Strategic bidding for multiple units in simultaneous and sequential auctions

Stéphane Airiau & Sandip Sen & Grégoire Richard
Mathematical & Computer Sciences Department
University of Tulsa
e-mail:stephane@utulsa.edu, sandip@kolkata.mcs.utulsa.edu, gregoirerichard@yahoo.fr

Abstract

A consumer may be interested in buying a bundle of items, where any one item in the bundle may not be of particular interest. The emergence of online auctions allows such users to obtain bundles by bidding on different simultaneous or sequentially run auctions. Because the number of auctions and the number of combinations to form the bundles may be large, the bundle bidding problem becomes intractable and the user is likely to make sub-optimal decision given time constraints and information overload. We believe that an automated agent that can take user preferences and budgetary constraints and can strategically bid on behalf of a user can significantly enhance user profit and satisfaction. Our first step to build such an agent is to consider bundles containing many units of a single a item and auctions that sell only multiple units of one item type. We assume that users obtain goods over several days. Expectations of auctions and their outcome in the future allow the agent to bid strategically on currently open auctions. The agent decides how many items to bid for in the current auctions, and the maximum price to bid for each item. We evaluate our proposed strategy in different configurations: number of items sold, number of auctions opened, expected closing prices, etc. The agent produces greater returns in situations where future auctions can provide better profit, and where not too many agents use our proposed strategy.

Introduction

Auction is an appropriate mechanism to reach an economically efficient allocation of goods, services, resources, etc. An amazing variety and quantity of goods, services are traded everyday in online auctions. These trades occur between and among businesses and consumers. The growth of online auction market, in size and variety, provides new challenges for buyers. A buyer's goal is to obtain the best deal possible, and to achieve this, he must keep track of multiple auctions at many different sites. Even if the buyer follows only one item in multiple, simultaneous auctions, it might lead to information overload and associated sub-optimal purchase decisions. To confound the problem, a buyer may want to buy not just one, but a bundle of items. In this case, the buyer is interested only in obtaining all the items in the bundle and not any proper subset. On a larger scale, this situation also describes an industrial producer who needs raw material and services for production.

Different aspect of the problem of bidding in sequential and/or simultaneous auctions has been studied; increasing the performance of simultaneous auctions (Priest, Byde, & claudio Bartolini 2001), defining strategy when auctions do not use the same mechanism (Byde, Preist, & Jennings 2002), defining strategy to bid in sequential auctions using the history (Boutilier, Goldszmidt, & Sabata 1999; Tesauro & Bredin 2001). Combinatorial auctions, i.e., auctions that offer bundles of goods have received particular attention from researchers, and facilitate the purchase of bundles by users. The allocation of bundles to bidders so as to maximize total revenue for the sellers, however, is known to be an NP-complete problem. Various approximation schemes , as well as exact schemes for limited bid types have been investigated (H.M.Rothkopf, A.Pekec, & M.R.Harstad 1995; K.Fujishama, Y.Leyton-Brown, & Y.Shoham 1999; Parkes 1999). In general, however, online auction marketplaces, such as eBay, host simultaneous auctions that sell only one item at a time. Therefore, the buyers at such auctions have to put together bundles for themselves through buying the bundle constituents at different auctions.

The general *bundle bidding problem* involves bidding for multiple units of different items that constitute a bundle. The bidding problem for such individual buyers or producers become even more complex since the buyer needs to select the auctions, from the numerous auctions being held at various auction sites, to bid in to obtain a bundle, how much to bid and for what quantity in each of these auctions and how to factor in considerations of future auctions. Related considerations include possibilities of obtaining too many or too few items. Because of the dynamics of online auctions, time constraints may not allow the buyer to compute an optimal decision, and the buyer may have to accept sub-optimal results.

We believe the bundle bidding problem provides a key opportunity for applying intelligent agent technology. An agent can automate the task of bidding for bundles on behalf of the associate user, i.e., the agent can take preferences for bundles from the user and try to put together the bundle by bidding at multiple simultaneous or sequential auctions held at different online auction sites. The research goal is to develop a bundle bidding strategy that takes into account the user preferences or bundle valuations, budget constraints, etc. We believe the use of such an agent have great potential to enhance user's profit and satisfaction.

