An Adaptive Bidding Strategy in Multi-Round Combinatorial Auctions For Resource Allocation

Xin Sui, Ho-Fung Leung
Department of Computer Science and Engineering
The Chinese University of Hong Kong
Shatin, N.T., Hong Kong
{xsui, lhf}@cse.cuhk.edu.hk

Abstract

Combinatorial auctions, where bidders are allowed to put bids on bundles of items, are preferred to single-item auctions in the resource allocation problem because they allow bidders to express complementarities (substitutabilities) among items and therefore achieve better social efficiency. Although many works have been conducted on combinatorial auctions, most of them focus on the winner determination problem and the auction design. A large unexplored area of research in combinatorial auctions is the bidding strategies. In this paper, we propose a new adaptive bidding strategy in multi-round combinatorial auctions in static markets. The bidder adopting this strategy can adjust his profit margin constantly according to bidding histories to maximize his expected utility. Experiment results show that the adaptive bidding strategy performs fairly well when compared to the optimal fixed strategy in different market environments, even without any prior knowledge.

1 Introduction

With the increasing popularity of centralized and distributed computing technologies, the Internet has become a powerful computing platform where different users can use existing computing resources to perform their own tasks [6]. This kind of resource allocation problem, that is, how to allocate these resources among a group of users, is an important issue. Internet Auction is a good solution for this problem because it allocates resources to bidders who value them most and gets the efficient allocation from the view of economics [3].

Among all types of auctions, combinatorial auctions, where bidders are allowed to put bids on bundles of resources, receive much attention from researchers in both computer science and economics [5]. Combinatorial auctions can lead to more economical allocations of resources than traditional single-item auctions when bidders have complementarities (substitutabilities) among them. Such expressiveness can lead to an improvement of efficiency, which has been demonstrated in airport landing allocation and transportation exchanges [10][12].

In recent ten years, most works on combinatorial auctions concentrate on the winner determination problem and the auction design. Winner determination problem is to compute the optimal allocation of resources among bidders and is proved to be NP-hard [11]. Many works have been conducted to solve this problem, including finding optimal solutions and approximate solutions [13][7]. Auction design involves the designing of different protocols for combinatorial auctions, such as single-round versus multi-rounds and sealed-bid versus open-bid [9][4].

Another area of research on combinatorial auctions is the designing of bidding strategies. As combinatorial auctions are incorporated with the first-price sealed-bid auction protocol in many applications [4], we are especially interested in bidding strategies in this kind of auctions. However, unlike the Vickrey Auction, where bidding truthfully is the dominant strategy, no dominant strategy exists in first-price sealed-bid combinatorial auctions [8]. In a market composed of self-interested bidders, the general unwillingness of information sharing makes this decision problem complex. For example, a bidder is not able to get the information beforehand that will benefit himself when participating, such as bidding bundles and prices of others' and the number of bidders competing for resources in the market.

An intuitive solution to this decision problem is that in different market environments, it is better for a bidder to use different strategies. For example, when the demands are far greater than supplies, which means that bidders compete severely for resources, a bidder should not bid too aggressively, that is, to use a low profit margin to get a high winning probability. On the contrary, when demands are less than supplies, a bidder should use a relative high profit margin to gain a high winning utility.

In this paper, we consider a scenario where first-price sealed-bid combinatorial auctions are employed to distribute computational resources among a group of users. Based on the intuition described in the above paragraph, we propose a new adaptive bidding strategy in multi-round combinatorial auctions. The profit margin used by the bidder adopting this kind of strategy can be adjusted constantly according to bidding histories, and finally approaches to the optimal profit margin in the current market environment even without the bidder's prior knowledge. Experiment results show that the bidder using the adaptive strategy outperforms bidders using other strategies and achieves good utilities compared to the optimal strategy in different environments.

This paper is structured as follows. Section 2 presents related work on bidding strategies in combinatorial auctions. Section 3 presents some preliminaries. Section 4 describes the proposed adaptive bidding strategy. Experiment results are discussed in Section 5. Finally Section 6 concludes this paper and highlights some possible future work.

2 Related Work

In first-price sealed-bid combinatorial auctions, a bidder has exponential number of bundles to bid on. The problem of deciding which bundles to choose and how much to bid for them are referred to as the bundle strategy and the price strategy respectively. In the following, we present a survey on both strategies in combinatorial auctions.

