Auction Equilibrium Strategies for Task Allocation in Uncertain Environments

David Sarne[1], Meirav Hadad[2], and Sarit Kraus[1,3]

[1] Department of Computer Science, Bar-Ilan University, Ramat-Gan, 52900 Israel
[2] Caesarea Rothschild Institute University of Haifa Mount Carmel 31905 Haifa, Israel
[3] Institute for Advanced Computer Studies, University of Maryland, College Park, MD 20742

Abstract. In this paper we address a model of self interested information agents competing to perform tasks. The agents are situated in an uncertain environment while different tasks dynamically arrive from a central manager. The agents differ in their capabilities to perform a task under different world states. Previous models concerning cooperative agents aiming for a joint goal are not applicable in such environments, since self interested agents have a motivation to deviate from the joint allocation strategy, in order to increase their own benefits. Given the allocation protocol set by the central manager, a stable solution, is a set of strategies, derived from an equilibrium where no agent can benefit from changing its strategy given the other agents' strategies. Specifically we focus on a protocol in which, upon arrival of a new task, the central manager starts a reverse auction among the agents, and the agent who bids the lowest cost wins. We introduce the model, formulate its equations and suggest equilibrium strategies for the agents. By identifying specific characteristics of the equilibria, we manage to suggest an efficient algorithm for enhancing the agents' calculation of the equilibrium strategies. A comparison with the central allocation mechanism, and the effect of environmental settings on the perceived equilibrium are given using several sample environments.

1 Introduction

Distributed task allocation to self-interested agents is an important concept in Multi-Agent System (MAS) environments [10]. Though the case of task allocation to cooperative agents with a joint goal is always preferred over the allocation to self-interested agents, the latter method is significantly important in environments where a central manager does not own the agents or cannot fully control them. In such environments, the central planner, striving to achieve a specific goal, will try to enforce cooperation throughout task allocation protocols. Nevertheless, while both centralized and distributed cooperative allocation mechanisms assume a joint goal for all agents, this is not necessarily the case once self-interested agents are introduced into a model. The latter case may suggest conflict of interests, as each agent strives to maximize its own utility. This might result in a deviation of the self interested agents from the centralized proposed mechanism. Thus, having a clear methodology for finding the allocation strategies in equilibrium is a necessary condition for resolving the non-cooperative problem.

In this paper we study equilibrium strategies for self-interested agents situated in uncertain environments. We suggest a model in which different types of tasks arrive dynamically according to a given probability (see Figure 1). The tasks are allocated according to a pre-defined protocol set by a central manager. This central manager may be defined as a government, a municipality, a company, etc., operating in a dynamic environment and lacking the required resources to perform the tasks by itself. The protocol defines the rules for selecting a performer of the task and the appropriate payment for this task. A capability for performing a task depends on a specific world state. Each agent has a different set of capabilities, thus the agents differ in their cost for performing a given task, in a specific world state. The goal of the central manager is to maximize his expected utility, defined as a function of the number of tasks being performed and the total payment.

Our long-term goal is to supply the central manager with the optimal protocol which maximizes its utility, given specific environmental settings. However, deriving the central manager's perceived utility when applying a given protocol, requires understanding the agents' equilibrium strategies when such protocol is being used. In this paper, we demonstrate this methodology by proposing and studying a specific protocol. The protocol suggests that the central manager announces a reverse Vickrey auction upon the arrival of a new task. The agent asking the lowest price is the winner and is paid the price requested by the second lowest bid in the auction.

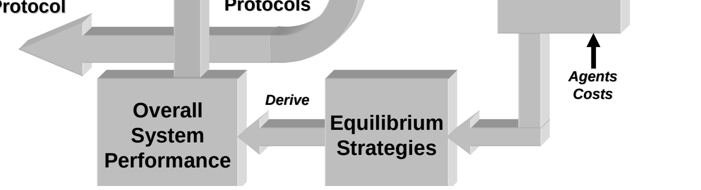

Fig. 1. The general framework for analysis

Given the above protocol, each agent sets its optimal bid for any world state. Winning the current auction will result in an immediate income, but the agent will need to allocate resources in order to perform the task, thus avoiding any additional auctions (possibly associated with better opportunities, e.g. better world states and/or fewer agents to compete with). Therefore the agent's bidding strategy must consider the tradeoff between an immediate gain from the current auction and the expected loss of future opportunities. The agent's evaluation of the above two measures is derived from the analysis of the other agents' strategies in current and future auctions.

