An Experimental Study on Agent Learning for Market-based Sensor Management

Viswanath Avasarala , *Member, IEEE* , Tracy Mullen , *Member, IEEE*,
David Hall, *Senior Member, IEEE*, and Sudheer Tumu

***Abstract*—Distributed Sensor management, the process of managing or coordinating the use of sensing resources in a distributed environment, is a multi-objective optimization problem. In our earlier work, we proposed MASM (Market-Architecture for Sensor Management), a market-based approach to allocate sensor resources in real-time to various resource requestors. MASM models the multi-objective sensor management problem as a combinatorial-auction based market where the network resources sell goods to the resource requestors. To allow the resource requestors to participate in the market, MASM grants "budgets" to these resource requestors based on their priority to the overall mission. However, for a given budget, self-interested resource requestors or buyers can learn from market-data and adapt their bidding behavior. This paper presents results of an initial experimental study, where the learning behavior of resource requestors is modeled and their effect on market performance is examined.**

I. INTRODUCTION

Sensor management which may be defined as "a process, which seeks to manage or coordinate the use of sensing resources in a manner that improves the process of data fusion and ultimately that of perception, synergistically"[1]. The recent advances in sensor and communication technologies have led to a proliferation of sensor network applications [2][3]. For example, distributed sensor networks that cater to multiple users which query the sensor manager using web-based communication are now feasible. Examples for this multi-sensor, multi-user sensor management scenarios include: (1) network-centric warfare, in which multiple sensing platforms, sensor nets, and individual soldiers with sensors interact to allow rapid tactical situation assessment and threat assessment [3, 4], and (2) condition monitoring of the environment via ground-based, airborne and space-based sensing systems for assessing different issues like weather, pollution etc. In these scenarios, the various users could be trying to accomplish different tasks using network resources. Distributed sensor management is a multi-objective optimization problem. For example, assume that two resource requestors *X* and *Y* are trying to use a sensor network to accomplish different two tasks. *X* is interested in estimating the position and velocity of a target (task T_x). *Y* is interested in searching for the presence of new targets in a certain region (task T_y). The sensor manager should allocate its resources including sensor schedules and communication bandwidth to optimize performance on both T_x and T_y. Approaches like dynamic goal lattices [5], decision theory [6], Bayesian networks [7] have had use in finding the priorities/weights of various objectives in the sensor management optimization. However, a comprehensive paradigm to take mission objectives and the state of the environment simultaneously and directly into consideration during resource allocation is not available. Furthermore, traditional sensor management approaches optimize allocations myopically, ignoring the future impact of current allocations. For example, consider the following simple scenario as illustrated in Figure 1. A sensor network spawns a particular area with varying coverage. Two targets, *A* and *B*, exist in the environment. The sensor network's task is to reduce the uncertainty associated with estimating target positions to below a certain threshold. The tasks with tracking *A* and *B* have the same utility. Target *A* is expected to move along the path shown in Figure 1 from region *R1* to *R2*. Target *B* is expected to move from region *R1* to *R3*. Since *R3* has much higher sensor coverage then *R2*, an optimal sensor should consider allocating more tracking resources to *A* than *B*, when both *A* and *B* are in *R1*. However, myopic sensor managers consider optimization only for the current schedule and therefore do not prioritize between *A* and *B* in *R1*.

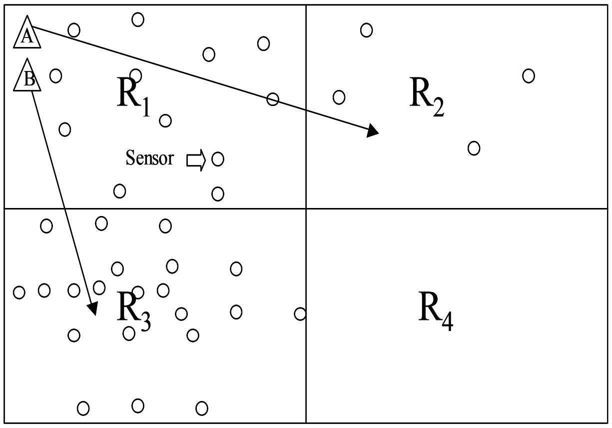

Figure 1 Sample Scenario for non-myopic Sensor Management

To address these issues, we have proposed a market-based resource allocation method for sensor management called MASM (Market-Architecture for Sensor Management) [8] in our earlier work. MASM is based on the market-oriented programming paradigm that uses market

978-1-4244-2764-2/09/$25.00 ©2009 IEEE

algorithms for resource allocation in distributed environments [9]. The genesis of market-oriented programming was in AI community's work in developing economic mechanisms for distributed problem solving and can be first traced to the contract net protocol [10]. Sandholm [11] extended the contract net protocol by integrating the concepts of cost and price. For a detailed review of Market-oriented programming, refer to [9, 12].