The work of Berhault et al. [2] focus on the bundle strategy in single-round combinatorial auctions. In their work, combinatorial auctions are employed to allocate unexplored terrains to robots distributed in a large field. Four bidding strategies are proposed for robots: Three-Combination, Smart-Combination, Greedy and Graph-Cut. Through experiments they show that combinatorial auctions achieve better efficiencies than single-item auctions and generate good results compared to optimal centralized allocations. They also show the influences of bundle strategies on team performances, where the Graph-Cut strategy clearly outperforms the other three.

An et al. [1] also study the bundle strategy in single-round combinatorial auctions. They propose two bundle strategies: Internal-Based and Competition-Based. Bidders using the former strategy only bid on bundles for which they have higher valuations, while bidders using the latter strategy only bid on bundles for which they have higher ratios of valuations to their competitors' according to their prior estimations. Simulation results show that wise bidders using these two strategies outperform naive bidders, who only submit single-item bids.

Schwind et al. [14] attempt to solve the computational resource allocation problem using multi-round combinatorial auctions. They study the situation where bidders use virtual currencies, which are obtained by selling unused resources, to get accesses to computational resources needed for accomplishing their own tasks. They propose price strategies for two kinds of bidders: impatient bidders and quantity maximizing bidders. Experiment results show that for the first kind of bidders, it is better to bid high prices to get fast accesses to resources, while the second kind of bidders had better bid low prices and keep on waiting for resources.

Although these existing works addressed price strategies in combinatorial auctions, their results focus on posterior analysis and the proposed price strategies are not adaptive. To the best of our knowledge, this paper is the first attempt to design an adaptive bidding strategy for combinatorial auctions, which is the main contribution of this work.

3 Preliminaries

3.1 Combinatorial Auction Mechanism For Resource Allocation

The combinatorial auction mechanism for the resource allocation problem can be described as follows. Multiple users (bidders) need some types of resources to perform their own tasks. There is a resource manager (auctioneer) who controls all resources. Suppose there are N users and M different types of resources. For each type of resource j, there is a capacity c_j denoting the total number of units that are available. There is a demand constraint $D = (d_1, d_2, ..., d_m)$, where d_j is the maximum number of units of resource j that each bidder can request for. Each user i submits a sealed-bid $b_i = (S, p_i)$, where $S = (s_1, s_2, ..., s_m)$ denotes a resource bundle, with s_j being the number of units of resource j user i needs, $0 \leq s_j \leq d_j, \forall j \in M$, and p_i is a positive number which denotes the price user i will pay for getting S. After receiving bids from all users, the resource manager solves the winner determination problem, that is to find a feasible allocation which maximizes the auctioneer's revenue. A feasible allocation means that for each type of resource, the total number of units allocated cannot exceed the capacity of that resource. Each winning user i pays p_i, gets accesses to the resources in the bundle he bids for, performs his own task, and then returns them to the resource manager. We refer to the process from the beginning of bid submission to the end of resource return as a *round* of combinatorial auctions. Because resources are reusable, the combinatorial auction can be repeated for multiple rounds before it is closed by the resource manager.

We list some assumptions used in this paper. First, we assume that the combinatorial auction market is static. A combinatorial auction market is said to be static if ratios of supplies and demands of different resources are kept constant in the whole process of the auction before it finishes. Second, each user only submits one bid in a round, denoting the number of different resources needed for the current task. This corresponds to the situation that each bidder has time sequences on tasks that the task which arrives earlier must be accomplished before later tasks are executed. Finally, in every round of combinatorial auctions, winners of the previous round submit new bids, while losers continue to submit the lost bids. A bidder will drop a bid if it has been lost for consecutive τ rounds and then submit a new bid for the next round, which means that the bidder will not keep on waiting for a resource bundle forever.

3.2 Valuation Model

The value of a resource bundle S, denoted as $v_i(s)$, comprises two parts: the sum of individual values of resources in the bundle and the synergy value among resources in the bundle, which are denoted as $v_{indi}(S)$ and $v_{syn}(S)$ respectively. The synergy value can be positive or negative. Resources in a bundle are said to have positive (negative) synergy, if their combined value for the bidder is larger (smaller) than the sum of their individual values. Suppose there is a bidder i who values a bundle containing 2 units of resource R1 and 3 units of resource R2 at \$19, but only \$4 and \$3 for each unit of R1 and R2. In this case, the combined value of these resources is greater the sum of individual values, which is 18 only, and there is a synergy value of positive 1 for this resource bundle.

