.$$

Sealed bid auctions In sealed bid auctions it is usually known when the auction closes. At any state the *closing probability* is therefore zero or one. In our approach, we want to allow that one can specify closing probabilities for sealed bid auctions. Thus the closing probability for sealed bid auctions is always given. The *winning probability* of an auction, when it closes, is given by the probability, that the Bid is higher than the highest opposing bid. Thus

$$Prg_i^{t+1}(1, cl|Bid) = Pr[Bid(i) \geq x_i(t+1)] = P_i^{t+1}(Bid(i)).$$

Rewards and costs

Now we want to focus on rewards and costs that are generated through the bids. The agent is only rewarded at the end of the planning horizon T, as it is not able to assign a value to immediately won goods. Thus, for every possible state at the end of the planning horizon the agent is rewarded the value of the bundle that he holds in this state [see also 4]:

$$e = u(h(z)),$$

where $u(h(z))$ denotes the value assigned to holdings in state z. The agent is assigned the price for every good the agent wins. One has to distinguish between different auction types. For the case of English and Dutch auctions the cost in case the agent is winning the auction is given by

$$c_i = x_i(t),$$

where $x_i(t)$ is the ask-price of auction i at stage t. Remember that we assumed English and Dutch auctions with fixed increments at every decision stage, thus the ask-price is a deterministic function of time.

For the case of sealed-bid-first-price auctions the immediate cost of winning a good results as

$$c_i = Bid(i).$$

As the cost assigned with the purchase of a good in a Vickrey auction equals the highest bid of an opponent player, the cost is given by the expected value of the opponent highest bid. If the agent wins the auction the expected cost in Vickrey auction is:

$$c_i(Bid(i)) = \frac{\int_0^{Bid(i)} x \cdot p_i^{t+1}(x)dx}{P_i^{t+1}(Bid(i))},$$

where $p_i^{t+1}(x)$ denotes the density function of $P_i^{t+1}(x)$.

Functional equation

Having described state- and action-space, transition-probabilities and rewards/cost for different auction protocols we will now develop the central functional equation of our model. The value-functions for any state are given in the final stage T by $V^T(z) = u(H(z))$. For each state in $t < T$:

$$Q^t(z, Bid) =$$
$$\sum_{w \subset O_z} \left\{ Prs_w^{t+1}(Bid, z) \left[\sum_{k \in K(w)} Prg_w^{t+1}(k|Bid, z) \cdot \left(V^{t+1}(h^z \cup k^{Anz}, k^S) - E\left[C(Bid, k)\right] \right) \right] \right\}.$$

We will now explain the components of the equation in detail. O_z denotes the set of auctions that are still open in state z. $Prs_w^{t+1}(Bid, z)$ denotes the probability that all auctions in w close at the next decision stage and is given by:

$$Prs_w^{t+1}(Bid, z) = \prod_{i \in w} Prs_i^{t+1}(Bid, z) \cdot \prod_{i \notin w} (1 - Prs_i^{t+1}(Bid, z)).$$

$K(w)$ denotes the set of possible wins of goods or the possible changes of statuses of any auction if all auctions in w close. k^{Anz} denotes the vector of won goods, and k^S denotes the vector of status changes. $k = (k^{Anz}, k^S)$ denotes the change of state, and $k(v)$ denotes the state change in auction v. For example, if auction 2 closes and the agent wins the auction, then $k(2) = (1, C)$. $Prg_w^{t+1}(k|Bid, z)$ results as

$$Prg_w^{t+1}(k|Bid, z) = \prod_{v \in w} Prg_v^{t+1}(k(v)|Bid, z) \cdot \prod_{v \in O_z \setminus w} \overline{Prg_v^{t+1}}(k(v)|Bid, z).$$