In situations like network-centric warfare, multiple resource requestors can use the sensor networks to accomplish different tasks. The sensor manager is responsible for allocating the network resources, so as to cater to these multiple resource requests. MASM models this multi-objective sensor management problem as a competitive market where the consumer agents (resource requestors) bid for and buy sensor network goods. The consumer agents are assumed to be self-interested and participate in the market to maximize their individual utility. The market abstraction simplifies the multi-objective problem into a simpler profit maximization problem. In our earlier work in MASM, we have established the market structure including the auction protocol and pricing mechanisms [8, 13] for this problem. However, we have ignored the crucial aspect of formulating bidding strategies for MASM consumers. Since the sensor manager uses bid prices as the primary means for discerning between the multiple objectives, formulating appropriate bidding strategies is of critical importance to MASM's performance. However, MASM uses a combinatorial auction based mechanism for resource allocation and therefore predicting the optimal bidding strategies is a difficult problem [14]. Furthermore, as illustrated in Figure 1, current sensor network allocations have significant future impact. Therefore, determining optimal bidding strategies involves the difficult problem of predicting future conditions of the market. In this paper, we propose a learning model for MASM agents that can be used to generate bid parameters, which are optimal, under certain market assumptions. We investigate the effect of this learning method on a simulation test bed. To establish the background for this problem, we describe the MASM architecture and the combinatorial-based mechanism used for sensor management in Section II. Section III explains our approach to model agent learning based on market data. Section IV shows the results and provides some concluding remarks.

II. MASM OVERVIEW

The design of MASM is shown in Figure 1. The Mission Manager (MM) is responsible for allocating task responsibilities and budgets to the various agents in the market. The actual details of the sensor network are invisible to the consumer agents. Instead, consumers are allowed to bid on high-level tasks, like "*Track Target X to an accuracy Y*". The sensor manager (SM) stipulates the type of tasks that the consumers can bid on, so that the consumer bids are restricted to tasks SM knows how to accomplish. Consumer bids are of the type <*t, p*> where *t* is the task description that includes the task type and final task quality desired by the consumer, and *p* is the price that the consumer is willing to pay. For example, the task description for a bid to search a particular grid, *x,* so that the uncertainty of target presence (as measured by entropy) < 0.001 is as follows:

(type: search
entity: grid no x
quality: (entropy < 0.001))

The sensor manager uses a combinatorial auction mechanism to find the optimal allocation given the consumer bids. However, the consumer bids are on high-level tasks and the sensor resources are raw sensor schedules or communication bandwidth etc. Thus, some method for bundling goods produced by various sellers is essential to create commodities that consumers are interested in and can bid for. The sensor manager uses a special-purpose protocol CCA (continuous combinatorial auction) to disintegrate the high-level consumer bids to bids for sensor resources. CCA has been designed to decrease computational and communication complexity of the market. The details of CCA are available in [13]. We briefly describe the salient features of CCA. CCA uses discrete time slots to schedule resources. For each time slot, a combinatorial auction is held to determine the optimal allocation of resources.

