Protocols and Strategies for Automated Multi-Attribute Auctions[1]

Esther David Rina Azoulay-Schwartz Sarit Kraus

Bar-Ilan University
Department of Computer Science
Ramat-Gan 52900
Israel
{saadone,schwart,sarit}@macs.biu.ac.il

ABSTRACT
In this paper, we consider a model of a procurement multi-attribute auction in which the sales item is defined by several attributes, the buyer is the auctioneer, and the sellers are the bidders. Such domains include auctions on task allocation, services, or compound items. The buyer announces a scoring rule, according to its preferences, before the auction starts, and each seller places a bid, which describes the attributes of the item it offers for sale.

First, we consider a variation of the first-price sealed-bid protocol, and we provide optimal and stable strategies for the buyer agent and for the seller agents participating in the multi-attribute auction. In addition, we analyze the buyer's expected revenue and suggest an optimal scoring rule that can be announced. Second, we consider four variations of the English auction for the case of a multi-attribute item, and we prove that, given some assumptions, they all converge to the same result. We also discuss which variation is preferred for different types of environments. Moreover, we show under which conditions, announcing the truth about buyer preferences is the optimal strategy for the buyer.

Categories & Subject Descriptors
Intelligent agents, Multiagent systems.

General Terms
Economics.

Keywords
Bidding and bargaining agents. Electronic commerce market.

1. INTRODUCTION
Recently, auction mechanisms have become a fundamental way of automating negotiation among agents [4,9,10,12]. The widespread research on automated auctions deals mostly with models in which price is the unique strategic dimension. However, in many situations, it is necessary to conduct negotiations on multiple attributes of a deal. For example, in task allocation, the attributes of a deal include the size of a task, starting time, ending deadline, accuracy level, etc. A service can be characterized by its quality, supply time, and the risk involved, in case the service is not supplied eventually. In the commercial domain, an item can be characterized by several attributes, such as size, weight, supply date, etc.. Auctions are efficient protocols for reaching agreements among agents, and they can also be used when the issue to be considered is associated with multi-attributes. However, several difficulties arise when trying to implement the traditional single-attribute auction protocols for multi-attribute items. In this paper, we address these difficulties and suggest how to design and implement automated auctions for multi-attribute items. We also provide the automated agents that participate in such auctions with stable and efficient strategies to be used in the auction.

We focus on markets in which an agent that wants to buy an item becomes the auctioneer. At the beginning of the auction, the buyer announces the required properties of the item, and then various sellers propose specific configurations that match its request. Finally, the buyer decides which bid it prefers, and the agent that suggested this bid, called the winner agent, is committed to providing it. Such markets appear in several situations, and currently there is no automated mechanism to deal with them. For example, telephone service providers and Internet portals, as well as video-on-demand suppliers, would like to rent extra storage capacity from suppliers over the Internet. The attributes of the required item in this domain are the storage capacity, the access rates to the data, the availability period and time limit, the level of security, etc. Printing a file is categorized by the quality of the print, the quality of paper, and the time deadline. Another particular domain of such markets is the International Logistics Supply Chain (ILSC). In this domain, the

[1]This work is supported in part by NSF Grant IIS 9820657 and IIS0208608.
The third author is also affiliated with UMIACS.

required service moves a given cargo from one location to another. The details of this service include arrival time, dispatching time, path length, weight, volume, etc.

- The questions relating to the agents' strategies in a multi-attributes auction are:
- What should the buyer (auctioneer) reveal at the start of an auction? Should it be all preferences, only part of them, or false preferences? How does the amount or the type of the announced preferences influence the auction results?
- Should the buyer be committed to the preferences it announces and should the seller agent be committed to its proposed bid? What should the winner selection criterion include?
- How should a seller agent formulate its bid considering the various attributes? What should the optimal bid/bids of each seller be, given the protocol, and its beliefs?
- Assuming that an ascending protocol is used, how can a bidder suggest a better bid than the current best bid, if it does not completely know the buyer's preferences?

In this paper, we address these various issues and propose ways to handle them. We suggest two types of protocols for multi-attribute contracts. First, we consider a variation of the first price sealed bid protocol, and we provide optimal and stable strategies for the buyer and for the sellers participating in such auctions. Second, we consider four variations of the English (ascending) auction protocol, and we prove that, given some assumptions, they all converge to the same result. We also discuss which of the above variations is preferred for different types of environments.

2. RELATED WORK

An auction is an efficient protocol for reaching agreements among agents [4,9,10,12]. There are several types of auctions which are used, including the English auction, first-price sealed-bid auction, second-price sealed-bid (Vickery) auction, and the Dutch auction [10]. In this paper, we consider the first-price sealed bid and the English auction. In the first-price auction, each bidder places a sealed bid. The winner agent will be the one with the best bid, and it will be committed to its bid. In the English auction, during the bidding process, each bidder can suggest a bid better than or equal to the last bid, and when no more bids are suggested, the last placed bid wins. In real world situations, each auction has its advantages and drawbacks.

Auctions can also be used when the issue to be considered includes more than one attribute. Researchers discuss auctions n which a bid is composed of several details, but mostly in the context of a special type of auction, called combinatorial auctions [1,6,8,9,11]. In this model, a set of available goods is given, and each bidder specifies bundles of goods and the prices it is willing to pay for each specified bundle. The problem that arises in this type of auction is how to determine which agents will obtain the bundles they ask for, since the amount of available goods is limited. Although in combinatorial auctions each bid is composed of multiple details, the problems that appear in our model are completely different, since, the auction discussed in our model includes one buyer of one multi-attribute item and several competing sellers. This diversity causes completely different problems, when trying to automize the auction mechanism. In this paper, we address these problems and suggest how to design and implement automated auctions for multi-attribute items. We also provide the automated auctioneer and bidders with stable and efficient strategies to be used in this type of auction.