Finally the expected cost are given by the sum of the cost resulting from every winning auction. Therefore

$$E\left[C(Bid, k)\right] = \underbrace{\sum_{i \in w \wedge i \in V} \frac{k(i)^{Anz} \cdot \int_0^{Bid(i)} x \cdot p_i^{t+1}(x)dx}{P_i^{t+1}(Bid(i)}}_{\text{cost in Vickrey auctions}} + \underbrace{\sum_{i \in w \wedge i \in F} k(i)^{Anz} \cdot Bid(i)}_{\text{cost in First-Price auctions}}$$
$$+ \underbrace{\sum_{i \in w \wedge i \in E} k(i)^{Anz} \cdot x_i(t+1)}_{\text{cost in english auctions}} + \underbrace{\sum_{i \in w \wedge i \in D} k(i)^{Anz} \cdot x_i(t+1)}_{\text{cost in dutch auctions}},$$

where $k(i)^{Anz}$ is 1 if in k auction i closes and the agent wins the good, or 0 otherwise. The agent's aim is to maximise the expected utility. Thus

$$V^t(z) = \max_{Bid} Q^t(z, Bid).$$

4 Conclusions and Outlook

In this paper we discussed bidding strategies for agents in heterogeneous auction scenarios. After a brief discussion of existing work from a decision theoretic point of view, we developed an MDP-based decision framework that agents can use to bid in multiple heterogeneous auctions. This paper was inspired by [4] and [5]. This work extends the state by providing an unified MDP-based framework for different auction-types, whereas arbitrary start and closing times may be incorporated.

So far, in our model, quasi-linear preferences and the absence of a budget-constraint was assumed. Future work will address the incorporation of these points with respect to computational tractability. Because MDP-approaches result in high complexity for realistic problem-instances the appliance of complexity-reducing methods is of high importance [3]. One may distinguish between methods of abstraction, aggregation, and problem decomposition. In the case of abstraction and aggregation, parts of the state space are ignored respectively aggregated to a common state variable. In the case of problem decomposition, the problem is divided into several smaller problems. In order to apply decision models to real-world problems, research regarding complexity-reduction methods is mandatory.

We are currently implementing the developed framework for the TAC scenario, which we briefly described in section 2. In this scenario, the problem may be decomposed into eight subproblems.

More precisely, for every customer, we modeled a separate MDP to compute the optimal bids in hotel auctions. The challenge in this case is to coordinate the subproblems in a way, that an acceptable solution of the entire systems may be reached. One state of the entire MDP corresponds to multiple possible combinations of sub-MDP states, because goods may be exchanged between customers. We are currently implementing a coordination mechanism that is based on [17], which was developed for the allocation of resources to independent tasks. In first experiments, we benchmarked this approach against an approach which uses historical averages of hotel closing prices and then computes the target auctions via integer programming. The bid value is then computed as the marginal value of each item as in [8]. On the TAC server-platform, we compared the new approach (DPAgent) with seven instances of the static approach (IPAgent). The following figure shows the performance of these approaches (50-run average). Entertainment tickets are ignored in these experiments.