Each bidder i has a private value v_{ij} for a single unit of resource j, $\forall j \in M$. For a resource bundle S, $v_{indi}(S)$ is defined as the sum of valuations of all types of resources in the bundle and $v_{syn}(S)$ is defined as the product of synergy seed and $v_{indi}(S)$. The synergy seed is a function of S $syn_{seed} : A_1 \times A_2 \times ... \times A_m \to [\mu, \nu]$, where A_j is integer of 0 or 1, that A_j=1 if and only if resource j is requested in bundle S, $\forall j \in M$. The values of μ and ν are the upper and lower bounds for the synergy seed.

According to this valuation model, the value of a resource bundle S is hence:

$$\begin{aligned} v_i(S) &= v_{indi}(S) + v_{syn}(S) \\ &= v_{indi}(S) + v_{indi}(S) \times syn_{seed}(S) \\ &= v_{indi}(S) \times (1 + syn_{seed}(S)) \\ &= \sum_{1 \le j \le m} v_{ij} s_j \times (1 + syn_{seed}(S)) \end{aligned} \tag{1}$$

From equation (1), we can see that the value of a bundle is $1 + \mu$ times the sum of individual values at least and $1 + \nu$ times the sum of individual values at most. In addition, if we refer to $|S|$ as the number of different types of resources in bundle S, we assume $syn_{seed}(S) = 0$ if $|S| = 1$, which means that there is no synergy value in bundle S when only one type of resource is contained. This valuation model conforms to our common sense that bundles containing same types of resources values more as the numbers of units for resources increase, e.g., a bundle with 2 units of resource R1 and 3 units of resource R2 values more than a bundle with 1 unit of resource R1 and 2 units of resource R2.

4 Adaptive Bidding Strategy

As mentioned above, each bidder will pay the price he bids for a resource bundle if he is a winner. The utility (profit) for bidder i of winning bundle S is computed as:

$$u_i(S) = v_i(S) - p_i(S) \tag{2}$$

where $u_i(S)$, $v_i(S)$ and $p_i(S)$ are his winning utility, valuation and price for bundle S respectively.

When bidding for a resource bundle, a rational bidder will use a bid value which is greater than zero and less than his valuation of that bundle. That is, suppose the valuation of bidder i for bundle S is $v_i(S)$, then his bidding price is $(1-pm)\times v_i(S)$, where $0 < pm < 1$. Here, we refer to the value of pm as profit margin. Note that pm is greater than zero to ensure that the bidder will have a positive utility and pm is less than one to make the bidding price positive. Combined with (2), the utility of bidder i is hence:

$$\begin{aligned} u_i(S) &= v_i(S) - p_i(S) \\ &= v_i(S) - (1 - pm) \times v_i(S) \\ &= pm \times v_i(S) \end{aligned} \tag{3}$$

A bidder faces a dilemma in deciding what profit margin to use in combinatorial auctions: bidding with a low profit margin will necessarily lead to a high chance of winning, but will only result in a low utility if he wins; on the other hand, bidding with a high profit margin will of course generate a high utility if he wins, but will decrease his winning probability. If a bidder can get some prior knowledge about the market, he can probably make the decision benefitting himself with the information. For example, consider a market where supplies are greater than demands. A bidder who has prior knowledge about the market will use a lower profit margin when bidding. This is because in this kind of market, bidders face little competition from others, and bidding with a low profit margin will lead to a winning utility while the winning probability is virtually unaffected. However, having prior knowledge is not always possible in a real market, because each bidder is self-interested and generally unwilling to share information with others who are also competing for resources. As a result, we make an assumption in this paper that the available information for each bidder is his bidding information only, that is, the information about his own bids in previous rounds.

4.1 Basic Concepts

In order to explain the adaptive bidding strategy, we first introduce some basic concepts.

Definition 1 A *reference record* of a bid b for bidder i is a tuple $rr_b = (S, v_i(S), pm_b, lose_b, win_b)$, where S is the requested bundle in b, $v_i(S)$ is the bidder's valuation of this bidder, pm_b is the profit margin for bid b, $lose_b$ is the number of rounds the bidder keeps on bidding before bid b is won or dropped and win_b is a integer of 0 or 1 denoting whether this bid is won or dropped. If win_b=1, then this bid is won, else it is dropped.

In the definition, the minimum value of $wait_b$ is 0, when the bidder wins the requested resource bundle at the first round after he submits it, and the maximum value of $wait_b$ is τ, when he keeps on losing all the time and finally drops the bid. Note that win_b=0 if and only if $wait_b$ equals to τ.