Gimenez-Funes et al. [4] developed trading agents for electronic auctions of multi-attribute items, where the seller suggested a given item for sale and the buyers compete by sending bids that are composed of the price of the item. They suggested that the agent apply the Case-Based Decision theory in order to decide which strategy to use in the auction. In this paper, we consider an auction in which the buyer is the auctioneer, and the sellers suggest items with multiple attributes. We have developed static and stable agents' bidding strategies that are based on their beliefs about the environment.

Bichler [2] has made an experimental analysis with professional bidders of multi-attribute auctions. He found that the utility scores achieved in multi-attribute auctions were significantly higher than those of single attribute auction. He found out also that, the second score auction format, which is a version of the second-price auction format that suits the case of multi-attribute auction, yields better results for the buyer than the English protocol format. In addition, he did not find evidence for revenue equivalence, for the auctioneer, between the various auction protocols. However, no theoretical explanation is given for these observations. In future work, we intend to analyze the second score protocol theoretically.

Very little theoretical work has been done on multi-attribute auction area, in which the work of Che [3] is the most advanced. Che considers an auction protocol in which a bid is composed of price and quality. In addition, he assumes that each seller is characterized by only one private cost parameter. In this paper, he proposed a design for first score and second score sealed bid auctions. In Section 4, we follow Che in the description of the first price sealed bid protocol. But, we have more than one quality dimension. Moreover, we consider the best strategy to be used from the seller's point of view.

3. THE MODEL

The auction model consists of one buyer, which is the auctioneer, and a fixed number of *n* sellers, which are the bidders. This type of auction is very useful in real world decisions about suppliers of services and in government decision-making processes. The buyer agent that needs a particular item (service or item) starts the auction process. At the beginning of the auction, the buyer announces its item request, which consists of the item's desired characteristics, the auction protocol, and a scoring rule, which describes its preferences concerning the item's properties. A seller agent, who decides to send a bid, has to specify the full configuration it offers.

Each buyer agent and each seller agent is characterized by a utility function that describes its preferences. The multi-attribute utility-functions we refer to are based on the Simple Additive Weighting (SAW) method [13]. A utility or a score in the SAW method is obtained by adding the contributions of each attribute. There are other methods for multi-attribute utility functions (e.g. multiplying the contributions of the various attributes). However, the SAW method suits the examples (e.g., ILSC domain or renting storage capacity) we refer to. The utility function of the buyer associates a value, which is the sum of the buyer's level of satisfaction from the various attributes' values, with each bid. The utility function of each seller associates a utility value reflecting the seller's cost and profit with each bid.

At the beginning of the auction, the buyer agent will announce a scoring rule that associates a score with each proposed offer, and the auction protocol dictates the winner (best scored bid) based on this scoring rule. The buyer agent tries to derive a scoring rule that maximizes its expected utility in a given auction protocol. Each participant knows its utility function, and time and bidding are not costly. In the following two sections we suggest different variations of auction protocols for the above model of the multi-attribute auction.

4. FIRST-SCORE SEALED-BID AUCTION

We start by suggesting a variation of the first-price sealed bid auction, for multi-attribute items. This type of auction is widely used in commercial and governmental auctions, because of its simplicity. It does not require a long process such as the English protocol, and it does not oblige the auctioneer to reveal the bids it obtained, as in the second-price sealed-bid auction.

According to the first-price sealed-bid auction protocol for one attribute, the bid with the highest price, in case of one seller and several buyers, wins, and the buyer is committed to its bid. However, if there are several attributes to negotiate, instead of referring only to the price, we need to refer to the overall offer.

4.1 Description of the Protocol and the Utility Functions

The extension of the first-price sealed-bid auction for a multi-dimensional auction, was developed by Che [3]. It is called the *first-score sealed-bid multi-dimension auction protocol.* According to this protocol, the buyer (auctioneer) initially announces a scoring rule, which will be used by the buyer (auctioneer) in order to evaluate the bids it will obtain. At the next step, each seller (bidder) submits a sealed bid that specifies the details of the item it suggests to supply. Finally, the winner agent is the seller that receives the highest score for its bid, according to the pre-announced scoring rule.

The auctioneer must be committed to its scoring rule. If a bid different than the best one was chosen, the bidder of the best bid can prove that it sent a better bid than the chosen one, and the auctioneer will have to change its decision. The winner agent is required to provide an item with the exact values of the bid it offered (e.g., the exact price, quality, delivery date, etc.).

In our model, each seller agent has private information about the costs of improving the quality of the item it sells, or its performance. In particular, each seller agent, S_i (bidder), is assumed to be characterized by a cost parameter θi, which is its private information. As θi increases, the cost to the seller for achieving an item of a high quality also increases; i.e., the seller is weaker. The buyer (auctioneer) knows only the distribution function of the other sellers' cost parameters, but has no information about the particular value of θi for each seller.

Similar to the model described by Che [3], we assume that θi is independently and identically distributed over $[\underline{\theta},\overline{\theta}]$ ($0<\underline{\theta}<\overline{\theta}<\infty$) according to a distribution function *F,* for which a positive, consciously differentiable density *f* exists. Because of complete symmetry among agents, the subscript *i* is omitted in the rest of the paper.

For simplicity, we analyze a specific model, which includes three attributes: the price *p* of the item, and two quality factors, q_1 and q_2 for which the preferences of the buyer and the sellers conflict. For example, in the ILSC domain, q_1 denotes the availability time, which is the time from within which a seller is committed to supply the item; and q_2 denotes the path length, which is the length of the way in which the item is transferred to its destination. The buyer prefers the availability time to be as early as possible. However, the seller prefers to supply the service as late as possible. Similarly, the buyer prefers the shortest possible path for security reasons, and the seller prefers a longer path, since this gives it more options in supplying the service.