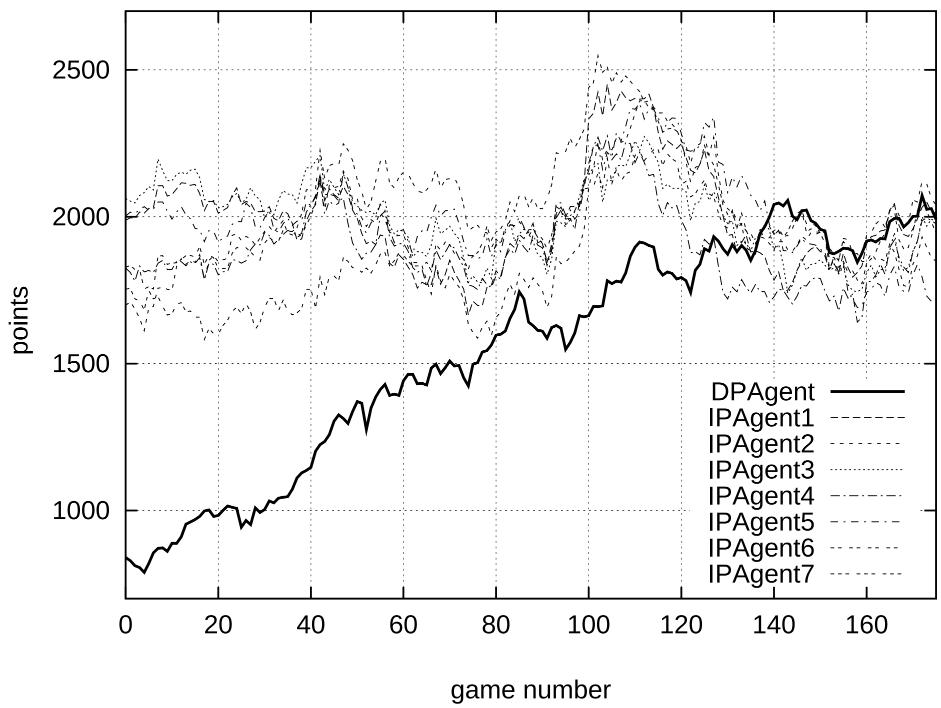

Figure 1: Chart of 50-run average.

The MDP-approach needs more time to learn closing price distributions as simple averages. When leaving away early runs, no statistical significant differences can be shown between these two approaches. We think that one reason that the MDP-approach did not outperform the "classical" approach is that flight tickets are currently not incorporated in the MDP-model. The agent buys flight-tickets according to a simple heuristic. Another reason is the fact, that the MDPs are coordinated only in a rudimentary way. Future research will focus on the coordination of the MDPs. It has to be figured out, whether the proposed mechanism is adequate for the TAC-scenario.

References

[1] Patricia Anthony and Nicholas R. Jennings. Developing a bidding agent for multiple heterogeneous auctions. *ACM Transactions on Internet Technology*, 3(3):185–217, 2003.

[2] Günther Bamberg and Adolf Coenenberg. *Betriebswirtschaftliche Entscheidungslehre*. Vahlen, München, 2002. 11., überarbeitete Auflage.

[3] Craig Boutilier, Thomas Dean, and Steve Hanks. Decision-theoretic planning: Structural assumptions and computational leverage. *Journal of Artifical Intelligence Research*, 11:1–94, 1999.

[4] Craig Boutilier, Moises Goldszmidt, and Bikash Sabata. Sequential auctions for the allocation of resources with complementarities. In *Proceedings of the 16th International Joint Conference on Artificial Intelligence (IJCAI-99)*, pages 527–534. Morgan Kaufmann Publishers, 1999.

[5] Andrew Byde. A dynamic programming model for algorithm design in simultaneous auctions. In Ludger Fiege, Gero Mühl, and Uwe Wilhelm, editors, *Electronic Commerce: Second International Workshop, WELCOM 2001*, volume 2232 of *Lecture Notes in Computer Science*, pages 152–163, 2001.

[6] Andrew Byde. A comparison among bidding algorithms for multiple auctions. In Julian Padget, Onn Shehory, David Parkes, Norman Sadeh, and William E. Walsh, editors, *Agent-Mediated Electronic Commerce IV. Designing Mechanisms and Systems: AAMAS2002 Workshop on Agent-Mediated Electronic Commerce Bologna. Revised Papers*, volume 2531 of *Lecture Notes in Artificial Intelligence*, pages 1–16, Berlin, 2002. Springer.

[7] Andrew Byde, Chris Preist, and Nicholas R. Jennings. Decision procedures for multiple auctions. In *Proceedings of the first international joint conference on Autonomous agents and multiagent systems*, pages 613–620. ACM Press, 2002.

[8] Shih-Fen Cheng, Evan Leung, Kevin M. Lochner, Kevin O'Malley, Daniel M. Reeves, L. Julian Schvartzman, and Michael P. Wellman. Walverine: A walrasian trading agent. In *Second International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-03)*, pages 465–472, 2003.