It is notable that e-commerce domains do not consider the modulation of competitors' future strategies when setting an agent's bid (e.g. [4, 8]). This is mainly due to the high rate of new agents entering the environment and the difficulty in assessing the new agents capabilities. Nevertheless, in our domain, the number of competing agents is relatively small, and their overall capabilities can be estimated with some probability. Thus an agent's strategy must consider the long term strategies of the other agents in the environment as part of its analysis.

Typical applications in this domain include exploration of remote planets, urban search, and rescue [5]. Consider for example an under-water exploration mission undertaken by a government. Different companies are invited to compete for different tasks (underwater surveys, inspections, mapping, pollution prevention, recovery, etc.) using their own Remotely Operated Vehicles (ROVs). Since the ROVs have been designed and evaluated by different companies, their capabilities to perform a task are dissimilar in different world states. An additional application can be found in an environment where self interested servers, with different configurations and changing loads, are competing for the execution of jobs arriving from an external source, such as universities. Even though each server might have information regarding the other servers' configuration, these servers capability to perform a given job is a function of their load at that specific time. Though each server can calculate its own capability for the performance of a new job, it can only assess the distribution of capabilities of the other servers for executing the job. In both applications, agents have to decide their strategies promptly, according to their current information of the world sate, and their evaluation of their competing agents' capabilities.

The concept of task allocation in a competitive environment is discussed in several works (e.g., [3, 7]). A core application in this domain is the contract net protocol [9]. The main focus of the works cited above is on the commitments and the communication problems that emerge in such an environment. A general architecture and applications are given in [5]. None of these works concern the concept of equilibrium and the modulation of other agents' future strategies. Several works from the adjacent domain of resource allocation involve equilibrium analysis [11, 1, 2, 12]. However, they do not suggest the full extent of changing capabilities and world states or the modulation of all future strategies of the other agents.

In the following section we present the formalization of the model. The equilibrium analysis is given in section 3. In section 4 we suggest an efficient algorithm for the distributed calculation of the agents' equilibrium strategies. Section 5 contains a discussion and computational demonstration of the affect environmental settings have on the perceived equilibrium strategies, as well as a comparison to central allocation. Conclusions and future directions for research are given in section 6.

2 Problem Formulation

We consider an environment with a set $\mathcal{A}$ of self interested agents. We denote an agent g by A_g. A measure for an agent's capability handling a given task, is the duration of time required to successfully complete it. The agent's required time to perform any given task, at time t, is derived from the world state, s_t. Thus, the duration required for agent A_g to perform a task, given a world state s_t at time t, is attained by the function $D^{A_g}(s_t)$. Due to the complexity of world's states, and the changing environment, we assume that for each world state s_t, the duration $D^{A_g}(s_t)$ is drawn from a probability function $P_D(x)$, defined over the interval $[D_{min}, \ldots, D_{max}]$.

The agent's decision must take into consideration two types of costs. The first reflects the cost of participating in an auction (this represents all the costs associated with preparing for the auction, possible auction fees set by the central manager, calculations and evaluation costs, etc.), denoted by C. The second cost, c, is the cost of operating the agent per time unit when performing a task (for simplification we assume all agents share the same c). Considering the servers application outlined above, C may be viewed as the resources a server has to spend in order to evaluate a job characteristics (possibly even performing a small pilot) and determine its capabilities to perform the job. The cost c in this example is the cost of operating the server. For each auction, of k competitors and a given $D^{A_g}(s_t)$, agent A_g calculates its bid, denoted by $B^k(D^{A_g}(s_t))$. The bid is limited by the maximum payment, M, the central manager is willing to pay per task.

An agent will leave the environment only upon winning an auction. The dynamic nature of the environment suggests possible entrance of new agents (either former auction winners once they have completed their tasks, or brand new ones). The number of agents entering the environment between two subsequent auctions is associated with a probability function $P_{new}(z)$, $z = 0, ..., m$, where m is the maximum number of new agents entering. We assume m, as well as $E[P_{new}(z)]$ are relatively small in comparison to the entrance rate in e-commerce.

We assume that all agents are acquainted with the total number of agents, k in the environment, at a given time. We also assume that all agents are familiar with the probability function $P_D(x)$, the cost parameters C and c, the maximum price M and $P_{new}(z)$. Within a given auction, each agent A_g can evaluate only its own duration $D^{A_g}(s_t)$ for performing this task.

3 Equilibrium Analysis

In this section we provide the equilibrium bids' structure and show that no single agent has an incentive to deviate from it.