We assume that there are fixed coefficients of each of the quality dimensions which are identical for all the sellers. Namely, 'a' is the coefficient of q_1, and 'b' is the coefficient of q_2.

The seller's cost function, is:

$C_s(q_1,q_2,\theta)=\theta(\frac{a}{q_1}+\frac{b}{q_2})$, where a,b>0. Based on the cost function, a seller's utility function is:

$$U_{si}(\theta)=p-\theta(\frac{a}{q_1}+\frac{b}{q_2}).$$

Notice that the utility function of the seller is the sum of its component, and as the payment it obtains increases, the utility increases. The above functions fit the case where as q_1 and q_2 increase and the quality of the item or service decreases, as in the ILSC example. Thus, as q_1 and q_2 increase, the cost of the seller decreases, and its utility increases. In domains where higher values of qi means a better item for the buyer, the coefficient of attribute *i* will be positive. For example, if q_1 means storage capacity, and q_2 means access speed to the data, Usi will be $p+\theta(1/q_1+1/q_2)$. Notice that the model can be easily extended to more than two quality dimensions, using a similar type of the utility function.

The influence of q_1 and q_2 is assumed to be independent, but they are not linear: as q_j increases, an additional increase will have a lower effect than will the initial increase. This assumption is valid in real world domains, since allowing more flexibility to the service supplied by the seller can have larger impact when moving from the most efficient service to something less efficient, though its impact will be lower as the efficiency is reduced. (For example, it may help the seller if it can transport cargo over a period of two days instead of one day, but the effect of 11 days instead of 10 is much less significant).

We also assume that the utility function of the buyer agent (the auctioneer) from the service is as follows:

$$U_{buyer}(p,q_1,q_2)=-p-W_1\cdot q_1-W_2\cdot q_2,$$

where W_1 and W_2 are the weights the buyer assigns to q_1 and q_2, respectively. (Without loss of generality, the function is assumed to be normalized by the weight of the price.) In fact, as the price decreases, the buyer's utility increases. The coefficients of $W_1 \cdot q_1$ and $W_2 \cdot q_2$ are negative in the ILCS models, but if a higher value of q_i means a higher quality item, then the coefficient of $W_i \cdot q_i$ will be positive. For example, if q_1 means storage capacity, and q_2 means access speed to the data, the buyer's utility will be $(-p + W_1 \cdot q_1 + W_2 \cdot q_2)$. In the rest of this section, we assume, without loss of generality, that as the value of q_1 and q_2, decreases, the quality of the item increases.

In the ILCS domain, the effect of the three attributes on the utility function of the buyer is linear. However, the effect of the availability time and path length is weighted by W_1 and W_2, respectively, where Wi can be smaller or larger than 1. Wi<1 means that the effect of qi is smaller than the effect of paying one additional unit of price, and vice versa.

4.2 Agents' Strategies

The buyer should announce a scoring rule, which is used for choosing among bids. The scoring rule of a buyer can be different than its real utility function, in the sense that the announced weights w_1, w_2 may be different than the actual weights W_1, W_2. In particular, we denote the scoring rule by: $S(p, q_1, q_2) = -p - w_1 \cdot q_1 - w_2 \cdot q_2$,

where w_1, w_2 are the weights that the buyer assigns to q_1 and q_2. From the scoring rule, we infer that the announced bid's value for the buyer is:

$V(q_1, q_2) = -w_1 \cdot q_1 - w_2 \cdot q_2$. The announced values of the weights w_1, w_2 can be equal to or different from the real values of weights W_1, W_2. For example, if $w_1 < W_1$, it means that, for some reason, the buyer declares a lower utility derived from each unit of q_1, with regard to its real utility from q_1. In the rest of this section, we will discuss how optimal announced weights can be determined.

Given the publicized scoring rule, each agent interested in selling the item will join the auction and send a sealed bid, describing its suggestion to satisfy the buyer agent's requirements. The sealed bid will be composed of three dimensions: price (p), q_1, and q_2.

In the following lemma, we reveal how each seller chooses the values of q_1 and q_2, given its type, and given a particular scoring rule.

Lemma 1

In first-score sealed-bid auction, the quality attributes q_1 and q_2 *are chosen independently of the price at* $q_1{}^*(\theta)$ *and* $q_2{}^*(\theta)$ *for all* $\theta \in [\underline{\theta}, \overline{\theta})$*, where,*

$$q_1{}^*(\theta) = \arg\max_{q_1} \{V(q_1, q_2) - C_s(q_1, q_2, \theta)\}, \text{ and}$$

$$q_2{}^*(\theta) = \arg\max_{q_2} \{V(q_1, q_2) - C_s(q_1, q_2, \theta)\}.$$

Lemma 1 describes how each bidder decides about a bid, given the announced scoring rule, and given the bidder's beliefs, about its cost parameter. The proof is similar to that of Che [3], but it considers the additional dimension of q_2. Because of space restriction, we do not include the proofs here. They can be found in [14].

From Lemma 1, we can infer that there is no loss of generality in restricting attention to $q_1{}^*(\theta)$ and $q_2{}^*(\theta)$ when searching for an equilibrium.

The following lemma explicitly specifies the values of p, q_1 and q_2, as proposed by the bidders, given the announced scoring rule, and given the model described in this section.