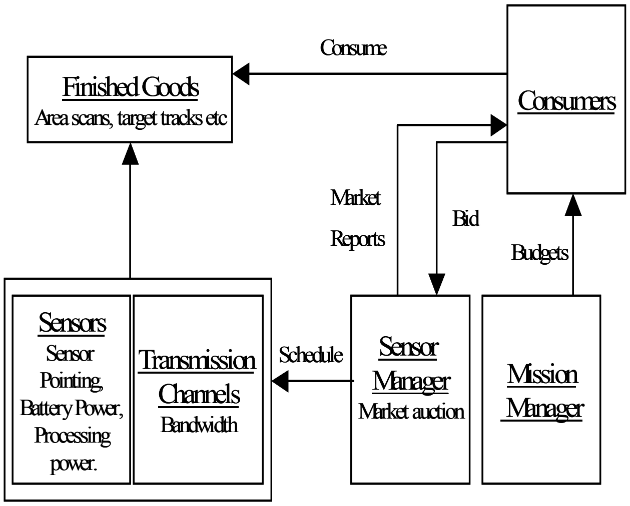

Figure 2 Architecture of MASM

As explained earlier, there is a dichotomy in what the consumers bid on (high-level tasks) and what resources the sensor network has. To address this dichotomy, the CCA uses a bid translation function to create low-level resource bids from high-level task bids. For each time slot *t*, the auctioneer constructs bids on each resource set, *S*, that can be allotted to task *X*. To construct these bids, the auctioneer needs to calculate the price associated with the bid on resource set *S*, for task *X*, based on the consumer bid price P. For this purpose, a novel mechanism of price calculation has been devised. For the task X, the auctioneer computes the bid price for a resource set S as the percentage of the consumer task completed by the resource set multiplied by the actual consumer bid price P. Computation of the percentage of task completed by a resource set S, is in terms of readings of a canonical sensor, A, as follows. Let task X require, on average, n_a consecutive schedules of the standard sensor A to be completed. Instead, if a resource bundle S is allocated to task X during the current round of scheduling, the expected number of standard sensor readings required changes to n'_a. The percentage of task completed by resource set S is equal to the percentage savings in the required number of sensor A readings.

$$f_{S,X} = (n_a - n'_a)/ n_a$$

Equation 1: Calculation of fraction of task computed based on readings of canonical sensor

Then, the bid-price for resource set S at time t for task X with bid price P is

$$p_{S,X} = f_{S,X} * P$$

Equation 2 : Calculation of bid-price for resource set S at time t for task X

The number of resource combinations that can be allocated to a task is exponential in the number of resources. For large sensor networks, we have formulated a special Genetic algorithm that solves the bid translation problem in polynomial time [13]. The representation schema used for the algorithms resembles the SGA algorithm that we developed for determining combinatorial auction winners [15]. First, let the number of bids be k and the number of sensors be n. The representation schema used by the Genetic algorithm is a string of length n, where each string element is a real-valued number between 1 and k. The j^{th} string member represents the bid to which sensor j is allocated. We obtain this string's overall utility as the sum of utilities obtained from individual bids, which are calculated as the sum of individual bid-prices calculated as shown in Equation 2. The Genetic algorithm is anytime and has polynomial complexity.

III. AGENT LEARNING FOR TASK PRIORITIZATION

Several methods based on techniques like dynamic goal-lattices [5], decision-theory [6] and Bayesian networks [7] have been proposed to determine the weights or priorities for the objectives in the multi-objective sensor management problem. However, none of these methods are directly applicable to a market-based scenario where the auctioneer uses bid-prices to differentiate task priorities. In MASM, the Sensor Manager(SM) accepts bids only on a set of pre-defined tasks. The consumer agent is responsible for deconstructing its high level tasks or goals into a sequence of SM acceptable subtasks on which it can bid. Furthermore, it is responsible for calculating appropriate bid prices for these sub-tasks in order to bid in the market. The resource requestors have utility for tasks that they are trying to accomplish but not necessarily for the tasks the SM accepts bids for. Therefore, the consumer agents have to formulate a bidding strategy to formulate the optimal prices for their auction bids. For example, assume that a consumer has utility u_t for destroying a target T. The high level task of destroying T might consist of the following two sub-tasks, for which SM accepts bids: a) search for target T, and b) track target T so that the uncertainty about its position is reduced. The consumer then has to divide its overall utility, u_t into utilities for the two sub-tasks, so that it can formulate bid prices. The optimal weights of the individual sub-tasks are dependant on the system conditions. For example, it might be difficult to search for targets in some environments whereas in others, tracking tracks to high accuracy might be the bottleneck. Utilities for the sub-tasks should take into consideration the competition for resources. For example, sensors that can be used for velocity estimation, like Doppler sensors, might be abundant in the network, making the tracking task a less competitive one. In this paper, a simple agent-learning approach is developed that allows consumer agents to use available market information to adapt dynamically to the system conditions and adapt their bidding strategies, so as to improve their performance.

To formulate bidding strategies, the agent learning approach requires two modules:

a. Data interpretation module: In MASM, consumers can use two key resources of data; 1) The periodic market reports generated by the SM, that provide the current price trends in the market and 2) Historic consumer performance data, which record the market response for previous consumer bids. The consumer agent uses this data to create price-QoS relationships for each sub-task. The constructed relationship should be in the form of a probability distribution, $Pr_{p,,t,\ i}\ (QoS)$, where $Pr_{p,,t,\ i}\ (QoS)$, is the probability that service at level QoS is offered by SM at price, p, at time, t for task, i. Clearly, estimating this distribution is a very difficult problem, since the distribution depends on a number of unpredictable factors like the bidding strategies of other consumers and their task load etc.

b. Response formulation: The agent learning technique should have an optimization routine to determine the optimal bid prices, based on the price-QoS mappings generated by the data-interpretation module. To formulate the optimization problem, it is assumed that the consumer agents' overall performance can be expressed in terms of its performance on the actionable tasks. Based on the mapping between price and quality of service generated by the data interpretation module, the overall agent's performance can be expressed as a function of the bid prices of the individual tasks.