Notice that the bid $B^k(D^{A_g}(s_t))$ set by agent A_g is determined solely by $D^{A_g}(s_t)$, and the current number of agents competing in the auction, k. Thus the agents' strategy is stationary, i.e., any agent A_1 associated with a duration $D^{A_1}(s_t)$ and k competing agents will bid the same as agent A_2 associated with $D^{A_2}(s_t)$ and k competing agents, where $D^{A_1}(s_t) = D^{A_2}(s_t)$ (regardless of the value t). From this point onward in the paper, we will refer to all durations $D^{A_g}(s_t)$ satisfying $D^{A_g}(s_t) = D_i \in [D_{min}, \ldots, D_{max}]$ ($A_g \in \mathcal{A}$) as D_i. Similarly, we denote the equilibrium bid $B^k(D_i)$ as B_i^k.

Consider a new task arriving at time t, where k agents are situated in the environment. Theorem 1 states that the agents' equilibrium bids when competing for a given task, weakly decrease in their required duration to perform it (given a specific world state).

Theorem 1. *For any two agents A_1 and A_2, having durations D_i and D_j respectively, satisfying $D_i \leq D_j$, the following holds: $B_i^k \leq B_j^k$.*

Proof: As both agents, if winning the auction, will return to the environment in the far future (according to the entrance probability function), their expected opportunity loss when winning the current auction is equal. Thus, any bid B_j^k encapsulating a positive revenue for agent A_2, will also suggest a positive revenue for agent A_1 which can perform the task with a lower cost. □

Consider an agent which is about to attend an auction with a total of k participating agents. The expected revenue of this agent is denoted by R^k. The expected revenue of the agent currently participating in an auction, where its duration for the proposed task is D_i is denoted by $R_{D_i}^k$. The expected revenue R^k is calculated as:

$$R^k = -C + \sum_{y \in [D_{min}, D_{max}]} R_y^k P_D(y) \tag{1}$$

The above equation considers the cost C of participating in the auction and the expected revenues given all the possible world states. The method for calculating R_y^k, ($y \in [D_{min}, D_{max}]$), is given in the following paragraph.

Notice that as the number of competitors an agent A_g has in an auction increases, it is less likely that this agent will win the auction. Thus, we can state the following theorem.

Theorem 2. *An agent's expected future utility from a given auction, monotonically decreases in respect to the number of participating agents in this auction. Formally stated: $R^k > R^{k+1}$ for any k.*

Proof: by induction. □

In order to extract the agents strategies, we first need to understand the influence of the agents' different bidding strategies on their expected utility functions. An agent winning an auction, when bidding B_i^k, will be awarded the mean of second bid values, denoted by $E_{D_i}[second]$. Otherwise, it will move on to the next auction where its expected revenue will be either (assuming k agents in last auction) $\sum_{j=0}^{m} P_{new}(j)R^{j+k-1}$, if one of the other agents won this auction; or $\sum_{j=0}^{m} P_{new}(j)R^{j+k}$, if all agents used a bid higher than M. For simplification, in the rest of this paper we will use: $R^{k+p(j)}$ to denote $\sum_{j=0}^{m} P_{new}(j)R^{j+k}$.

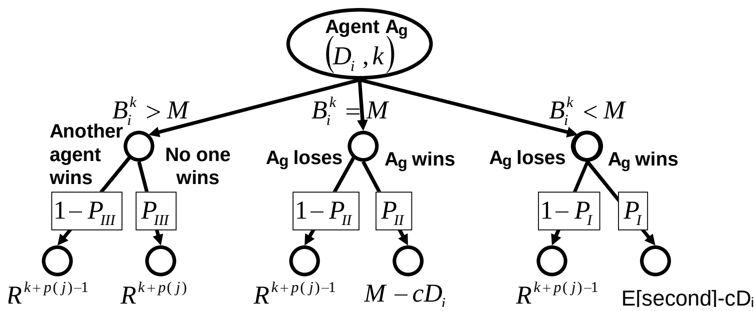

Fig. 2. Agents' bidding scenario

Figure 2, describes the different scenarios for an agent participating in a given auction. We distinguish between 3 types of bids within the equilibrium:

(I) $B_i^k < M$**:** Using bids of this type, the agent wins the auction with a probability P_I, resulting in an expected revenue of $E_{D_i}[second] - cD_i$. Otherwise (with a probability of $1 - P_I$) the agent loses the auction, thus its expected revenue is $R^{k+p(j)-1}$. For agents associated with this type, the bid B_i^k satisfying:

$$B_i^k = R^{k+p(j)-1} + cD_i \tag{2}$$

This bid guarantees their indifference to winning or losing the auction.

(II) $B_i^k = M$**:** Here, with a probability of P_{II} the agent wins the auction, obtaining a reward of M, thus its revenue is $M - cD_i$. Otherwise, the agent's expected revenue is $R^{k+p(j)-1}$.