Lemma 2

A symmetric equilibrium of a first-score auction in the model described above is one in which each seller bids:

$$q_1{}^*(\theta) = \sqrt{\frac{\theta \cdot a}{w_1}}, \quad q_2{}^*(\theta) = \sqrt{\frac{\theta \cdot b}{w_2}},$$

$$p^*(\theta) = (\sqrt{a \cdot w_1} + \sqrt{b \cdot w_2}) \cdot \left(\sqrt{\theta} + \frac{1}{(\overline{\theta} - \theta)^{n-1}} \cdot \int_{\theta}^{\overline{\theta}} \frac{(\overline{\theta} - t)^{n-1}}{\sqrt{t}} dt\right).$$

The seller agent decides about its bid according to its private costs in producing the item, the scoring rule, and its beliefs about the other sellers. We can see from Lemma 2 that its beliefs about other agents will influence only the price it proposes. For example, when there are more bidder agents, its price will decrease, and this is demonstrated in Figure 1, for specific values of parameters ($w_1 = w_2 = 1$, a=b=1, and the distribution F of θ is uniform, where $\underline{\theta} = 0.5, \overline{\theta} = 1$).

However, the buyers' decision about the quality of the service is independent of these beliefs. As the publicized weight wi (i=1,2) increases, the quality of the goods, concerning q_i, increases (since $q_i{}^*$ will be smaller), but the price of the bid increases, since p* increases. Also, as the private cost parameter θ increases, and the seller is less efficient, it will suggest lower quality items (with higher values of $q_1{}^*$ and $q_2{}^*$), for higher prices (higher value of p*). This can be inferred from the formulas of $q_1{}^*, q_2{}^*$ and p*: as θ increases, the nominators of $q_1{}^*$ and $q_2{}^*$ increase, so their values increase. Also, as θ increases, it is also clear that p* increases, since this has two positive effects on p* (a decrease of the denominator, and an increase of the root of θ).

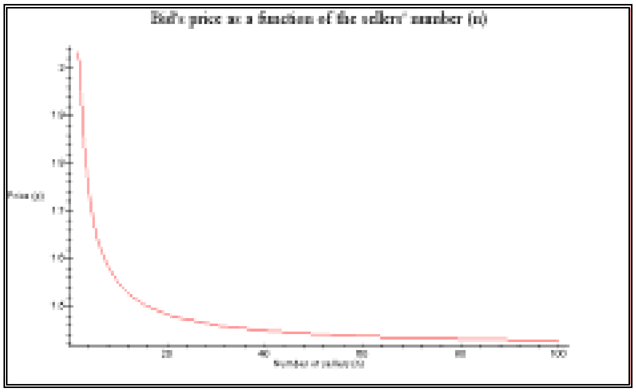

Figure 1. Bid price as a function of the number of bidders.

In Lemma 3 below, the value of the buyer's expected revenue of the first-score sealed-bid auction ER^1 is determined given its beliefs about the types of sellers (bidders).

Lemma 3

The buyer's expected revenue ER^1, given W_1, W_2, and the announced values w_1, w_2 of the scoring rule, is:

$$ER^1(\underline{\theta},\overline{\theta}) = \frac{-n\cdot(\sqrt{a\cdot w_1}+\sqrt{b\cdot w_2})}{(\overline{\theta}-\underline{\theta})^n}\cdot\int_{\underline{\theta}}^{\overline{\theta}}\left(\sqrt{t}\cdot((\overline{\theta}-t)^{n-1}+(\int_t^{\overline{\theta}}\frac{(\overline{\theta}-z)^{n-1}}{\sqrt{z}}dz)\right)dt - \frac{n}{(\overline{\theta}-\underline{\theta})^n}\cdot(\frac{W_1\cdot\sqrt{a}}{\sqrt{w_1}}+\frac{W_2\cdot\sqrt{b}}{\sqrt{w_2}})\cdot\int_{\underline{\theta}}^{\overline{\theta}}\sqrt{t}\cdot((\overline{\theta}-t)^{n-1}dt.$$

Sketch of Proof

The buyer's expected revenue ER^1 is simply the expected *actual* value of the highest bid among the n sellers. Computation is based on [5]. ■

Based on Lemma 3, it is clear that the values of the announced w_1 and w_2 will influence the expected revenue of the buyer. In the following theorem, we specify the optimal values of w_1 and w_2 to be announced, given the actual weights W_1 and W_2, and given the distribution of θ.

Theorem 1

The optimal values for the announced weights w_1, w_2 are:

$$w_i(\overline{\theta},\underline{\theta}) = W_i\cdot\frac{\int_{\underline{\theta}}^{\overline{\theta}}\sqrt{t}\cdot(\overline{\theta}-t)^{n-1}dt}{\int_{\underline{\theta}}^{\overline{\theta}}(\sqrt{t}\cdot(\overline{\theta}-t)^{n-1}dt+\int_{\underline{\theta}}^{\overline{\theta}}(\int_t^{\overline{\theta}}\frac{(\overline{\theta}-z)^{n-1}}{\sqrt{z}}dz)dt}$$

Sketch of proof

In order to find the weights $w1,w2$ that maximize the expected revenue, we first derive the function of ER^1 by w_1 and find its maximum value. Then, we derive this function by w_2 in a similar manner since the weights w_1, w_2 are independent. ■

Based on the above results, given the number of sellers, the distribution of θ, and the sellers' optimal strategies, the buyer can announce the optimal scoring function that will optimize its expected revenue from the auction's result. Notice that the ratio between w_1 and w_2 remains the same as the ratio between $W1$ and $W2$. Due to this property, for each bid, the ratio of q_1 and q_2 remains equivalent to their ratio, given the actual weights (according to Lemma 2). The only difference is in the prices with regard to the qualities. If $wi<Wi$, the price will be lower than the price given the actual weights, but the qualities will be lower, too (i.e., higher cardinal values), and vice versa.