We frame this problem as follows. Let φ_i ($i = 1$ *to* m) be the set of tasks that the consumer has a utility for. Let u_i be the consumer utility of accomplishing the i-th consumer task. Let ϕ_j $\{j = 1$ to $n\}$ be the set of tasks that the sensor network accepts bids for and can accomplish. We assume that each consumer task φ_i consists of a collection of sensor network tasks ϕ_j accomplished in a certain sequence. The sequence of sensor tasks for the i-th consumer task is denoted by χ_i. We also assume that there is one to many mapping between consumer tasks and sensor network tasks. That is, each sensor network task can be a sub-task for one and only one consumer task. For example, in the MASM simulation scenario, consumer has a utility for destroying targets (φ_1 = *'destroy targets'*). To accomplish this task, the consumer has to use the sensor network resources to first search for and detect targets {ϕ_1 = *'search for targets'*}. Then, the consumer has to track targets,{ϕ_2 = *'track targets'*}. In this case,

χ_1 = {*'search for targets'*, *'track targets'*}.

Clearly, establishing the optimal bid prices involves solving a stochastic, multi-period optimization problem. An established method to solve this optimization problem is to frame it as an approximate dynamic programming problem (see [16-19] for related work). These methods solve the equation of the form:

$$V_t = g_t(\Lambda_t) + \alpha E(V_{t+1}(S_{t+1})/S_t, \Lambda_t)$$

where $g_t(\Lambda_t)$ is the utility of taking action Λ_t at time t, S_t is the *"state"* of the system at time t, $E(V_{t+1}(S_{t+1})/S_t)$ is the expected future value of current actions and α is the time discount factor

Equation 3: Recursive Equation for Approximate Dynamic Programming

However, using this formulation is difficult for MASM consumers. This is because modeling the "state" variable involves modeling the state of MASM market and the state of the competing market consumers. In a MASM market, this information is not visible to a consumer. For example, let M_t be the state of market during the scheduling period t. For the combinatorial auction market, M_t can be modeled as having two parts $\{B_t, P_t\}$, where B_t is the set of bids that the sensor manager is handling at t and P_t is the set of resource prices at t. In the MASM market, both B_t and P_t are not visible to the consumer. Similarly, modeling the state of the competing agents is intractable. Moreover, solutions based on forecasting of market state and the competing consumer agents' states will be unreliable because of the involved uncertainty. Therefore, instead of solving this problem, we solved a simpler problem making the following assumptions:

a. Consumers model the market as a fixed price market. In a fixed price market, the different commodities have fixed prices and the only choice consumers have is whether to buy them or not [20]. For fixed-price markets, consumers can construct the price-QoS mapping as a deterministic relationship.

b. Consumers model the price-QoS to be independent of time. That is, consumers assume that the current price-QoS mapping will persist throughout the sensor network operation.

c. Consumers optimize bidding prices under the assumption that they use a constant price for each sensor network task throughout the sensor network operation.

Since the consumers assume that the market is a fixed price market, they can model the market behavior with a series of monotonically increasing functions, γ_{ϕ_j}, $\{j = 1$ *to* $n\}$ where $\gamma_{\phi_j}(p_{\phi_j})$ represents the fraction of sensor task ϕ_j that will be completed in any round of scheduling if the consumer pays a price p_{ϕ_j}. The number of schedules for completing the sensor network task ϕ_j using p_{ϕ_j} as the bid price is

$$\tau_{\phi_j} = [1/\gamma_{\phi_j}(p_{\phi_j})] ,$$

where $[x] = \min\{n \in Z \mid n \geq x)$, Z being the set of integers.

Equation 4: No of Schedules for completing sensor network task ϕ_j

The consumer earns a utility u_i for completing a consumer task φ_i. If χ_i is the set of sensor tasks that comprise φ_i, then the estimated number of bidding cycles to complete consumer task φ_i can be calculated as

$$\Gamma_{\varphi_i} = \sum_{k=1}^{|\chi_i|} \tau_{\chi_i^k}$$

where χ_i^k is the k-th of element of χ_i.

Equation 5: No of Schedules for completing consumer task φ_i

Let P_{χ_i} be the vector of bidding prices $p_{\chi_i^k}$ ($k = 1$ *to* $|\chi_i|$) for the sensor network tasks that comprise φ_i. To calculate the optimal set of bidding price for the i-th consumer task, consumers have to solve the optimization problem:

$$\max_{P_{\chi_i}} \left(u_i - \sum_{k=1}^{|\chi_k|} p_{\chi_i^k} \tau_{\chi_i^k}\right) / \Gamma_{\varphi_i}$$

Equation 6: Optimization equation for profit maximizing consumers.