(III) $B_i^k > M$**:** When using bids of this type the agent will never win the auction. Thus, its expected revenue will be $R^{k+p(j)}$ (if all other agents, as well, bid higher than M, with a probability of P_{III}), otherwise if any of the other agents win the current auction, it will be $R^{k+p(j)-1}$.

In the rest of this paper we refer to the above types as type(I), type(II) and type(III), respectively.

The following theorem suggests several important characteristics of the different equilibrium bid types.

Theorem 3. *Consider an agent* A_g*, with a duration* D_i*:*
(a) if $cD_i < M - R^{k+p(j)-1}$ *holds, then the agent's equilibrium bid is inevitably according to equilibrium type(I). Otherwise, the agent will bid according to type(II) or type (III). (b) if the agent's equilibrium bid is* M *(type (II)), then any other agent* $A_{g'}$ *with* $D_{i'} < D_i$*, not complying with part (a) of the theorem, will bid* M *as well.*

Sketch of Proof: (a) Agents associated with type (I) strategy in equilibrium, gain a better utility by using $B_i^k < M$, than by using $B_i^k = M$. Thus considering Figure 2, the following must hold:

$$P_I(E_{D_i}[second] - cD_i) + (1 - P_I)(R^{k+p(j)-1}) \geq \\ \geq P_{II}(M - cD_i) + (1 - P_{II})R^{k+p(j)-1} \quad (3)$$

Manipulating the above equation we obtain:

$$P_I E_{D_i}[second] - P_{II}M \geq (P_I - P_{II})(R^{k+p(j)-1} + cD_i) \quad (4)$$

Notice that $E_{D_i}[second] \leq M$, as a bid greater than M will never win. Thus substituting $E_{D_i}[second]$ with M, we obtain $M > R^{k+p(j)-1} + cD_i$. Similarly, we can prove that these agents do not have an incentive to move towards type (III) strategy.

(b) As agent $A_{g'}$ is not associated with type (I) strategy our only concern is to prove that this agent's expected benefit when bidding $B_{i'}^k = M$ is greater than the expected benefit when using $B_{i'}^k > M$. Agent A_g is associated with type (II) strategy, for which the following must hold:

$$P_{II}(M - cD_i) + (1 - P_{II})(R^{k+p(j)-1}) \geq (1 - P_{III})R^{k+p(j)-1} + P_{III}R^{k+p(j)} \quad (5)$$

Manipulating the above equation we obtain:

$$P_{II}(M - cD_i) \geq P_{III}R^{k+p(j)} - (P_{III} - P_{II})R^{k+p(j)-1} \quad (6)$$

The above equation will be valid for agent $A_{g'}$, where $D_j \leq D_i$. Thus agent $A_{g'}$ is also associated with equilibrium strategy of type (II), i.e., $B_{i'}^k = M$. □

Fig. 3. Division into types

The above theorem divides the agents into 3 continuous groups as demonstrated in Figure 3. We denote the shortest duration of an agent bidding M, as $\underline{D}$. The longest duration, is denoted $\overline{D}$. Utilizing the above theorem, and the new notations, we can now formulate the expressions for the agents' expected revenues R_i^k, according to the three different types.

Type (I): The expected revenue is composed of 3 components: (a) The agent is a sole best bidder (awarded the expected second bid) ; (b) the agent is the best bidder along with other agents with equal bids (awarded its own bid with a probability equal to the others) ; (c) the agent loses the auction, moving on to the next one. The above is formulated as follows:

$$R_{D_i}^k = \sum_{y \in [i+1, max]} (min(B_y^k, M) - cD_i)(P_D(D \geq y)^{k-1} - P_D(D > y)^{k-1}) + \\ + P_{equal}(B_i^k - cD_i) + (1 - P_D(D \geq D_{i+1})^{k-1} - P_{equal})R^{k+p(j)-1} \quad (7)$$

where $P_{equal} = \sum_{j=1}^{k-1} \binom{k-1}{j} \frac{P_D(D_i)^j P_D(D > D_i)^{k-j-1}}{j+1}$ is the probability the agent will win the auction when one or more additional agents have the same duration D_i.

Type (II): The expected revenue is composed of 2 components: (a) The agent wins the auction with a probability similar to all other agents offering M (awarded M) ; (b) The agent loses the auction, moving on to the next one. The above is formulated as follows:

$$R_{D_i}^k = (M - cD_i)P_{II} + (1 - P_{II})R^{k+p(j)-1} \quad (8)$$

where $P_{II} = \sum_{j=0}^{k-1} \binom{k-1}{j} \frac{P_D(D_i)^j P_D(\underline{D} \leq D \leq \overline{D})^j P_D(D > \overline{D})^{k-j-1}}{j+1}$ is the probability the agent will win the auction when bidding M.