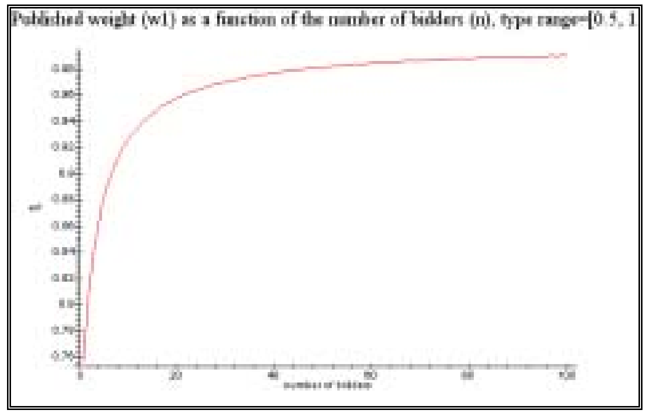

Figure 2. w_1 as a function of the number of bidders, types of sellers in the range ([0.5, 1]).

Figure 2 demonstrates the influence of n on $w1$, as $W1$=1, and the values of θ are between 0.5 and 1. It can be shown that then as n increases, w_1 increases, and it approaches 1 for higher values of n; i.e., as the number of bidders increases, the buyer is more motivated to announce a scoring function close to its utility function, so the announced values of w_1 and w_2 converge to the real weights W_1 and W_2.

The specific values of θ do not influence the scoring rule, but the relationship between $\overline{\theta}$ and $\underline{\theta}$. As the relationship increases, the slower the scoring function converges to the real utility function. For example, in Figure 2, the relation between $\overline{\theta}$ and $\underline{\theta}$ ([0.5, 1]) is 2, while in Figure 3 the relation between $\overline{\theta}$ and $\underline{\theta}$ ([0.1, 0.5]) is 5. Note that in Figure 2, the announced values of w_1 and w_2 converge closer to 1 than do the values of w_1 and w_2 in Figure 3.

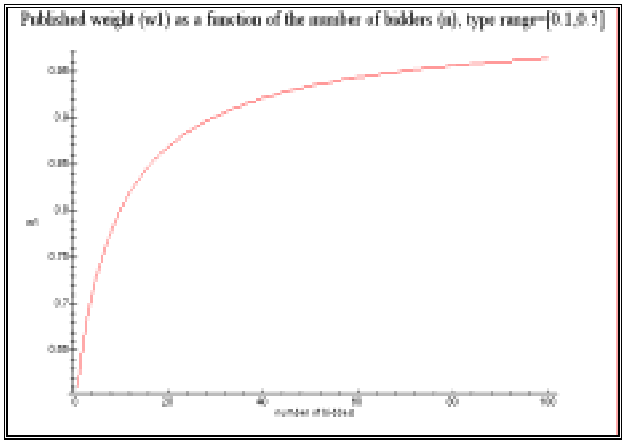

Figure 3. w_1 as a function of the number of bidders, types of sellers in the range ([0.1, 0.5]).

In Figure 4, we demonstrate the influence of the ratio between $\overline{\theta}$ and $\underline{\theta}$, on w_1. The parameters set to n=4, $\underline{\theta}$ =1, W_1=1, while the values of $\overline{\theta}$ varied from 2 to 20. As expected, w_1 decreases as the ratio between $\overline{\theta}$ and $\underline{\theta}$, increases.

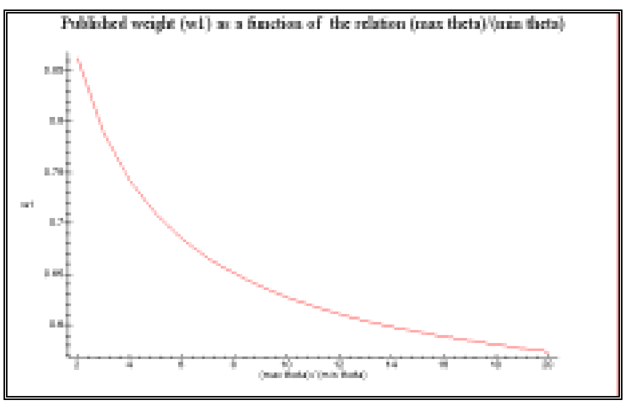

Figure 4. w_1 as a function of the ratio between $\overline{\theta}$ and $\underline{\theta}$.

To summarize, in this section we propose a sealed bid auction protocol for an automated auction in which the auctioneer is the buyer. We specify the utility function of the buyer and that of the seller. We reveal how each bidder can bid optimally, given its cost parameter θ_i and the distribution of θ. We proceed by evaluating the buyer's expected revenue, and based on it, we reveal the optimal scoring function to be announced by the buyer, in order to maximize its expected revenue.

The above protocol and strategies can be used for automated auctions on multi-attribute items in real world domains. Another type of well-known auction protocol is the English auction. In the next section, we will consider several variations of the English auction and discuss the optimal strategies of the buyer and the sellers, given this protocol.

5. THE ENGLISH AUCTION PROTOCOL

In this section, we suggest an additional type of protocol, based on the English protocol. The advantages of the English protocol are: (a) The bidders have no motivation to manipulate and change bids according to their beliefs about other agents. (b) English auction is the best preferred protocol if the valuation problems of the agent are difficult [7], and finding the best possible multi-attribute bid can be a difficult problem. (c) Some variations of the English protocol allow the buyer's preferences regarding the required item to remain partially hidden.

In this section, we define the number of attributes to be unlimited, but this number is fixed and known in advance. We discuss four variations of the English protocol, and we present some interesting results concerning these variations.

5.1 Protocols Description

The different variations of the English protocol differ in the auction details. Two protocols are *full-information-revelation protocols*, and two are *directed protocols*. In the full-information-revelation protocol, the auctioneer announces its scoring rule, similar to the sealed bid auction of the previous section, and it is committed to this scoring rule. Thus, each seller can evaluate the score of each bid according to the announced scoring rule and, accordingly, decide about its bids. This protocol, combined with the scoring rule, becomes similar to the traditional English auction.