In the above equation, we have assumed that consumers are profit-maximizing agents. Instead, in a cooperative environment where consumers are unselfish and intend to maximum the number of tasks completed subject to budget constraints, the optimization problem faced by the consumers is of the form:

$$\min_{P_{\chi_i}} (\Gamma_{\varphi_i})$$

subject to

$$(u_i \geq \sum_{k=1}^{|\chi_k|} p_{\chi_i^k} \tau_{\chi_i^k})$$

Equation 7: Optimization equation for altruistic consumers

For the above formulation, we assumed that the consumer can simultaneously pursue multiple consumer tasks and one sensor task pertaining to each consumer task is active at any time. For example, in the MASM simulation scenario involving target destruction explained previously, this translates to the assumption that consumers can either search for targets or track a particular target at any given time, but cannot do both simultaneously. These assumptions were guided by the design of consumer agents in the MASM simulation environment (see Section IV). For alternate scenarios, Equation 6 and Equation 7 have to be reformulated accordingly.

Though this optimization cannot be solved using closed form optimization methods, it is easily solved, compared to Equation 3. This is because, the equations do-not involve expectation of random variables. Therefore, straightforward heuristic-based approaches like genetic algorithms can be used to solve the optimization problems specified in Equation 6 and Equation 7. However, the results of this optimization are dependant heavily on the price-QoS mapping γ_{ϕ_i} , which is constructed based on current market conditions.

Directly using the values obtained by maximizing Equation 6 leads to extensive reliance on current market conditions. For example, a particular consumer agent that is tracking a highly important target might bid aggressively for tracking resources. As a result, the estimate of price-QoS mappings for the track subtask might show a temporary shift. If all the consumer agents adapt rapidly to the new auction outcomes, they cause a permanent price increase in the market. To avoid speculative behavior, we implemented Widrow-Hoff learning to buffer consumer responses, similar to the approach used by Cliff and Bruten's ZIP traders [20].

ZIP traders were originally designed to extend the Zero intelligence agent based simulations of Gode and Sunder [21]. The Widrow-Hoff learning rule with momentum [22] is used by the consumers to adapt their bid prices, based on their bid prices in the previous auction round and the calculated optimal prices. Assume that the bidding price used by the consumer for task ϕ_i in the previous round of scheduling is $p_{\phi_i}^{t-1}$. Let the optimal bidding price calculated according to Equation 6 or Equation be $\bar{p}_{\phi_i}$. The current bidding price for task budget, $p_{\phi_i}^{t}$, is calculated as follows:

$$e = (\bar{p}_{\phi_i} - p_{\phi_i}^{t-1})$$

$$\tau_{\text{updated}} = \kappa * \tau_{\text{current}} + (1- \kappa) * e$$

$$p_{\phi_i}^{t} = p_{\phi_i}^{t-1} + \lambda * \tau_{\text{updated}}$$

Equation 8: Widrow Hoff's based learning

The learning rate parameter, λ, determines the rate at which the budget is changed. During each iteration, the search budget is updated using the current momentum, τ_{current}, which is a weighted sum of the momentum during the previous iteration and the current error. Current error is defined as the difference between the current search-budget and optimal search-budgets. The momentum rate parameter, κ, determines the weight given to the past changes in the calculation of momentum.

IV. EXPERIMENTAL STUDY

A. Simulation Platform

We developed a multi-sensor, multi-consumer, multi-target platform to a serve as a test bed for MASM. The design of the sensor network and the communication channel are derived from [23]. The complete details of the simulation environment are available in [13]. The simulation environment represents a two-dimensional area where targets are uniformly distributed. These targets move around the search area with constant velocity. The search area has several different kinds of sensors, including sensors that provide range and bearing, bearings only sensors, electronic support measure (ESM) sensors. The simulation has a set of software agents that search for and destroy targets. The agents are not provided any sensing resources and they depend on the sensor network for obtaining information about the environment. They bid for sensor resources and update their status based on information provided by the sensor manager. They use the sensor network's resource to search for potential targets and if the probability of target existence within their range exceeds a certain threshold, initialize target tracks. Once a target track is initialized, the agents attack the target if their confidence in the target position is greater than a certain threshold. This is again accomplished by buying sensing resources from the sensor network. Agents are assumed to have a utility u_t *(=1.0)* for destroying a target. To divide the overall utility into utilities for search and track tasks, agents initially use equal priorities. During the simulation run, agents update the search to track budget ratio using the learning described in section III.