We assume that our agent has some expectation of the number of auctions selling a given item type in the near future. In addition, the agent has expectations of closing prices, or valuations of other bidders, of the items in those auctions. Based on previous auctions, one may design a probabilistic model for future auctions. We are interested in buyers who want to obtain a certain number of a given item type over a given time period. The buyers may not have an exact number of item to buy; rather, they have a utility function which evaluate their own value of obtaining a given number of units of the item of interest. We consider different English auctions that sell multiple units of this item type and that run simultaneously or sequentially over a period of time. Because of the number of auctions, the possible combinations to obtain a number of items may be large, and the bidding problem may quickly turns out to be intractable for a buyer. The use of agent technology offers a solution for the buyer. Expectations on the future auctions and their closing price allow an automated agent to bid strategically on behalf of a user. The agent chooses the auctions to place a bid, decides how many items to bid for in an auction and the corresponding bid price.

The focus of this work is to run simulations to verify our intuition. In particular, when assumptions are violated and the expectations are no longer correct, we show that the agent still behave well compared to others.

Simulation framework

For our simulation, we designed the following scenario: on each day, many auctions, each selling multiple units of a given item take place. We consider buyers who need to buy a number of items over several days.

Each auction is an English ascending auction selling a known number of units, all of the same type. The current ask price of every unit in each auction is known to all bidders, but the name of the buyer who has each bid is unknown. Each bidder has access to his current number of active bids.

The auctions taking place in the same day are simultaneous, they have the same opening time, and they terminate when no buyer place a bid in any of the auctions. To simulate this, the auction house picks randomly one buyer from the set of buyers. This buyer has the right to place one bid in each auction. To place a new bid, a buyer announces how many units he wants to obtain, and how much he bids. The bid is valid when for each unit, when the bid is a minimum increment above the the ask price for the auction. Invalid bids are rejected. When all his bids are placed or rejected, the auction house picks another buyer and gives him the hand.

Let us emphasize two details here. Each buyer can know how many bids are active in a given auction, and he has access to the current price. Thus, if there are more bids than he has at the same price, he cannot know whether his bids have been placed first or not. If he wants to place a bid to obtain more units in the same auction, he may place a new bid with only one unit and wait to see whether he has out-bid one of his bid. As the auction house only allow one bid per auction, this is not a good strategy. As a consequence, if a buyer wants to obtain more units from one auction, he may have to out-bid his own bids to ensure having the desired number of active bids. Also, as the auction house can keep track of who got the hand in the past, it can detect when all the participants have successively refused to place new bids. Then, the auction house conclude that everybody is satisfied with the current state and close all the auctions.

In the simulations, we are primarily interested in the buyers who want to buy a certain number of units over a given number of days. Each strategic buyer has a valuation function val, where $val(k)$ returns the value of obtaining k units for the buyer. The bidding strategy for strategic buyer will be discussed in the following section. We also use non-strategic, dummy buyers with specific valuation functions in each auction. Each dummy buyer bids on the auctions of a particular day. These dummy buyer's bid up to their valuations, where these valuations are drawn from a probability distribution, specific to a particular day. The same probability distributions are used by strategic buyer's to form expectations of closing prices of auctions on a particular day.

Bidding Strategy

Strategies

The goal of each buyer is to maximize the difference between the valuation function and the cost of obtaining the units. To achieve this goal, each buyer may use the expectation of the closing price of auctions in the future. We designed three different strategies corresponding to how far an agent looks ahead to the future:

- *oneday* is a strategy where no expectation on the future is made. It optimizes the benefit over the current day, if other units may be bought, this can be done another day.
- *twodays* is a strategy where the buyer looks one day ahead. The buyer using *twodays*, unlike one using *oneday* can optimize his benefit over two days.

- *threedays* strategy looks two days ahead, and optimize the benefit over three days.

Let assume a buyer has expectations over the next d days. This means that he has some expectations on the number of units for sale in the next d days, and for each unit, he has an estimated closing price. More precisely, a buyer who has expectation on the k^{th} day has a global picture of all the auctions of day k at closing time.

Moreover, a buyer has access to the current price of all units for sale in the current auctions. To bid strategically, we break the problem into two decisions: how much does it cost to buy k items now and how much does it cost to buy l items in the different cases in the future? Having answered these two problems, it is easy to find k and l such that the buyer optimizes its benefit over the d days, i.e., maximize the difference between the valuation of buying $k+l$ items and the cost to obtain these $k + l$ items.