[9] Didier Dubois and Henry Prade. Possibility theory as a basis for qualitative decision theory. In Chris Mellish, editor, *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence*, pages 1924–1930, San Francisco, 1995. Morgan Kaufmann.

[10] Marlon Dumas, Lachlan Aldred, Guido Governatori, Arthur ter Hofstede, and Nick Russell. A probabilistic approach to automated bidding in alternative auctions. In *Proceedings of the eleventh international conference on World Wide Web*, pages 99–108. ACM Press, 2002.

[11] Peyman Faratin, Carles Sierra, and Nicholas R. Jennings. Negotiation decision functions for autonomous agents. *Int. Journal of Robotics and Autonomous Systems*, 24(3-4):159–182, 1998.

[12] Pere Garcia, Eduard Giménez, Lluís Godo, and Juan A. Rodríguez-Aguilar. Possibilistic-based design of bidding strategies in electronic auctions. In *The 13th European Conference on Artificial Intelligence (ECAI-98)*, pages 575–579, 1998.

[13] Amy Greenwald. The 2002 trading agent competition: an overview of agent strategies. *AI Magazine*, 24(1):83–91, 2003.

[14] Minghua He and Nicholas R. Jennings. Southamptontac: An adaptive autonomous trading agent. *ACM Transactions on Internet Technology*, 3(3):218–225, August 2003.

[15] Paul Klemperer. Auction theory: A guide to the literature. *Journal Of Economic Surveys*, 13(3): 227–284, 1999.

[16] R. Preston McAfee and John McMillan. Auctions and bidding. *Journal Of Economic Literature*, 25(2): 699–738, 1987.

[17] Nicolas Meuleau, Milos Hauskrecht, Kee-Eung Kim, Leonid Peshkin, Leslie Pack Kaelbling, Thomas Dean, and Craig Boutilier. Solving very large weakly coupled markov decision processes. In *Proceedings of the fifteenth national/tenth conference on Artificial intelligence/Innovative applications of artificial intelligence*, pages 165–172. American Association for Artificial Intelligence, 1998.

[18] Chris Preist, Claudio Bartolini, and Andrew Byde. Agent-based service composition through simultaneous negotiation in forward and reverse auctions. In *Proceedings of the 4th ACM conference on Electronic commerce*, pages 55–63. ACM Press, 2003.

[19] Chris Preist, Claudio Bartolini, and Ivan Phillips. Algorithm design for agents which participate in multiple simultaneous auctions. In Frank Dignum and Ulises Cortés, editors, *Agent-Mediated Electronic Commerce III : Current Issues in Agent Based Electronic Commerce Systems*, volume 2003 of *Lecture Note in Artificial Intelligence*, pages 139–154, Berlin, 2001. Springer.

[20] Chris Preist, Andrew Byde, and Claudio Bartolini. Economic dynamics of agents in multiple auctions. In *Proceedings of the fifth international conference on Autonomous agents*, pages 545–551. ACM Press, 2001.

[21] Chris Preist, Andrew Byde, Claudio Bartolini, and Giacomo Piccinelli. Towards agent-based service composition through negotiation. *AISB Journal*, 1(1):109–124, 2001.

[22] Martin L. Puterman. *Markov decision processes: discrete stochastic dynamic programming.* Wiley, New York, 1994.

[23] P. Stone, J. Schapire, M. Littman, J. Csirik, and D. McAllester. Decision theorethic bidding based on learned density models in simultaneous, interacting auctions. *Journal Of Artificial Intelligence Research*, 19:209–242, 2003.

[24] Michael P. Wellman, Daniel M. Reeves, Kevin M. Lochner, and Yevgeniy Vorobeychik. Price prediction in a trading agent competition. *Journal Of Artificial Intelligence Research*, 20:1–17, 2004.