B. Implementation Details

1) Market Reports

In the current simulation, SM generates market reports once, every ten rounds. The market report contains the mean, the upper and lower quartiles of bid prices for both search and track tasks and the corresponding Quality-of-Service (QoS) offered to the consumers. The *QoS* is measured in terms of the percentage of the task that is completed.

2) Agent Learning

As explained in III, consumer agent learning involves constructing a price to QoS mapping for each sub task based on market conditions. This simulation implements a simple technique where each consumer agent use the market report provided by the SM to construct an approximate relationship between *QoS* and bid price for each task. Market report in MASM consists of the lower quartile (Q_1), the mean quartile (Q_2) and the upper quartile (Q_3) of bid prices for each sensor network task and the corresponding *QoS* (F_1, F_2, F_3) offered by the SM at that price to the consumer. A linear interpolation of values, provided in the market report, generates an approximate price-QoS mapping. For example, Figure 3 shows the price-QoS mapping(γ_{ϕ_i} as explained in section III) , constructed by a consumer for the search task, during a simulation experiment. MASM consumers are modeled as being unselfish and solve Equation 7 for deciding their bid prices.

Since the simulation framework has only two sensor network tasks, we varied the values of $p_{\phi_{search}}$ and $p_{\phi_{track}}$ in increments of 0.1 between 0 and u_t and used a simple exhaustive search technique for determining the optimal bidding parameters for a given price-QoS mapping. We used Widrow Hoff's learning as described in Equation 8. to buffer consumer response to unsystematic market fluctuations. We used a value of 0.05 for the learning parameter λ, and a value of 0.1 for the momentum rate parameter, κ.

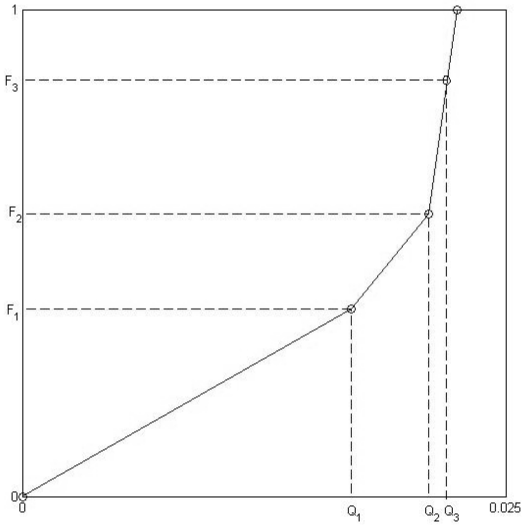

Figure 3 Approximate price-QoS mapping generated by consumer for search task, during a simulation experiment

3) Results

To determine the optimal search to track task budget ratios, simulation experimentations were conducted, varying the percentage of the total budget that is used for search task. Since all the consumer agents are similar in the simulation experiment, they will have similar optimal track to search budget ratios. Figure 4 shows the average number of targets destroyed for different values of the percentage of the search budget in the overall task budget. It is clear that the optimal search budget percentage is around 0.05.

To determine the efficiency of the learning algorithm, the search budget of the five consumers in the simulation were initialized to different values. The learning algorithm was allowed to modify the search budget, after each market report is available. Since the overall utility of target destruction in the simulation is equal to one, search budgets were initialized to the following values.

*search-budget of i-th consumer = 0.1*i;*
*track-budget of i-th consumer = 1 – 0.1*i;*

Figure 5 shows the pattern in which the learning algorithm of each consumer adapts the search-budget as the simulation progresses, for a sample run. Search budgets of the consumers converge to the optimal value for all consumers within 350 rounds of scheduling for the chosen experimental parameters.

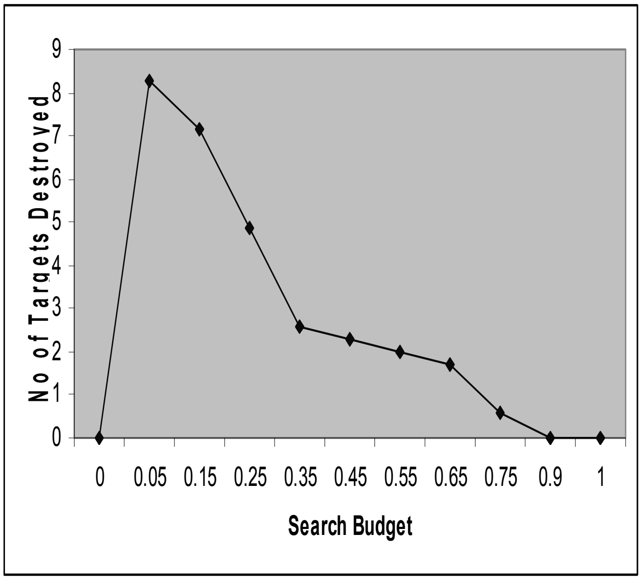

Figure 4 Number of targets destroyed versus search budget (averaged over 10 simulation

experiments)