Definition 2 A *bidding history* of a bidder, denoted as bh, is the sequence of recent κ reference records.

For different bids in the bidding history, a bidder can use different profit margins, e.g., for two reference records rr_{b_m} and rr_{b_n} in a bidding history, their profit margins pm_{b_m} and pm_{b_n} can be different.

Definition 3 A *consistent bidding history* of a bidder, denoted as cbh, is a bidding history in which all reference records share the same profit margin.

Every time when a bid is won or dropped, a new bidding history is formed. However, we say that it is consistent only when the all reference records share the same profit margin. Consider a bidder who uses a fixed profit margin for all reference records, then all his bidding histories are consistent. On the contrary, if a bidder never uses the same profit margin for two consecutive reference records, none of his bidding history is consistent.

Definition 4 The *expected utility function* of bidder i on a consistent bidding history *cbh, denoted as* $u_{ex}(cbh)$*,* is defined as:

$$u_{ex}(cbh) = pm_{cbh} \times \frac{\sum_{rr_b \in cbh} win_b}{\sum_{rr_b \in cbh}(win_b + wait_b)} \tag{4}$$

where pm_{cbh} is the common profit margin used in this consistent bidding history, and $wait_b$ and win_b are the same as in the definition of reference record.

The reason why we refer to this function as expected utility function is explained as follows. When the value of κ is infinitely large, the second factor in the function is the winning chance of the bidder if he uses the profit margin of pm_{cbh} to bid in the auction. If we multiply it by the profit margin, the product is thus the scaled utility of the bidder in the consistent bidding history.

4.2 Adaptive Strategy

Based on the basic concepts defined above, we describe the adaptive strategy. The general idea is that every time when a new consistent bidding history is formed, the profit margin used by the bidder is increased or decreased according to the bidder's $1st$ and $2nd$ most recent consistent bidding histories. The new profit margin is used by the bidder when he bids in subsequent rounds until the next consistent bidding history is formed. This process is referred to as an adaptation of the profit margin. Through adaptations, the profit margin used will converge to the optimal profit margin. Here, the optimal profit margin is the profit margin that will maximize the expected utility when κ is infinitely large.

We refer to the increase and decrease of the profit margin as a positive and negative adjustment respectively, and use a 0 or 1 variable δ to indicate the previous adjustment of the profit margin: if $\delta = 1$, then the previous adjustment is positive, otherwise negative. We use u and u' to denote the expected utilities of the $1st$ and $2nd$ most recent consistent bidding histories. We also use pm to denote the current profit margin, and use pm' to denote the profit margin before the previous adjustment.

The adaptive strategy is illustrated in Algorithm 1.

We first give a general view of two functions: DecreaseStep (line 7) and ProfirMarginReset (line 19) before the adaptive strategy is illustrated. In function DecreaseStep, $step$ is decreased under some conditions, and in function ProfitMarginReset, pm is reset to a value according to recent consistent bidding histories when certain conditions hold. We will introduce them later in the paper.

The adaptive strategy can be explained as follows. At the beginning, pm, $step$, δ and u' are initialized. During the process of the auction, the value of pm is used by the bidder to bid in the auction, and is changed every time when a new consistent bidding history is formed. On deciding how to change this value, the bidder first computes the expected utility of the newly formed consistent bidding history, which is denoted by $u_{ex}(cbh)$, record this value in u, and then changes the profit margin according to the value of δ and the relationship between u and u' (line 8-12). If the previous adjustment of the profit margin, which is recorded by δ, leads to a decrease of the expected utility, an opposite adjustment will be made (line 8-9), otherwise, a same adjustment will be made (line 10-11). An opposite adjustment

Algorithm 1 Adaptive strategy

$pm = \eta$, $step = \theta$, $\delta = 1$ and $u' = 0$.
while auction does not finish **do**
  Use profit margin of pm to bid for the current round
  if a new consistent bidding history cbh is formed and $step > \epsilon$ **then**
    Compute $u_{ex}(cbh)$.
    $u = u_{ex}(cbh)$ and $pm' = pm$.
    DecreaseStep ();
    if $u < u'$ **then**
      $pm = pm - \delta \times step$
    else if $u \geq u'$ **then**
      $pm = pm + \delta \times step$
    end if
    if $pm > pm'$ **then**
      $\delta = 1$
    else if $pm < pm'$ **then**
      $\delta = -1$
    end if
    $u' = u$
    ProfitMarginReset ();
  end if
end while

means that the previous and next adjustment of the profit margin are different, e.g., one is positive and the other is negative, and a same adjustment means that both the previous and next adjustment are positive or negative. The value of pm is changed from time to time in the auction and will gradually converge because of the decrease of $step$. The adaptation is stopped when $step$ is smaller than a threshold ϵ and the bidder use the profit margin at that time for all subsequent rounds until the auction finishes.