But, generally, agents are not willing to expose their private information. Thus, we propose two additional variations, called *directed* protocols, in which the auctioneer is required to announce some basic information about its scoring rule, termed the *minimal set of constraints,* to enable the sellers to form relevant bids. Most of the buyer's preferences remain hidden, and during the auction process, the buyer directs the auction using its hidden preferences, by announcing the winner in each step. Notice that this protocol is meaningless in the case of a single attribute auction, since the auctioneer's preferences are clear; e.g., it prefers the price to be as low as possible. The participants can evaluate the bids by themselves. Only in the case of multi-attribute auction, if the bidders do not have full information about the buyer's preferences, is this protocol relevant.

For example, in the ILSC domain the minimal set of constraints is comprised of: the cargo's weight, the cargo's source, the cargo's destination, and the arrival time deadline. The hidden preferences consider the maximum price, whether swapping planes is allowed, and which transportation means are preferred. In the computer storage rent domain, the set of constraints includes the required storage capacity and the time for which the service is required. The hidden properties are the minimum required speed and the maximum available price. The advantages of a full-information-revelation protocol are that it takes less time, fewer computation efforts, and fewer communication costs, and that the received bids are much more relevant.

In automated auctions, participants are not physically located in the same place and usually they are connected by a communication medium. Therefore, we consider two different rules that define the order of bidding. In a *sequential* protocol, each seller is assigned a serial number and it can bid when its turn arrives. In the *simultaneous* protocol, all interested sellers can bid simultaneously in each step. In summary, the four protocols are: (1) the simultaneous full-information-revelation auction, (2) the sequential full-information-revelation auction, (3) the simultaneous directed auction, and (4) the sequential directed auction.

At the beginning of each auction the buyer announces the length of the time interval, during which, if no new bid is made, the auction is closed. This interval is termed the *closing time*. Given the buyer agent's scoring function, we consider *two bids to be equivalent* if they yield the same utility to the buyer. In order to avoid any possible influence of the serial number a seller receives on whether or not he wins, and to prevent the bidders from being motivated to manipulate and change bids according to their beliefs about other agents, bidders are allowed to make a bid that is equivalent to a proposed bid in each step.

We define the *selected bid* to be the bid which is selected by the auctioneer in each step. In a sequential protocol, if there are sellers who want to propose a bid equivalent to the previous bid, they can propose it during the closing time, from the moment that the first equal bid was proposed. In case of a simultaneous protocol, the buyer waits for another closing time from the moment a selected bid is announced, in order to enable other sellers to bid equivalent bids. If no higher bid w.r.t. the auctioneer's announced scoring rule is proposed in the meantime, then the auctioneer makes a decision and announces which of the equal bids becomes the selected bid.

For example, we describe in detail the sequential directed protocol. In this protocol the, auctioneer directs the auction in each step by determining whether the new proposed bid is better than the selected bid.

Sequential Directed English auction:

Step 1: The buyer sends each of the interested sellers: (a) a minimal set of constraints, (b) the closing time, and (c) an individual serial number for each of the interested sellers.
Step 2: Each seller can submit a new bid at its turn. Each time that a bid is submitted the auctioneer must choose the best bid among the current *selected bid* and the new proposed bid.
Step 3: If a *closing time* has passed and no new bid was proposed, the auction is closed. The winning seller is the owner of the *selected bid.*

5.2 Strategies of the Sellers and the Buyer

In this section, we provide the seller agents and the buyer agent with tools to participate efficiently in one of the English protocols we consider. We start by describing the sellers' optimal strategy, given the buyer's scoring function.
We assume that each seller has a list of bids ordered in a non-increasing order according to the given seller's utility function. That is, for seller S_i with utility function U_{S_i}, there is a list of proposals $(P_{i1}, P_{i2}, \ldots, P_{ik_i})$, where $U_{si}(P_{ij}) \geq U_{si}(P_{i\,j+1})$ for $1 \leq j < k_i$. This list is termed the *bid-list of seller* S_i. In case of a full-information-revelation protocol, the *next bid* of seller S_i is defined to be the next bid in seller S_i's bid list that is better or equal to the *selected bid* with respect to the buyers scoring rule. (Note that since each item is multi-attribute, the order according to *Usi* is not necessarily completely opposite to the order of the bids according to Ubuyer). In case of a directed auction protocol, the *next bid* is defined to be the next bid in the seller's bid list, which has not yet been proposed, and which is not worse for the buyer than the *selected bid*, according to the *minimal set of constraints*. The following lemma describes the optimal strategy for the sellers participating in one of the four versions of the English protocols.

Lemma 4:

Given one of the English auction variations, at each time point when bids can be placed, the best strategy of seller S_i is to bid its next bid in case it is not the bidder of the current selected bid. If the current selected bid belongs to S_i, the best strategy is not to bid at all.

See Example 1 below, which illustrates the differences among the various protocols. Recall that each seller S_i has a bid-list $(P_{i1}, P_{i2}, \ldots, P_{ik_i})$ ordered in a non-increasing order by *Us*. Denote ***Lall*** to be the list of t bids, where t is: $\sum_{i=1}^{n} k_i = t$, ordered by the scoring rule function in a non-decreasing order: *{P1, P2, ...,Pt}*, where for each bid P_j there exist *i* and *l* such that $1 \leq i \leq n$ and $1 \leq l \leq k_i$ that $P_j = P_{il}$. Notice that if the scoring rule is identical to the buyer's utility function then Pt is the best bid from the buyer's point of view. Otherwise, Pt is the best bid according to the announced scoring rule, which the buyer is committed to.
We define the ***Tail_of_Lall*** to be the sequence of bids that belong to the same seller and which appear at the end of Lall, such that it does not include a bid that has another bid equal to it before this sequence. We define the rest of Lall to be the ***Head_of_Lall***. That is, ***Head_of_Lall*** *U* ***Tail_of_Lall*** *= Lall.* If the bid at the end of ***Lall*** has an equal bid that belongs to another bidder, then the ***Tail_of_Lall*** is empty. We define the set of equal bids appearing at the end of the ***Head_of_Lall*** (if they exist) to be the ***equal-set***. The following theorem specifies the winning bid, given Lall.