Obtaining k more items at the same time during a day

- **Bidding for one more item in an auction:** Assume that there are n units for sell in a given auction. A buyer knows the current ask price for each unit, and knows how many active bids he has. δ denotes the increment to add to place a new bid, and $AP(i)$ denotes the i^{th} cheapest ask price of the auction. There can be three distinct cases when an agent tries to bid for one additional item;
  - if the buyer does not have any active bid, then the additional cost to get one more unit is the lowest ask price, $AP(1)$, plus the increment δ.
  - the buyer already has m active bids. As stated, since he cannot know the order of the bids, he must out-bid his own bids to make sure he will out-bid the $(m+1)^{th}$ cheapest bid. By doing so, he ensures to obtain an additional active bid. The additional cost to obtain this active bid is

$$AP(m+1) + \delta + \sum_{i=1}^{m} \left(AP\left(m+1\right) + \delta - AP\left(i\right)\right).$$

  - the buyer owns all the active bids, there is no more items he can get from that auction.
- **Bidding for one more item in the auctions of the day:** Since we can compute for each auction of the day the price of getting an additional item, we can find the auction to bid on that will minimize the cost to obtain the additional item. If the buyer already have all the active bids in all auctions on that day, he can not buy an extra unit that day.
- **Bidding for k more items at the same time during a day:** To find the cost of obtaining k additional items, one needs to repeat k times the process of calculating the cost of obtaining one additional item, where each iteration is performed with the updated auction states after simulating the placement of the last bids.

Obtaining l more items in the future Let us assume the buyer has maximum valuation of other buyer's for items in auctions to be held over the next few days. This translates into closing price expectations when the agent is not bidding in that auction. We assume that adding l bids above the expected closing price will enable the agent to win those l items in that auction, i.e., the adding of bid by this agent does not prompt other agent's to raise their bids in those auctions.

If one applies the same technique for the current day's auction for the future auctions, then a buyer has an estimate of the price he is likely to pay to obtain l more items in the future. But this cost is associated with two uncertainties: the buyer's expectations about other's valuations may be wrong and other agents may raise their bids in response to the buyer's bids above the predicted closing prices.

Experiments

The question we want to answer in the experiments is the following: if one believes the market will evolve in a certain way, what is the best strategy to bid now. It is obvious that if one expects the price to rise with time or remain the same, there is no question about considering to buy in the future, now is the best moment. Consequently, we focused on cases where one expects the prices to vary significantly over days.

Let us motivate our experiments with a realistic example. One illustration of the problem we are considering may be to bid on behalf of provider of a fine restaurant. Assume that the restaurant specializes in providing fresh fruits. Based on past weather, supply disruption news, etc. the owner may expect the quality or the quantity of the produces in the market to be high or low, the expected price varying consequently. We also consider that the restaurateur cannot look ahead too far: he may guess the state of the market for the next day, or for the two next days. His aim is to buy a certain quantity of fruit where a very small or a very large quantity is of less interest to him. This information is represented by the valuation function v of the provider. We used a valuation of this form:

$$v(n) = \frac{c}{1 + e^{\frac{n_1 - n}{n_2}}},$$

where c is the maximum amount of money that a buyer is ready to pay for obtaining goods, n_1 and n_2 controls how many units are wanted, and how tight this number is. For the following experiments, we have used $c = 1500$, $n_1 = 20$ and $n_2 = 15$.

The owner has contracted a supplier (a role which will be played by our agent) to buy the produces, etc. over the next D days (we use $D = 5$). He may attend a market where producers sell their produces in English ascending auctions. The auctions for a given day are held simultaneously. At the end of the D days, the supplier obtains n units, each unit i for a price of $c(i)$.

The goal of the supplier is to maximize the restaurant owner's utility (we will also refer the utility as his gain) i.e. $Gain = v(n) - \sum_{i=1}^{n} c(i)$.

In the setting of our experiments, we consider 5 auctions each day, each auction selling k units. We choose the valuations for the dummy buyers from Gaussian distributions that have different means on the different days. the means are chosen such that prices can go down significantly in the future. This variation allows strategic buyers to benefit by looking ahead to the future. The average valuations of the dummy buyers over different days are described in Figure 2.

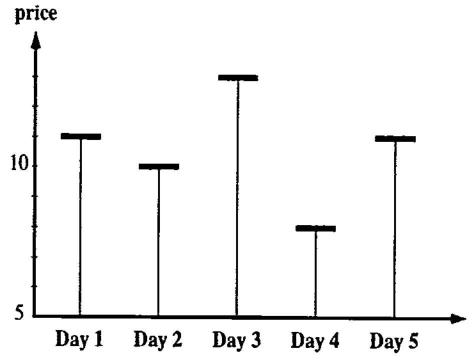

Figure 2: Setting for the dummy buyers.