Next, we will describe the two functions in the adaptive strategy in detail: DecreaseStep and ProfitMarginReset.

4.2.1 DecreaseStep

As mentioned above, the second factor in the expected utility function is the winning probability if κ is infinitely large. From Algorithm 1, we can see that the bidder using this strategy will change his profit margin every κ reference records. In this case, κ cannot be set to a large value because the bidder also needs to adapt to the environment in a timely manner. By such constraint, the second factor can be only regarded as an approximation of the winning probability. If we refer to $p_{win}(pm)$ as the winning probability of the bidder who bids with the profit margin of pm, then the larger the value of κ is, generally the closer to $p_{win}(pm)$ the second factor is. Thereby, we say that the second factor is vulnerable to noises and the expected utility is inaccurate when κ is small.

We need an algorithm that is robust against noises to decide when to decrease $step$. Before it is introduced, the concept of profit margin history is defined as the notion of goes towards is given.

Definition 5 A *profit margin history*, which is denoted as pmh, is a sequence of λ real numbers, in which the kth element, pmh^k, is the profit margin used for the kth most recent consistent bidding history.

We also give the notion of "*going towards*" as follows. The profit margin pm is said to *go towards* a value π if 1) $pm<\pi$ and the next adjustment for pm is positive or 2) $pm>\pi$ and the next adjustment for pm is negative.

The function of DecreaseStep is given in Algorithm 2. On whether or not to decrease $step$, the bidder first computes $mean$, the mean value of elements in pmh. For each pmh^k, there is a variable ω^k that equals to 1 if $|pmh^k - mean| \leq step$, otherwise 0 (line 2-8). Then $step$ is decreased by γ if all the three conditions are satisfied.

Algorithm 2 Function: DecreaseStep

1: Compute $mean$ = $\frac{1}{\lambda}\sum_{k=1}^{\lambda} pmh^k$.
2: **for** $k = 0$ to λ **do**
3:   **if** $|pmh^k - mean| \leq step$ **then**
4:     ω^k = 1
5:   **else**
6:     ω^k = 0
7:   **end if**
8: **end for**
9: **if** $\sum_{k=1}^{\lambda} \omega^k \geq \phi$
  and $\omega^1 = 1$
  and pm goes towards $mean$ **then**
10:   Decrease $step$ by γ
11: **end if**

We illustrate these three conditions as follows. When the first condition that $\sum_{k=1}^{\lambda} \omega^k \geq \phi$ is satisfied, it means that the profit margins in pmh fluctuate around $mean$. In this case, the value of $mean$ is regarded as an approximation of the optimal profit margin. The second condition that $\omega^1 = 1$ and the third condition that pm goes towards $mean$ guarantee that the optimal profit margin can be further approached in the next adaptation if $step$ is decreased.

4.2.2 ProfitMarginReset

Just as its name implies, this function resets the value of pm if it deviates too much from the optimal profit margin. Here, pm is said to *deviate from* a value π if 1) $pm < \pi$ and the next adjustment for pm is negative or 2) $pm > \pi$ and the next adjustment for pm is positive. Because the optimal profit margin is not known in advance, we use the same way as in Algorithm 2 to get its approximation.

The reason for the deviation can be explained as follows. From Algorithm 1, we can see that given δ, the next adjustment of the profit margin is determined by the relationship between u and u'. However, when $step$ is small, this relationship mainly depends on the relationship between the second factors in the expected utility function, which are vulnerable to noises as mentioned above.

However, the profit margin should not be reset all the time. For example, if the profit margin has been consecutively reset from higher or lower values for a number of times, then it is quite possible that the value, which the profit margin is reset to, is not close to the optimal profit margin. This corresponds to our common sense that if we always get wrong answers to a problem with a solution, we tend to believe that the solution itself may be improper and needs to be corrected.

Before we describe the function of ProfitMarginReset condition, the concept of *profit margin history center* is given as follows.

