Equilibrium Strategies for Task Allocation in Dynamic Multi-Agent Systems

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

1 Introduction

We address a model of self interested 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. The tasks are allocated according to a pre-defined protocol set by a central manager such as a government, a municipality, a company. The central manager lacks the required resources to perform the tasks by itself. The protocol defines the rules for selecting a performer for a task and the appropriate payment to this performer. 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. 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. Hence, a major challenge in task allocation process for self-interested agents in such environments is to identify the agents' equilibrium strategies, for a given protocol and environmental settings.

We suggest a methodology for calculating the agents' equilibrium strategies. Specifically, we focus on a protocol in which, upon arrival of a new task, the central manager starts a Vickrey auction, and the agent who bids the lowest cost wins. In our domains, 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 own bid strategy determination. An example of such an application can be found in an environment where self interested servers, with different configurations and changing loads, compete for the execution of jobs arriving from an external source. The servers set their strategies on-the-fly, according to their current information of the world sate, and their evaluation of their competing agents' capabilities. The main difference from e-commerce domains is that in e-commerce the rate of new agents entering the environment is relatively high so it is unfeasible to consider the modulation of competitors' future strategies when setting an agent's bid (e.g. [2]).

The concept of task allocation in a competitive environment is discussed in several works (e.g., [1, 3]). The main focus of these works is on the commitments and the communication problems that emerge in such an environment. None of them 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 (e.g., [4]). 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.

[1] Department of Computer Science, Bar-Ilan University, Ramat-Gan, 52900 Israel email: sarned@cs.biu.ac.il
[2] Caesarea Rothschild Institute University of Haifa Mount Carmel 31905 Haifa, Israel email: meirav@cs.haifa.ac.il
[3] Department of Computer Science, Bar-Ilan University, Ramat-Gan, 52900 Israel email: sarit@cs.biu.ac.il and Institute for Advanced Computer Studies, University of Maryland, College Park, MD 20742

2 A General Description of the Model

We suggest a model in which different types of tasks arrive dynamically according to a given probability. We consider an environment with a bounded set of self interested agents competing to perform these tasks. A measure for an agent's capability of handling a given task is the duration of time required to successfully complete it. A capability for performing a task depends on a specific world state. Due to the complexity of world states, and the changing environment, we assume that for each world state the duration is drawn from a probability function, $P_D(x)$ where x is in the interval $[D_{min}, \ldots, D_{max}]$. Thus, each agent has a different set of capabilities in a specific world state. The agent's decision must take into consideration two types of costs. The first reflects the cost of participating in an auction (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).

For each auction, of k competitors and a given $D_i \in [D_{min}, \ldots, D_{max}]$, the agent calculates its equilibrium bid denoted by B_i^k. The agents' strategy is stationary, as the required duration for performing a task is derived from $P_D(x)$. 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). We assume that the number of agents entering the environment between two subsequent auctions is associated with a probability function and can be evaluated by the agents. This evaluation is denoted by $p(j)$.

All the agents are acquainted with the total number of agents, k in the environment at a given time, the cost parameters C and c, the maximum price M and $p(j)$. Also, all the agents are familiar with the probability function $P_D(x)$ (though they have no specific information regarding the duration D_i of any of the other agents).

3 Equilibrium Analysis

Given the Vickery auction 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.

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. Once the task associated with the auction is revealed, the agent can evaluate its own required duration D_i for performing the task in the current world state. The expected revenue of the agent in an auction where its duration for the proposed task is D_i is denoted by $R^k_{D_i}$. The expected revenue R^k is calculated as:

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

Consider a new task arriving at a given time, where k agents are situated in the environment. An agent winning an auction, when bidding B^k_i, 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) $R^{k+p(j)-1}$, if one of the other agents won this auction; or $R^{k+p(j)}$, if all agents used a bid higher than M. In order to compute the $R^k_{D_i}$, we distinguish between 3 types of bids within the equilibrium (based on a theorem we prove, which divides the agents into 3 continuous groups in the interval $[D_{min} \dots \underline{D} \dots \overline{D} \dots D_{max}]$).