Type (III): In this case the agent inevitably loses the auction thus the only consideration is the number of agents it will compete with in the next auction (affected by whether or not one of the other agents wins the current auction):

$$R_{D_i}^k = P_D(D > \overline{D})^{k-1} R^{k+p(j)} + (1 - P_D(D > \overline{D})^{k-1}) R^{k+p(j)-1} \quad (9)$$

At this point, we have all the necessary equations to calculate the equilibrium bids and the appropriate perceived revenues for each agent. Solving a system of simultaneous equations of types (1-2, 7-9), yields the appropriate strategy parameters. However, in the current structure of the problem, this would be extremely difficult as we need to solve a set of $2 * N + K$ complex equations, where N denotes the number of discrete durations in the interval $[D_{min}, D_{max}]$.

In the next section we show that for an important applicable variant of the above model, a simple algorithm with a complexity $O(N^2 k)$ can be used to calculate the equilibrium bids, and thus the agents' equilibrium revenues.

4 Bounded Size Environments

In this section, we focus on bounded size environments, where no new agents enter the environment ($P_{new}(0) = 1$). The number of available agents never increases, thus once an agent is awarded a task, the number of available remaining agents always decreases by one. This scenario is mostly common in disaster environments (where task durations are relatively long and the arrival rate of new tasks is quite high) or when the group of agents is physically isolated as in the application for the exploration of remote planets [5]. Similarly, we can identify such a scenario in the proposed self-interested servers application (see section 1). Consider a scenario in which servers are competing for the execution of night jobs (assuming they have idle resources only during night). A typical execution of such a job lasts several hours, thus preventing the executing server from competing for additional jobs during the night run. The entire application will start over the next night as all servers will be available again to compete for incoming jobs.

In the bounded size environment, a situation will eventually occur where one of the agents is left alone. Nevertheless, unlike in the general model where new agents might join this agent in subsequent auctions, here the agent will have no future competition. In this situation the agent will undoubtedly bid M, having no other agents to compete with. This is a unique scenario that can not be found in general cases. However, even in this situation the agent might not be interested in winning any given auction. As no competition is expected in future auctions, it might be more beneficial for it to wait for a better world state, in which its capabilities allow it to complete the task in a shorter duration, and thus with a lower cost. In the absence of the cost of participating in an auction, C, the agent will wait until it reaches a world state in which its capability to perform the task is D_{min}. However, the introduction of cost C requires a cost-effective analysis. The agent's optimal strategy, when left alone, is stationary, i.e., an agent refusing to bid in an auction with a duration of D_i will always repeat this strategy. Thus the agent will use a reservation value strategy, bidding M in all world states where its duration D_i is smaller or equal to its reservation value. By denoting the reservation value of the agent as D_r^1 we obtain:

$$R^1(D_r^1) = C + \sum_{y=D_{min}}^{D_r^1} (M - yc) P_D(y) + P_D(D > D_r^1) R^1(D_r^1) \tag{10}$$

This recursive equation is derived from classical *search theory* ([6], and references therein), in which the searcher, having a fixed cost per search stage and a distribution of benefits from possible opportunities, seeks to maximize its overall utility (opportunity utility minus search cost). The above modification, results in:

$$R^1(D_r^1) = \frac{C + \sum_{y=D_{min}}^{D_r^1} (M - yc) P_D(y)}{P_D(D \le D_r^1)} \tag{11}$$

The optimal reservation value, D_r^1, is the one where the agent is indifferent to obtaining an immediate revenue M and to continuing the search with an expected future payoff of $R^1(D_r^1)$. Formally stated:

$$R^1(D_r^1) = \frac{C + \sum_{y=D_{min}}^{D_r^1}(M - yc)P_D(y)}{P_D(D \le D_r^1)} = M - D_r^1 c \qquad (12)$$

Thus $R^1 = R^1(D_r^1)$.

Since the analysis given in section 3, for the case of having more than one agent in an auction, remains unchanged, except for replacing $R^{k+p(j)}$ with R^k, we will not repeat it in this context. Though the modifications of equations (1-9) for the bounded size environment do not take into consideration an increase in the number of agents, the solution process still involves solving a set of $2N + k$ complex simultaneous equations. However, using the analysis given for the case where $k = 1$ and Theorems (1-3) we suggest an efficient algorithm that the agents can use for calculating their equilibrium bids.