Definition 6 The *profit margin history center* of a profit margin history is the mean value of elements in pmh if the condition that $\sum_{k=1}^{\lambda} \omega^k \geq \phi$ in line 9 of Algorithm 2 is satisfied, otherwise, it does not exist.

We use cen' to denote the most recent profit margin history center, and use res_l and res_h to denote the number of times that the profit margin is consecutively reset from lower and higher values respectively.

The function of ProfitMarginReset is described in Algorithm 3.

Algorithm 3 Function: ProfitMarginReset

1: Compute d = pm - cen'.
2: **if** $|d| > \psi \times step$ **then**
3:     **if** $d>0$ **then**
4:         Set res_l = 0 and res_h = res_h + 1.
5:         **if** $res_h < \chi$ **then**
6:             Set pm = cen'
7:         **else**
8:             Set cen' = pm and res_h = 0.
9:         **end if**
10:     **else**
11:         Set res_h = 0 and res_l = res_l + 1
12:         **if** $res_l < \chi$ **then**
13:             Set pm = cen'
14:         **else**
15:             Set cen' = pm and res_l = 0.
16:         **end if**
17:     **end if**
18: **end if**

The algorithm can be explained as follows. First, the bidder computes the difference between pm and cen'. If the absolute value of this difference is more than $\psi \times step$, it is regarded that the profit margin has deviated too much from the optimal one. Values of res_l and res_h are updated according to the value of d (lines 4 and 11). If the upper bound of consecutively reset time χ has not been reached, the profit margin is reset to cen', the most recent approximation of the optimal profit margin (lines 6 and 13). Otherwise, it is regarded that this approximation is inappropriate, in which case it is replaced by pm (lines 8 and 15).

5 Experiment Evaluation

To evaluate the performance of the adaptive strategy, two sets of experiments are conducted. In the first set of experiments, the performances of different fixed strategies are compared in different markets. A fix strategy is a strategy that a same profit margin is used by the bidder for all reference records. In the second set of experiments, the performances of the random strategy (RS), the adaptive strategy (AS) and the best fixed strategy (BFS) are compared. Random strategy is a strategy that a random profit margin is used for each reference record. Best fixed strategy is the fixed strategy that generates the highest utility among all fixed strategies used in the first set of experiments. We refer to the best fixed profit margin as the the profit margin used by the best fixed strategy. In addition, we also show the typical adaptation process of the profit margin in a single run in different markets.

5.1 Experiment Setup

In our experiments, each combinatorial auction is repeated for 500 rounds and an iteration of 500 rounds is referred to as a *run*. Motivated by other works [1][14], in each run, we have one test bidder using strategy X and others bidding their true valuations. Here, X can be the adaptive strategy, the random strategy or any fix strategy. The performances of different strategies are compared through accumulated utilities of the test bidder in a static market in 100 runs.

Settings of these experiments are as follows. Four types of resources with capacities of $60, 40, 40, 20$ respectively are provided by the resource manager. Numbers of units that a bidder requests for different types of resources are integers randomly drawn from uniform distributions [0, 3], [0, 2], [0, 2] and [0, 1] respectively. That is to say, the demand constraint is $D = (3, 2, 2, 1)$. At the beginning of each run, each bidder initializes his valuations for all bundles: his valuations for a single unit of different types of resource are real numbers randomly drawn from uniform distributions

[3, 6], [4, 8], [4, 8] and [6, 10] respectively, and his synergy seeds for different bundles are real numbers randomly drawn from a uniform distribution [-0.2, 0.2].

We use the ratio of total supplies and demands to denote a market type and a market is said to be a $1{:}n$ market if such ratio equals to $1{:}n$. In our experiments, because the total supplies are fixed, we use different value of N to denote different market types. For example, when N=40, the expected total demands are 60, 40, 40 and 20, which equal to the total supplies, and we say that this is a 1:1 market. We use four values of N=30, 40, 50 and 60 to denote the 1:0.75, 1:1, 1:1.25 and 1:1.5 market respectively.

Parameters used in experiments are showed in Figure 1.

Parameter	Value	Description
τ	3	maximum lost round
κ	5	length of a bidding history
η	0.05	initial value of pm
θ	0.1	initial value of $step$
ϵ	0.01	threshold for $step$ to stop adaptation
λ	10	length of a profit margin history
ϕ	6	see Algorithm 2 and Definition 6
γ	2	degree of decrease for $step$
ψ	3	see Algorithm 3
χ	3	maximum consecutive reset time

Figure 1. Parameters used in experiments

5.2 Experiment Results and Analysis

Figure 2 shows the results of the first set of experiments. In Figure 2, each curve represents a different market type, and for each market type, the accumulated utilities of the test bidder using 10 different fixed strategies with profit margins of pm_i, $i = 1, 2, ...10$, where $pm_i = (i-1) \times 0.1 + 0.05$ are compared.