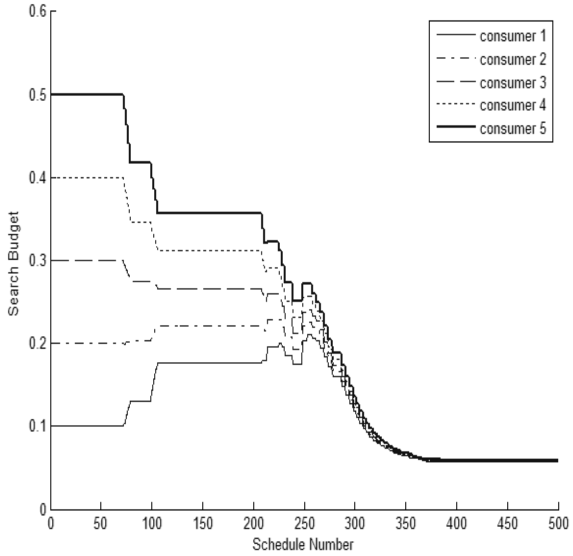

Figure 5 Convergence of search budget to optimal value, based on Widrow-Hoff learning

Clearly, the convergence nature of the algorithm cannot be generalized. The simulation environment has many simplifying design elements, including; a. the symmetry of the market-consumers and b. the small number of tasks. In spite of its simplicity, the learning approach proposed in this paper has some important advantages, compared to conventional static approaches like goal lattices. The learning algorithm requires the agent to analyze market data and compute the optimal weights (prices) for subtasks, based on market conditions and state of the environment in which the agent operates, during each scheduling. As a result, the sub-task weights adapt dynamically to changed system conditions. For example, assume that the target density in the environment decreases suddenly, due to some targets leaving the simulation area. The search task becomes more difficult and the optimization routine automatically increases the search-budget. On the other hand, hard-coded sub-task weights do not respond adequately to the new situation.

Another critical advantage is the non-myopic nature of market-based allocations. Though the proposed method uses myopic market assumptions, it optimizes consumer-bidding behavior considering future schedules also. To understand this feature, consider the scenario illustrated in Figure 1. If MASM is used for resource allocation, it will learn to price the sensors in region *R2* higher than sensors in region *R3* after a certain period of operation, since they are scarcely available (refer to [13] for MASM's pricing mechanism). Consequently, the price-QoS relationship for tracking *A* constructed using historic data will yield a lower *QoS* for the same price as compared to the price-QoS relationship for tracking *B*. Therefore, bids on task *B* will be priced lower than bids on task *A* since bidding agents will learn that task *B* can be completed at a overall lower cost. Consequently, in *R1*, the tracking tasks involving targets that are expected to move to *R2* get preferential resource allocation than tracking tasks that involve targets expected to move to *R3*.

Our algorithm has successfully simplified the highly complex, stochastic multi-period optimization problem involved in market-based scheduling to a simpler deterministic single-period problem based on myopic assumptions about market behavior. To mitigate the effect of myopic market assumptions, we used Widrow-Hoff learning to buffer consumer responses to abrupt market changes. In future, we plan to compare our MASM based allocation approach to algorithms that do-not use price and cost based attractions like contract-net protocol. We are also currently in the process of implementing an extended version of MASM that does not use a centralized SM on a real sensor network based on Crossbow motes. Our experiences from this work should provide more insight into the performance of the proposed agent learning on resource allocation.

The agent-learning technique introduced in this paper is intended to be a preliminary study. A more rigorous approach should investigate approaches that can relax the strong myopic market assumptions that we used. Also, we have ignored the effects of strategic bidding in this paper, by assuming that the agents bid honestly. In a non-cooperative environment with self-interested agents, agents could indulge in dishonest utility revelations. Game-theoretic analysis is required to understand market performance under strategic agent behavior. A preliminary study of the effect of strategic behavior on MASM has been conducted in [13]. However, we note that since MASM hides the actual details of task execution from the consumers, impact of strategic behavior could be reduced.

References:

[1] J.M.Manyika and H. F. D. Whyte, *Data Fusion and Sensor Management: A Decentralized Information Theoretic Approach*. New York: Ellis Horwood, 1994.