units obtained by *twodays*

day	a1	a2	3	a4	a5	$\bar{p}$
1	0/1	1/1	7/7	2/7	7/7	12.20
2	2/7	10/10	1/10	1/1	1/1	10.18
3	0/3	0/2	0/5	0/1	0/1	13.03
4	0/5	0/7	0/5	0/5	0/15	8.70
5	0/7	0/7	0/7	0/1	0/1	11.10

units obtained by *threedays*

day	a1	a2	a3	a4	a5	$\bar{p}$
1	1/1	0/1	0/7	5/7	0/7	12.20
2	1/7	0/10	0/10	0/1	0/1	10.18
3	0/3	0/2	0/5	0/1	0/1	13.03
4	0/5	3/7	5/5	0/5	15/15	8.70
5	0/7	0/7	0/7	0/1	0/1	11.10

Table 1: Repartition of the units obtained, competition *twodays* vs *threedays*.

The goal of this setting is to encourage the spread of the purchase of units over several days. A buyer who is considering only the current or first day will try to buy as much as he can during this day. An agent capable of looking at one or two days ahead may find better to wait for future opportunities, or to buy some units now and buy more later. From Figure 2 we see that at the start of the second day, an agent who is looking one day ahead knows the prices will be high on the next day, and hence has the incentive to buy as many items as possible on the second day itself. The situation, however, is quite different for a buyer who looks two days ahead, as it can predict that the market will be comparatively cheaper on day 4, and hence it does not rush to complete purchases on day 2. From this example, we see having knowledge of the future can be of advantage to strategic bidders. Of course, the advantage of such strategic bidding is lost if the information is incorrect or if a significant percentage of bidders in the marketplace bids strategically.

Results

We run a simulation of the auctions over the five days with different participants. We monitor the "smart buyer" with different lookahead capabilities. For instance, the next example describes a market where two strategic agents were present, an agent using the *twodays* strategy, and an agent using the *threedays* strategy. The final utility of the *twodays* bidder is 674.62, and that for the *threedays* bidder is 707.46. The buyer who looks two days ahead performs better since during day 2, he foresees that day 4 will be cheap, so he waits for this day to have better deals. The buyer who considers only the next day has to compete to obtain as many units as he can during day 2 since he only sees that day 3 will be much more expensive than day 2. When day 3 comes, he has already bought enough units. The table 1 gives an example of the repartition of the units bought by each agent in each auction a_i at the end of the five days. The entries of the form p/q represent the p units bought by the agent over the q units for sell in that auction. The column $\bar{p}$ contains the mean of the price of one unit during the day.

If we repeat the same experiment with the same participants and the same expected prices, we are not likely to find exactly the same results. If the maximum price that a "smart buyer" is ready to pay is close to the maximum price of a dummy buyer, the one who obtains the good is the one who had the hand first. Having or not this item may change the strategy of the "smart buyer" in the future. Consequently, the following results are averages over a certain number of repetitions

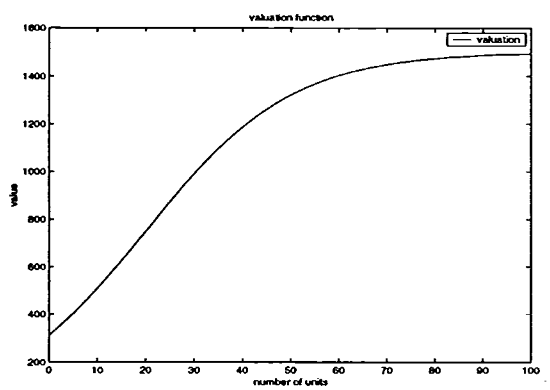

Figure 1: Valuation function used by "smart buyers".

of the setting. The standard deviation is around 30 for the gain, around 3 for the number of units obtained.

The valuation functions are identical for all strategic bidders and the price expectations are also the same, except that some look ahead farther into the future than others. We should obtain the best utility that a smart agent can reach.

In the first set of experiments, we study the performance of one "smart agent" competing against dummy buyers. In that case, if the strategic agent bids above the expected ask price, dummy buyers will not react. Thus, the gain obtained by the strategic agent is optimal.

	utility	# units
oneday vs dummies	642	32
twodays vs dummies	736.7	37.4
threedays vs dummies	803.5	39.3

Table 2: Experiment set 1: utility and number of units obtained.