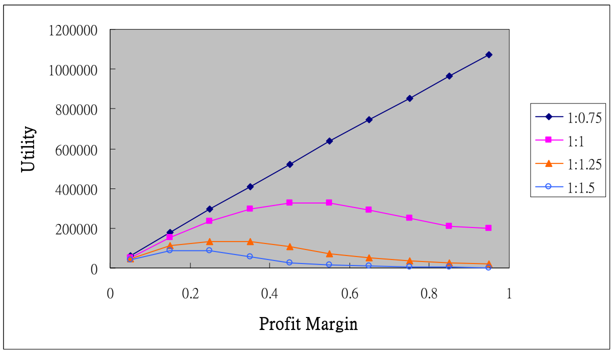

Figure 2. Utilities of the test bidder using 10 different fixed strategies in different markets

From Figure 2, we can see that the more competitive the market is, the smaller the value of the best fixed profit margin is. For example, in the $1 : 0.75$ market, where bidders face few competitions from others, the best fixed profit margin is 0.95, and in the $1 : 1.5$ market, where bidders face fierce competitions from others, the best fixed profit margin is 0.25. This corresponds to our common sense that it is better for a bidder to use different profit margins in different markets: in a market that is short of competition, it is better for a bidder to use a high profit margin to obtain a high utility, while in a market that is highly competitive, it is better for a bidder to use a low profit margin to beat others.

Figure 3 shows the results of the second set of experiments. Here, RS corresponds to the random strategy, AS corresponds to the adaptive strategy, and BFS corresponds to the optimal fixed strategy.

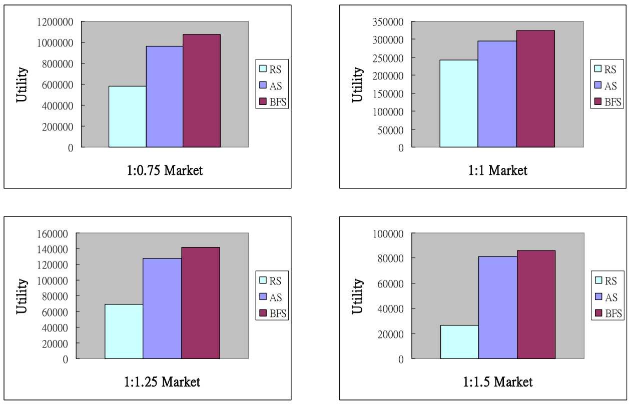

Figure 3. Utilities achieved by the test bidder using strategies of RS, AS and BFS

From Figure 3, we can see that the adaptive strategy performs fairly well when compared to the best fixed strategy, and outperforms the random strategy much in different market environments. As described above, the best fixed strategy is the strategy that performs best among all fixed strategies. The bidder using the best fixed strategy should be regarded as having prior knowledge about the market environment and is able to use the best fixed profit margin to obtain a high utility. On the contrary, the bidder using the random strategy can be regarded as not having any prior knowledge about the market and will use a random profit margin for every reference record. Therefore, it is impressive that the bidder using the adaptive strategy, can still obtain utilities that is about 90% compared to the utilities obtained by the bidder using the best fixed strategy in different market environments. As the bidder using the adaptive strategy does not need to know the market type in advance, we can draw the conclusion that the performance of the adaptive strategy is good, even without any prior knowledge about the market.

In addition, we also show the typical adaptation processes of the profit margin in a single run in different markets. For each type of market, the horizontal line represents the optimal fixed profit margin in that type of market.