For the bounded size environment, we can formulate equation (1) by dividing the expected utility, $\sum_{y \in [D_{min}, D_{max}]} R_y^k P_D(y)$, into:

$$R^k = -C + \sum_{y \in [D_{min}, \overline{D}]} P_D(y) R_y^k + \qquad (13)$$
$$+ P_D(D > \overline{D})(1 - P_D(D > \overline{D})^{k-1})R^{k-1} + P_D(D > \overline{D})^k R^k$$

This way, we can extract R^k, once we have R^{k-1}, as the following algorithm suggests.

Algorithm 1 *An algorithm for calculating the equilibrium bids.*
Input:
k - *Number of initial agents*
M - *The maximum price*
$D[1 : N]$ - *Vector of the discrete durations*
$P_D[1 : N]$ - *Vector of probabilities, associated with D[1:N].*
Output: $B[1 : k][1 : N]$ - *An array of equilibrium bids.*

```
init: for (i=1;i≤N;i++) B[1][i]=M;
calculate R[1][N+1] using Equation (12);
for (i=2;j ≤ k;i++) {
   for (j=N;j ≥ 1;j--) calculate B[i][j] using Equation (2);
   for (Indx=1;(B[i][Indx] ≤ M)and(Indx ≤ N);Indx++) ;
   for (j=Indx;j ≤ N;j++) {
      calculate R[i][N+1] using Equation (13);
      if R[i][N+1] calculated using Equation (8) is greater than
         when calculating using Equation (9) then {
        B[i][j]=M;}
      else return(B[1 : K][1 : N]);
   }
}
return(B[1 : K][1 : N]);
```

The algorithm follows the rules of distinction between the 3 agent types as suggested in the former section, particularly making use of the characteristics described in Theorem 3.

Theorem 4. *(a) Algorithm 1 will always terminate in finite time. (b) If an equilibrium exists for the environment*[4] *the array $B[1:K][1:N]$ will store the equilibrium bids after the algorithm execution is completed. (c) The complexity of the algorithm is $O(N^2k)$.*

Proof:
(a) The loops *01*, and *03-06* are finite.
(b) After executing line *02*, the equilibrium revenue for the case of a single agent, R^1, will be stored in R[1][N+1]. Lines *03-13* are for calculating the equilibrium bids for each scenario of more than one agent in an auction. First, the algorithm isolates the agents that will unavoidably bid less than M. This is done using Equation (2). A temporary R^k (stored in R[k][N+1]) is calculated using Equation (13), based on the bids stored in array B[][]. Then (lines *08-09*) the algorithm starts scanning the agents that currently bid more than M, but have the incentive to deviate towards bidding M, given the values of the current R^k, and the R^{k-1} formerly calculated. The agent characterized with the shortest duration in the group of agents bidding more than M is assessed first, and then the other agents according to their durations. Once there is no additional agent to be added to the group of agents bidding M, a final R^k can be calculated for the case of k competing agents. The value of R^{k+1} is calculated in the same manner, thus, starting with $k = 1$, all R^k values (for $k > 1$) can be calculated recursively.
(c) The complexity of the algorithm is $O(N^2k)$, as it uses two loops bounded by N, for all k agents. The computations made at each stage are immediate. □

In addition to the algorithm being used by each agent for situations where M, k and C are fixed, it can also be utilized by the central manager. By using the algorithm, the central manager can efficiently calculate the expected equilibrium strategies for different environments controlling the above parameters. Thus the expected benefit can be derived for each of these environments, and the central manager's utility can be optimized when the second price auction is used as the allocation protocol.

5 Computational Examples and Discussion

In this section we explore the behavior of the agents' expected revenues and the central manager's total expenses for performing tasks in different environmental settings. Our goal is mainly to emphasize and illustrate some of the characteristics of the equilibrium as analytically developed in the previous sections.

[4] A scenario where an equilibrium does not exist might occur only in discrete environments and this is very rare. Nevertheless, the suggested algorithm can be extended to handle such a scenario.

Calculations were performed using Algorithm 1. We use an environment where the durations are uniformly distributed in the interval $[1, 10]$ with 100 discrete values, and the parameters: $c = 5$, $M = 100$.