A second set of experiments uses a small number of strategic agents in the auctions. The presence of other strategic agents violate the assumptions of other bidders' valuations and that the others are not going to respond to higher bid by one agent. In terms of strategy, the expected prices of the next days should be thought of as the minimum price of the auctions for the next days. The uncertainty lies on how much these prices will increase. A buyer looking two days ahead performs at least as well as a buyer who is looking only one day ahead. But the competition is likely to adversely affect the utility of all the buyers. However, since the expected closing price may be significantly smaller than the actual closing price, one agent may be forced to give up bidding and wait another day, which may allow him to take advantage of new opportunities in the future.

The table 3 presents the average gain of strategic bidders in different competitive environments. The percentage of loss for an agent using strategy *kdays* is the ratio between its gain and the gain obtained by a agent using *kdays* but competing only against dummy buyers (using the results contained in table 2). As expected, a buyer using the *threedays* strategy performs better than an agent using the *twodays* or *oneday* strategy. The utility is the lowest when two buyers using the same strategy compete against each other. They will behave the same way and increase competition during the same days. However, the loss of gain encountered when two agents using different strategy compete against each other may not be as important. In that particular setting, provided no other agent uses the same strategy, the gain of a buyer using the *threedays* strategy does not decrease significantly (no more than 2.4% of the gain obtained in the first set of experiments). When two agents using different strategies are competing, if one decides to give up bidding, the other may not take the

	gain		% of loss
1day vs *1day*	473.5		26.2
2days vs *2days*	640.2		13.1
3days vs *3days*	647.8		19.4
1day vs *2days*	avg	666.9	
	1day	618.5	3.7
	2days	715.1	3
1day vs *3days*	avg	698.3	
	1day	606.4	5.4
	3days	790.2	1.7
2days vs *3days*	avg	765	
	2days	731	0.7
	3days	799.5	0.5
1day vs *2days* vs *3days*	avg	680.5	
	1day	609.3	5.1
	2days	647.1	12.2
	3days	785	2.3

Table 3: Experiment set 2: utility and number of units obtained.

same decision, and thus it obtains the units. There are two main effects to this, since one agent gets some units, he may not compete intensively to obtain more units in the future. Also, this allows the other agent to take advantage of new opportunities later. Consequently, the increase of competition does not necessarily implies an important loss of gain for the buyers.

Future Work

The goal of the current work was to develop a bidding strategy that can utilize knowledge of valuations of other bidders in future auctions to enhance buyer utility when purchasing multiple units of an item from several auctions. We have developed a strategic bidding agent that uses a user valuation function for different quantities of an item, and the knowledge of valuations of other buyers in auctions to be held in the near future. In particular, we have studied strategic agents with knowledge of different time horizons. We demonstrate that strategic agents with longer lookahead perform better not only when competing against non-strategic agents, but also outperform agents with shorter lookahead. Moreover, an increase of the competition does not necessarily implies a significant decrease of the gain of buyers using different strategies.

In the current work, an agent uses the expected highest valuation of other bidders as the closing price in a future auction. We believe a more effective decision mechanism needs to use a probability distribution over possible closing prices or buyer valuations. We plan to investigate this modification to our work in the future. The use of entire probability distribution provide additional options in the decision process. There can be a tradeoff between choosing a strategy which is more likely to better results versus another strategy which has the potential for higher average gain. Thus, we can model risk adverse, risk neutral or risk seeking strategies by choosing the maximum likelihood, the maximum expected utility, or mixed strategies.

Another extension of this work may be the use of learning and modeling strategies to estimate the probability distribution or expected prices in future auctions.

References

Boutilier, C.; Goldszmidt, M.; and Sabata, B. 1999. Sequential auctions for the allocation of resources with complementarities. In *IJCAI*, 527–523.

Byde, A.; Preist, C.; and Jennings, N. R. 2002. Decision procedures for multiple auctions. In *1st Int Joint Conf. on Autonomous Agents and Multi-Agent Systems, Bolgna, Italy. (to appear).*

H.M.Rothkopf; A.Pekec; and M.R.Harstad. 1995. Computationally manageable combinatorial auctions. *DIMACS Technical Report* 95(9).

K.Fujishama; Y.Leyton-Brown; and Y.Shoham. 1999. Taming the computational complexity of combinatorial auctions: Optimal and approximate approaches. In *Proceedings of the Sixteenth International Conference on Artificial Intelligence*, 548–553.

Parkes, D. 1999. ibundle: An efficient ascending price bundle auction.

Priest, C.; Byde, A.; and claudio Bartolini. 2001. Economic dynamics of agents in multiple auctions. In *Proceedings of the 5th International Conference on autonomous Agents*, 545–551.

Tesauro, G. J., and Bredin, J. 2001. Strategic sequential bidding in auctions using dynamic programming. *IBM Research, technical report.*