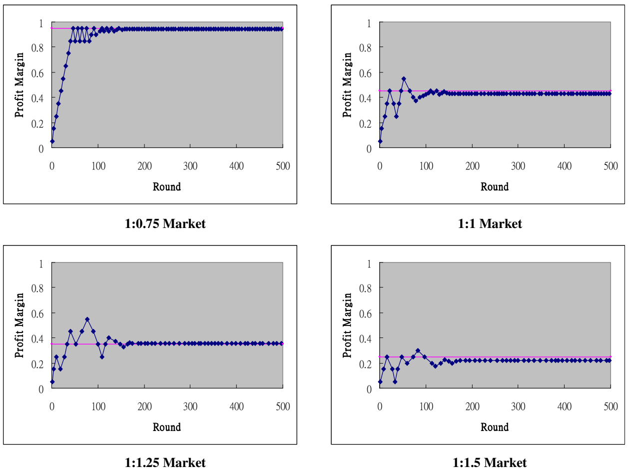

Figure 4. The typical adaptation process of the profit margin in a single run in different markets

From Figure 4, we can see that for each type of market, the profit margin in the adaptive strategy has converged to a values, which is very close to the best fixed profit margin in that market type. This means that the adaptive strategy is capable of adapting in different markets. In addition, the convergence speed is fast: for each market type, the profit margin has converged at about the $150th$ round and a value that is close to the optimal fixed profit margin is found. This value is used by the bidder to bid in subsequent rounds, which guarantees that the adaptive strategy can generate a very good utility when compared to both the optimal fixed strategy and the random strategy.

6 Conclusion and Future Work

In this paper, we propose a new adaptive bidding strategy in multi-round combinatorial auctions for the resource allocation problem in static markets. The bidder adopting this strategy can adjust his profit margin constantly according to bidding histories and finally adapts to the market environment. Experiment results show that 1) the adaptive strategy performs fairly well compared to the optimal fixed strategy and the random strategy in different market environments. 2) the bidder using the adaptive strategy can still obtain a very good utility, even without any prior knowledge about the market. 3) the adaptive strategy is capable of adapting to different markets and the convergence speed is fast.

In the future, we intend to design an adaptive strategy for dynamic markets, where the ratios of supplies and demands change from time to time. In this type of market, an intelligent bidder should perceive and respond to the environment in a more timely manner, which makes the designing of the adaptive strategy more complex. In addition, from the results we can see that although the adaptive strategy performs fairly well, there is still room for improvement. We also intend to explore the influences of different parameters on the performance of the adaptive strategy.

References

[1] N. An, W. Elmaghraby, and P. Keskinocak. Bidding strategies and their impact on revenues in combinatorial auctions. *Journal of Revenue and Pricing Management*, 3(4):337–357, 2005.

[2] M. Berhault, H. Huang, P. Keskinocak, S. Koenig, W. Elmaghraby, P. Griffin, and A. Kleywegt. Robot exploration with combinatorial auctions. *Intelligent Robots and Systems, 2003.(IROS 2003). Proceedings. 2003 IEEE/RSJ International Conference on*, 2, 2003.

[3] S. Clearwater. *Market-Based Control: A Paradigm for Distributed Resource Allocation*. World Scientific, 1996.

[4] P. Cramton, Y. Shoham, and R. Steinberg. *Combinatorial Auctions*. MIT Press, Cambridge, Massachusetts, 2006.

[5] S. de Vries and R. Vohra. Combinatorial Auctions: A Survey. *INFORMS Journal On Computing*, 15(3):284–309, 2003.

[6] I. Foster. The Anatomy of the Grid: Enabling Scalable Virtual Organizations. *Euro-Par 2001 Parallel Processing: 7th International Euro-Par Conference, Manchester, UK, August 28-31, 2001: Proceedings*, 2001.

[7] Y. Fujishima, K. Leyton-Brown, and Y. Shoham. Taming the computational complexity of combinatorial auctions: Optimal and approximate approaches. *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 548–553, 1999.

[8] V. Krishna. *Auction Theory*. Academic Press, 2002.

[9] D. Parkes and L. Ungar. Iterative combinatorial auctions: Theory and practice. *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 74–81, 2000.

[10] S. Rassenti, V. Smith, and R. Bulfin. A combinatorial auction mechanism for airport time slot allocation. *Bell Journal of Economics*, 13(2):402–417, 1982.

[11] M. Rothkopf, A. Pekec, and R. Harstad. Computationally Manageable Combinational Auctions. *Management Science*, 44(8):1131–1147, 1998.

[12] T. Sandholm. An implementation of the contract net protocol based on marginal cost calculations. *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 256–262, 1993.

[13] T. Sandholm, S. Suri, A. Gilpin, and D. Levine. CABOB: A Fast Optimal Algorithm for Winner Determination in Combinatorial Auctions. *Management Science*, 51(3):374–390, 2005.

[14] M. Schwind, T. Stockheim, and O. Gujo. Agents' bidding strategies in a combinatorial auction controlled grid environment. In *TADA/AMEC*, pages 149–163, 2006.