[2] M.Balazinska,, A. Deshpande ,M. Franklin, P.Gibbons, J.Gray, M Hansen, M.Liebhold, S. Nath, A. Szalay and V.Tao, "Data Management in the Worldwide Sensor Web, " *IEEE PervasiveComputing*, Vol. 6(2), pp. 30-40, 2007.

[3] D. L. Hall, S. A. H. McMullen, *Mathematical Techniques in Multisensor Data Fusion,* second ed. Norwood, MA: Artech House Publishers, 2004.

[4] D. L. Hall and J. Llinas, *Handbook of Multisensor Data Fusion*: CRC Press, 2001.

[5] K.J. Hintz and J. Malachowski, "Dynamic goal instantiation in goal lattices for sensor management," *Proc. SPIE Signal Processing, Sensor Fusion, and Target Recognition XIV*, vol. 5809, pp. 93–99, 2005.

[6] F. Dormon, V. Leung, D. Nicholson, E. Siva, M. I. Williams, "Information-based decision making over a data fusion network," *Proc. SPIE Signal Processing, Sensor Fusion, and Target Recognition XIV*, vol. 5809, pp. 100–110, 2005.

[7] K. Veeramachaneni, L. Osadciw, and P.K. Varshney, "An adaptive multimodal biometric management algorithm," *IEEE Trans. Systems, Man, and Cybernetics*, pp. 344–356, Aug. 2005.

[8] Tracy Mullen, Viswanath Avasarala, David L. Hall, "Customer-driven sensor management," *IEEE Intelligent Systems* 21(2): 41-49 (2006).

[9] M. Wellman, W. Walsh, P. Wurman, and J. MacKie-Mason, "Auction protocols for decentralized scheduling," *Games and Economic Behavior*, vol. 35, pp. 271-303, 2001.

[10] Davis R,, Smith, R.G, "Negotiation as a metaphor for distributed problem solving," *Artificial Intelligence*, 20, pp 63-109, 1983.

[11] Sandholm T, "An implementation of the contract net protocol based on marginal cost calculations, " *Proc. of the National Conference on Artificial Intelligence,*pp 256-262, AAAI, 1993.

[12] Wellman M.P., "Review of Huberman," *Artificial Intelligence*, 52, pp 205-218, 1991.

[13] V. Avasarala, "Multi-agent systems for data-rich, information-poor environments," in *College of Information Sciences and Technology*, Ph.D. University Park: The Pennsylvania State University, 2006.

[14] P. Cramton,Y. Shoham, and R. Steinberg. Introduction to combinatorial auctions. In P. Cramton, Y. Shoham, and R. Steinberg, editors, *Combinatorial Auctions*. The MIT Press, Cambridge, MA, 2005.

[15] V.Avasarala, H.Polavarapu, T.Mullen: "An approximate algorithm for resource allocation using combinatorial auctions,".*Proc. of IAT 2006*: 571-578.

[16] D. A. Castañon, "Approximate dynamic programming for sensor management," *Proc. 36th Conference on Decision and Control. IEEE*, pp. 1202–1207, December 1997.

[17] Washburn, R., M. Schneider, and J. Fox. "Stochastic Dynamic Programming Based Approaches to Sensor Resource Management." *Proc. 5th International Conference on Information Fusion*, pp 608-615, 2002.

[18] J.L. Williams, J.W. Fisher III, and A.S. Willsky. "An approximate dynamic programming approach to a communication constrained sensor management problem," in *Proc. 8th Int. Conf. Information Fusion*, July 2005.

[19] M. K. Schneider and C.-Y. Chong, "A rollout algorithm to coordinate multiple sensor resources to track and discriminate targets," *proc. SPIE Conference on Signal Processing, Sensor Fusion and Target Recognition*, vol. 6235, April 2006.

[20] D.Cliff and J.Bruten, "Minimal-intelligence agents for bargaining behaviors in market-based environments," Technical Report, HPL-97-91, HP Labs, 1997.

[21] D.K Gode, S.Sunder, "Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality," *Journal of Political Economy*, vol.101, pp. 119-37, 1997.

[22] B. Widrow and M.E.Hoff, "Adaptive switching circuits," *IRE WESCON Conv. Rec.*, vol. 4, pp.96-104, 1960.

[23] G. A. McIntyre and K. J. Hintz, "Sensor management simulation and comparative study," *Proc. SPIE Signal Processing, Sensor Fusion, and Target Recognition VI - The International Society for Optical Engineering*, vol. 3068, Orlando, FL, 20-25, pp 250-260, Apr 1997.