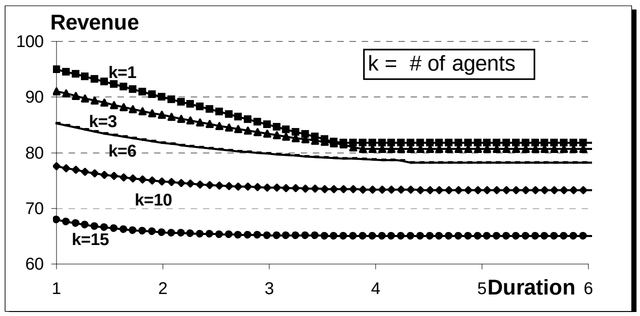

Fig. 4. Expected revenue as a function of duration (Di)

Our first goal is to demonstrate how the number of agents competing in an auction, k, affects each agent's equilibrium revenue and the cost for the central manager. Figure 4, presents the agents' expected revenue, as a function of the duration D_i for different numbers of competing agents, k. As predicted, an agent's expected revenue decreases in its duration (agents characterized with shorter durations are more likely to win an auction and their cost for performing the task is lower). The curves decrease as the agent bids lower than the maximum price and from a specific transition point they become constant as the agents bid higher than the maximum price. The relatively steep change in the agent's expected revenue that can be noticed at the transition point is a measure of the interval in which agents set their equilibrium bids to the maximum price (types II). Notice that unlike the agents bidding higher than the maximum allowed bid (all having the same expected revenue, as they will never win the auction), agents setting their bid to the maximum price seldom win the auction. Thus agents of the latter type differ in their expected revenue as their cost for performing the task is different.

Next, we demonstrate the affect of C (cost of participating in an auction) and the number of agents, k, on the expected expenses of the central manager. The affect is illustrated in Figure 5, depicting the average cost per task, for different C and k values. Simple intuition suggests that increasing the number of agents and decreasing the cost C will enhance competition. Thus the central manager's average expenses per task should decrease. Figure 5 demonstrates that the above is true, though the improvement is limited. This is simply because adding more agents extends the average number of auctions an agent needs to participate in, prior to winning a task (see Figure 6, describing the expected number of auctions until all tasks are allocated, for different values of C and k). At some point, as the number of agents increases, the expected future revenue becomes negative

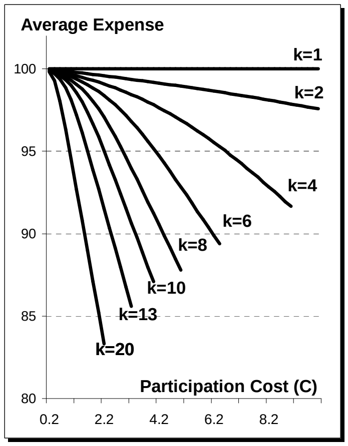

Fig. 5. Average cost per task

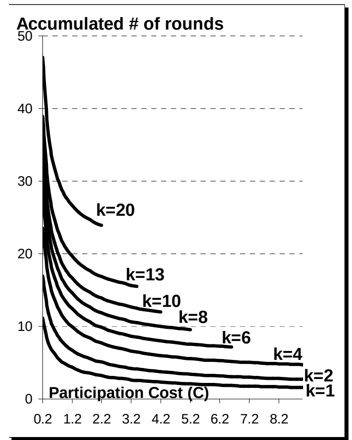

Fig. 6. Accumulated rounds

for any agents participating in this type of auction sequence. In the latter case, the protocol is not feasible as the agents initially prefer not to participate in any of the auctions. The same holds for the increase in cost. If the central manager can control C and k, it will certainly select the combination that will produce the lowest feasible expected cost (in our example, $k = 20$, $C = 2.2$). The behavior presented in Figure 6 can be explained as follows: for small values of C the agents tend to wait for better world states. As C increases, the agents tend to compromise and prefer taking a task even in non-optimal world states, rather than paying C for an additional auction.

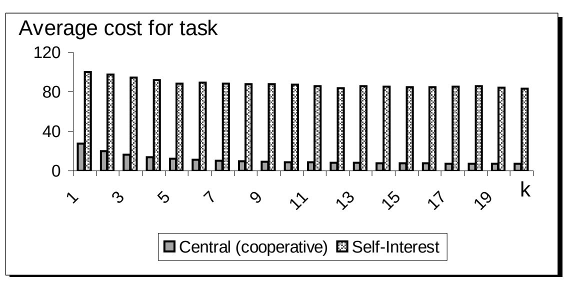

Fig. 7. Comparison with cooperative

We also compare the performance of the distributed allocation to the self-interested agents model with the alternative method where agents are cooperative, and the central manager uses a centralized allocation. For this purpose we consider an allocation method where the central manager allocates each arriving task to the agent with the lowest duration. The expected average expense for a task given k available agents, denoted as Q^k, in this case will be:

$$Q^k = \sum_{y \in [D_{min}, D_{max}]} cy(P_D(D \geq y)^{k-1} - P_D(D > y)^{k-1}) \tag{14}$$

Obviously, we can not obtain such a theoretically expected expense in a distributed environment of self-interested agents. However, this presents a reference point for comparison. Figure 7 compares the two methods, when for each number of agents, k, the optimal C (as derived from Figure 5) was used by the central manager to supply the best average expected cost in the distributed self-interested agents allocation. The graph suggests that the expected expense is significantly smaller in the central allocation method, however this is mainly a function of the parameters C, c and the distribution function. Notice the expenses associated with the central allocation as formulated in Equation (14), give only the "net" cost of agent's operation, c. However, if we also take into consideration leasing fees the central manager needs to pay for having full control over the agents, then for some leasing fee values it will be more economic to use the allocation of tasks to self-interested agents.