Theorem 2
Assume that the sellers participating in an auction according to one of the protocols described above, follow the strategy defined in Lemma 4. Then, the result of the auction is independent of the auction protocol. Thus, it follows that:

If ***Tail_of_Lall*** $=\phi$*, then winning_bid* $\in$ *equal_set.*

If ***Tail_of_Lall*** $\neq\phi$*, then suppose that* ***Tail_of_Lall*** $\subseteq bid\_list_{sj}$ *(of seller j). In this case, (a) or (b) holds.*

(a) winning_bid $\in$ ***Tail_of_Lall****, such that* $\forall p_{jk} \in$ ***Tail_of_Lall****,*

Usj(winning_bid) $\geq$ *Usj(pjk).*

(b) winning_bid $\in$ *equal-set s.t. winning_bid* $\in bid\_list_{sj}$*.*

Theorem 2 indicates that the winning bid is located in the ***Tail_of_Lall***, or in the *equal-set*. In case the winning bid is located in the tail, we show that the winning bid is the bid that yields the best utility for its seller among the bids in the tail. That is, suppose that the ***Tail_of_Lall*** belongs to seller S_j and it includes bids *{Pjk1, Pjk2 ...Pjkn}* where kn is the length of the tail. Suppose that the best bid in the tail according to *Usj* is *Pjkm* where m is in the range of *[1,..,n]*. Then, the winner is seller S_j with bid *Pjkm*. In case the winning bid is located in the equal–set, then if the tail is empty, the winning bid is one of the equal bids and if the tail is not empty, the winner is the bidder of the bids in the ***Tail_of_Lall***.
Notice that we have actually showed that the winner is determined independently of the order of the bidding in the case of a sequential protocol and also independently of the sellers' utility functions. That is, the buyer's scoring rule function

actually directs the auction process. Based on Theorem 2, if the sellers follow their proposed strategies in Lemma 4, then all four protocols will converge to the same point. Thus, all of them yield the same profit for the buyer and for the sellers.
In the following example, we demonstrate a case in which the auction process converges to a bid in the ***Tail_of_Lall***.

Example 1:

There are one buyer and two sellers, S_a and S_b. Assume that the following non-increasing bid-lists are ordered by the utility functions of the sellers accordingly:

$$bid\_list_{Sa} = \{p_{a1}, p_{a2}, p_{a3}\}, bid\_list_{Sb} = \{p_{b1}, p_{b2}\}.$$

Assume that the following ***Lall*** is ordered by the announced scoring rule: $L_{all} = \{p_{a1}, p_{b1}, p_{b2}, p_{a2}, p_{a3}\}$, where

$$U_{buyer}(p_{a2}) \neq U_{buyer}(p_{b2})$$
$$\Rightarrow Tail\_of\_L_{all} = \{p_{a2}, p_{a3}\}.$$

Table 1 shows the auction result for the different auction protocols. Notice that the winning bid is *pa2*, which is indeed located in the ***Tail_of_Lall*** and is better than the other bids in the tail with regards to the utility function of seller S_a.

Table 1. The differences among the protocols and the winning bid.

Auction protocol	Proposed bids	Selected bid	Winning bid
Simultaneous full information	*{(pa1,pb1), (pa2)}*	{pb1,pa2}	p_{a2}
Sequential full-information	$\{p_{a1}, p_{b1}, p_{a2}\}$	$\{p_{a1}, p_{b1}, p_{a2}\}$	p_{a2}
Simultaneous directed	*{(pa1,pb1),(pa2), (pb2)}*	$\{p_{b1}, p_{a2}, p_{a2}\}$	p_{a2}
Sequential directed	*{pa1,pb1,pa2, pb2}*	$\{p_{a1}, p_{b1}, p_{a2}, p_{a2}\}$	p_{a2}

The next question to be addressed is: which scoring rule will yield the best result for the buyer? The following lemma characterizes the cases in which announcing a scoring rule different from its real utility function may yield a better result for the buyer.

Lemma 5:
Given a scoring rule=buyer's utility function, if $|Tail\_of\_L_{all}| > 1$*, then by announcing a*

$$Scoring\_rule \neq Buyer's\_utility\_function,$$

the buyer may reach a better deal than if it announces its utility function.

The following example demonstrates a case in which the best strategy for the buyer is to use a scoring rule different from the real utility function.

Example 2:

There are one buyer and two sellers *Sa* and *Sb*. Assume that the bid includes two attributes: price (p) and quality (q).
Given the following utility function of the sellers and two versions of scoring function,

$$Usa = p - 0.5 \cdot q;$$
$$Usb = p - q;$$
$$scoring\_rule = Ubuyer = q - p$$
$$scoring\_rule \neq Ubuyer = 0.5 \cdot q - p,$$

the bid list of *Sa* is *{pa1,pa2,pa3}*, and the bid list of *Sb* is *{pb1,pb2}*, where the explicit bids are the sellers' utility values and their scores are as follows:

Table 2. The values of the bids' attributes considered in Example 2.

Bids' ID	price	quality	Usa/ Usb	Score= Ubuyer	Score≠ Ubuyer
Pa1	10	4	8	-6	-8
Pa2	5	2.5	3.75	-2.5	-3.75
Pa3	3	1	2.5	-2	-2.5
Pb1	9	5	4	-4	-6.5
Pb2	4	1	3	-3	-3.5

Assume an auction that follows the sequential full-information-revelation protocol. The following table describes the winning bid for two different scoring rules.