(I) Bidding less than the maximum price set by the central manager, i.e., $B^k_i < M$. In this case the expected revenue is composed of 3 components: (a) The agent is the 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 and moves on to the next one. The above is formulated as follows (P_{eq} is the probability the agent will win the auction when one or more additional agents have the same duration D_i):

$$R^k_{D_i} = \sum_{y\in[i+1,max]} (min(B^k_y, M) - cD_i)(P_D(D\geq y)^{k-1} - P_D(D> y)^{k-1}) + P_{eq}(B^k_i - cD_i) + (1 - P_D(D\geq D_{i+1})^{k-1} - P_{eq})R^{k+p(j)-1} \qquad (2)$$

(II) Bidding exactly the maximum price, i.e., $B^k_i = M$. The expected revenue in this case 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 (P_{II} is the probability the agent will win the auction when bidding M):

$$R^k_{D_i} = (M - cD_i)P_{II} + (1 - P_{II})R^{k+p(j)-1} \qquad (3)$$

(III) Bidding more than the maximum price, i.e., $B^k_i > M$. 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):

$$R^k_{D_i} = P_D(D > \overline{D})^{k-1} R^{k+p(j)} + (1 - P_D(D > \overline{D})^{k-1}) R^{k+p(j)-1} \qquad (4)$$

where the lowest duration of an agent bidding M, is denoted as $\underline{D}$ and the highest duration as $\overline{D}$.

Solving a system of simultaneous equations of types (1-4), 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 $[D_{min}, D_{max}]$.

We suggest a solution for an important applicable variant of the above model where no new agents enter the environment ($p(j) = 0$). Considering the proposed self-interested servers application (see section 1), for example, we can identify such a scenario, in which servers are competing for the execution of night jobs (assuming they have idle resources only at nighttime). 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 this case, once an agent is awarded a task, the number of available remaining agents, for the next auction, always decreases by one.

In this scenario, once a single agent remains in the environment, it will undoubtedly bid M, having no other agents to compete with. 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 at a lower cost. In the absence of the cost of participating in an auction, the agent would have to wait until it reaches a world state in which its capability for performing the task is optimal. However, the introduction of cost C requires a cost-effective analysis. Thus the agent will use a reservation value strategy, bidding M in all world states where its duration is smaller or equal its reservation value. The analysis of such problems can be deduced from classical *search theory*, 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.

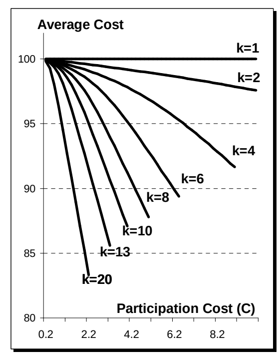

Figure 1. Cost per task

For this case a simple algorithm with a complexity $O(N^2k)$ can be used to calculate the equilibrium bids and revenues. Using the algorithm, we explored the behavior of agents' expected revenues and the central manager's expense for different environmental settings. We investigated several properties of the equilibrium including the affect of cost C for participating in an auction and the number of agents, k, on the expected expenses of the central manager. As shown in Figure 1 increasing the number of agents and decreasing cost C enhances competition and the average expenses of the central manager per task decreases. However, at some point, as the number of agents increases, the expected future revenue becomes negative for any agent participating in this type of auction sequence. This is simply because adding more agents extends the average number of auctions an agent needs to participate in, prior to winning a task, and 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 select the combination that will produce the lowest feasible expected cost (e.g., $k = 20$, $C = 2.2$).

Acknowledgements: This research was supported in part by NSF under grant #IIS-0208608 and by ECCAI travel grant.

REFERENCES

[1] 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).

[2] A. Byde, C. Preist, and N. Jennings, 'Decision procedures for multiple auctions', in *AAMAS'02*, (2002).

[3] 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.

[4] W. Walsh and M. Wellman, 'Efficiency and equilibrium in task allocation economies with hierarchical dependencies', in *IJCAI'99*, (1999).