6 Conclusions

In this paper we have developed a theoretical framework for analyzing the task allocation process for self-interested agents in a dynamically changing environments. The main challenge in these types of environments is to identify the agents' equilibrium strategies, for any given protocol and specific environmental settings. Typical environments of such domains (e.g., exploration of remote planets, servers competing for executing jobs during idle periods) are characterized by a relatively small number of agents with partial knowledge regarding other agents' capabilities. Thus each agent attempts to maximize its revenue by taking into consideration both the other agents' long term strategies and the influence changes in its own strategy will have on the other agents. This imposes a significant computation complexity as large sets of complex simultaneous equations need to be solved in order to derive the equilibrium. We focused on a specific allocation protocol, where the central manager initiates a second price reverse auction for each arriving task. We developed the equilibrium equations for this environment and identified important characteristics of the equilibrium. Based on the analysis, we were able to produce an efficient algorithm for calculating the agents' equilibrium bids when no new agents appear in the environment. The complexity associated with the proposed algorithm suggests a significant improvement in comparison to the complexity of solving the sets of equilibrium equations. In addition to the usage of the algorithm by the agents, the central planner can now use it to fine tune the parameters of the proposed protocol to maximize its utility.

In future work we intend to use the methodology developed in this paper for exploring the equilibrium of the task allocation process for self-interested agents in respect to additional protocols. This will further enhance the central manager's capabilities to affect its utility by also controlling the type of protocol to be applied.

Acknowledgements: This research was supported in part by NSF under grant #IIS-0208608. Preliminary results of this research were published in the proceedings of ECAI 2004. We would like to thank Uri Yechiali for enlightening discussions.

References

1. E. Altman and R. Hassin. Non-threshold equilibrium for customers joining an m/g/1 queue. in Proceedings of 10th International Symposium on Dynamic Game and Applications, July 2002. Saint-Petersburg, Russia.
2. C. Boutilier, M. Goldszmidt, and B. Sabata, 'Sequential auctions for the allocation of resources with complementarities', in *IJCAI'99*, pp. 527–523, (1999).
3. F. Brandt, W. Brauer, and G. Weiss, 'Task assignment in multiagent systems based on vickrey-type auctioning and leveled commitment contracting', in *CIA'00*, pp. 95–106, (2000).
4. A. Byde, C. Preist, and N. Jennings, 'Decision procedures for multiple auctions', in *AAMAS'02*, (2002).
5. B. Dias and A. Stentz, 'Traderbots: A market-based approach for resource, role, and task allocation in multirobot coordination', Technical Report CMU-RI -TR-03-19, Robotics Institute, CMU, PA, (2003).
6. J. McMillan and M. Rothschild. Search. In Robert J. Aumann and Amsterdam Sergiu Hart, editors, *Handbook of Game Theory with Economic Applications*, pages 905–927. 1994.
7. T. Sandholm and V. Lesser, 'Issues in automated negotiation and electronic commerce: Extending the contract net framework', in *ICMAS'95*, pp. 328–335. The MIT Press: Cambridge, MA, USA, (1995).
8. O. Shehory, 'Optimal bidding in multiple concurrent auctions', *International Journal of Cooperative Information Systems*, **11**(3-4), 315–327, (2002).
9. R. G. Smith., 'The contract net protocol: High-level communication and control in a distributed problem solver', *IEEE Transactions on Computers*, **29**(12), 1104–1113, (1980).
10. T. Smith, T. Sandholm, and R. Simmons. Constructing and clearing combinatorial exchanges using preference elicitation, 2002. AAAI workshop on Preferences in AI and CP: Symbolic Approaches.
11. W. Walsh and M. Wellman, 'Efficiency and equilibrium in task allocation economies with hierarchical dependencies', in *IJCAI'99*, pp. 520–526, (1999).
12. M. Wellman and W. Walsh, 'Auction protocols for decentralized scheduling', *Games and Economic Behavior*, **35**, 271–303, (2001).