Table 3. Comparing two different scoring rules.

scoring rule	*Lall* increasing lists ordered by the scoring rule	*Tail_of_ Lall*	winning bid
True	*{pa1,pb1,pb2,pa2,pa3}*	*{pa2,pa3}*	*pa2*
False	*{pa1,pb1,pa2,pb2,pa3}*	*{pa3}*	*pa3*

According to the above table, if the buyer announces its utility function as the scoring rule, the winning bid will be *pa2*, which is less preferred by the buyer than the best possible bid. However, the buyer can distort its scoring rule in order to achieve *pa3*, which is its best possible bid. In the following lemma, some cases are given in which announcing a scoring rule, which is identical to the buyer's utility function is optimal.

Lemma 6:
In the following two cases, the buyer has no incentive to announce a scoring rule that is different than its utility function: (a) The tail contains one bid and the equal set does not contain a bid of the seller to whom the tail belongs. (b) The tail is empty.

In other cases, speculative scoring rules can be announced. In Section 4.2 we discuss how to evaluate a scoring rule for the first-score sealed-bid protocol. Similar methods can be used in order to find the best scoring rule for the English protocol.
In summary, we have proposed four auction protocols, which are variants of the English auction. We suggest how the agents should behave during the auction, and we have proved that the four auction protocols converge to the same result and that they are equivalent in the sense of the winning bid. Finally, we have found

situations in which the optimal buyer's strategy is to announce a scoring rule, which is identical to the buyer's utility. It seems that from the buyer's point of view, using the directed protocols is better, especially in case of commerce or other domains in which privacy is important. In case of many participants (e.g. in the Internet environment), or many possible offers, the choice of a simultaneous protocol is better than a sequential one. In the following table, we summarize the advantages and the disadvantages of the different auction types we have considered.

Table 3. Advantages and disadvantages of the 4 variations of the English protocol.

Protocol type	Advantages	Disadvantages
Sequential	Fewer computations and communications	More time
Simultaneous	Less time	More computations and communications
Full information revelation	Less time, fewer computations and communications	Privacy is not saved
Directed	Privacy is kept	Expensive in time, computations and communications

6. CONCLUSION

In this paper, we consider the problem of multi-attribute auctions, in which the auctioneer is the buyer of an item or a service and in which different sellers bid and offer different configurations of the desired item or service. This type of auction is widely used in services (such as cargo deliveries), by government auctions, etc., and may be very useful in solving resource allocation problems. However, because of its complexity, it is not yet automated. In this paper, we present the first attempt to automate this type of auction. We suggest two main protocols for the multi-attribute auction. The first is a variation of the first-price sealed-bid auction, and the second includes four possible variations of the English (ascending) auction protocol. We describe each protocol in detail and find stable and beneficial strategies for the buyer agent and for the seller agents. Our protocols and strategies can be used automatically in electronic markets of multi-attribute items or services, and they are stable and efficient.

In future work, we intend to consider a situation in which two private cost parameters, θ_1 and θ_2, associated with each seller. We also intend to find the optimal scoring function to be announced by the buyer at the start of English auctions, given a full information revelation protocol, and also to find its optimal behavior in the directed auction. Finally, we will compare the first price auction and the English auction with respect to the auctioneer's utility.

REFERENCES

[1] Andersson, A. Tenhunen, M. and Ygge, F. *Integer programming for combinatorial auction winner determination*. In *ICMAS-00*, Boston, MA, 2000, 39-46.

[2] Bichler, M. An experimental analysis of multi-attribute auction. *Decision Support Systems* Vol. 29, 2000, 249-268.

[3] Che, Y. K. Design competition through multidimensional auctions, *RAND Journal of Economics*, Vol. 24, 1993, 668-680.

[4] Gimenez-Funes, E., Godo, L., Rodriguez-Aguilar, J. A. and Garcia-Calves, P. Designing bidding strategies for trading agents in electronic auctions. In *ICMAS-98,* Paris, France, 1998, 136-143.

[5] Flavio, M. and Monteiro, P. Auctions with endogenous participation*, Rev. of Economic Design* 5, 2000, 71-89.

[6] Nisan, N. Bidding and allocation in combinatorial auctions. In *Proc. of the ACM Conference on Electronic Commerce (ACM-EC)*, Minneapolis, MN, 2000, 1-12.

[7] Parkes, D. C., Ungar, L. H., and Foster, D. P. Accounting for cognitive costs in on-line auction design. In *Agent Mediated Electronic Commerce (LNAI 1571),* P. Noriega and C. Sierra, Eds. Springer-Verlag, 1999, 25-40.

[8] Parkes, D. C., and Ungar, L. H. Iterative combinatorial auctions: Theory and practice. In *AAAI 2000*, 74-81.

[9] Rosenschein, J., and Zlotkin, G. *Rules of encounter: designing conventions for automated negotiation among computers*. Cambridge, Mass. MIT Press, 1994.

[10] Sandholm, T.W. Limitations of the Vickrey Auction in Computational Multiagent Systems. In *proc. of ICMAS-96* 1996, 299-306.

[11] Sandholm, T.W. *An algorithm for optimal winner determination in combinatorial auctions*. In *IJCAI-99*, 1999, 542-547.

[12] Walsh, W.E., Wellman, M. P., Wurman, P.R. and MacKie-Mason, J. K. Auction protocols for decentralized scheduling. In *Proc. 18th International Conference on Distributed Computing Systems*, 1998.

[13] Yoon, K. and Hwang, C. *Multiple attribute decision making: an introduction.* Thousand Oaks: Sage. 1995.

[14] A long version of the paper exists at: http://www.cs.biu.ac.il/~schwart/papers/Auction-full-paper.doc