A Consensus Approach to the Assignment Problem: Application to Mobile Sensor Dispatch

Tichakorn Wongpiromsarn, Keyou You and Lihua Xie

***Abstract*— We propose a consensus protocol to solve the assignment problem in a completely distributed manner without any central management or global knowledge. Our protocol is based on the auction algorithm and ensures convergence to an assignment that is within $n\epsilon$ of being optimal where n is the number of agents and ϵ is a design parameter. We then apply the proposed protocol to the mobile sensor dispatch problem to let each mobile sensor determines, in a completely distributed manner, the failed static sensor it should recover so that the total distance traveled by all the mobile sensors is minimized.**

I. INTRODUCTION

The assignment problem, also regarded as the bipartite Maximum Weighted Matching (MWM) problem, is a fundamental problem in combinatorial optimization [1] and provides several applications, especially in the areas of multi-agent coordination, distributed computing and distributed manufacturing. In particular, it has been shown that optimal task and resource allocation can be reduced to an instance of the assignment problem [2]. Another important application is sensor dispatch, an emerging problem in wireless sensor networks where a subset of mobile sensors need to decide how to move in order to optimize certain global objectives while maintaining the coverage ability [3], [4].

Roughly speaking, the assignment problem is the problem of assigning n objects to n agents such that the resulting assignment is optimal according to some predefined cost or benefit. The remarkable Kuhn's Hungarian method [5] solves the assignment problem in polynomial time and is based on a primal-dual method. A price for each object and an (incomplete) assignment of agents and objects are maintained throughout the execution. Starting with an empty assignment and a zero price for each object, in each iteration, the algorithm either adds more assignments or raises object prices to maintain the *complementary slackness* (CS) condition, which essentially asserts that at an optimum, each agent is assigned to the most "profitable" object. The validity of the algorithm relies on a well-known optimality condition which states that an assignment-price pair (x, p) solves the primal and dual problems, respectively, if and only if x is complete (i.e., each of the agents is assigned to an object) and satisfies the CS condition together with p.

This work was partially supported by AFOSR.

Tichakorn Wongpiromsarn is with the Division of Engineering and Applied Science, California Institute of Technology, Pasadena, CA 91125, USA nok@caltech.edu

Keyou You and Lihua Xie are with the School of Electrical and Electronic Engineering, Nanyang Technological University, Singapore 639798, Singapore {youk0001, elhxie}@ntu.edu.sg

The classical version of the Hungarian method is serial and centralized in nature. Bertsekas proposed the auction algorithm that solves the assignment problem in an asynchronous, parallelizable manner [6]. The algorithm can be interpreted as a Jacobi-like relaxation for solving a dual problem and terminates with a sub-optimal solution in polynomial time. It uses the notion of independent single node price changes and *ϵ-complementary slackness* (ϵ-CS), a relaxation of the CS condition by allowing agents to be assigned to objects that come within ϵ of being optimal. As will be further discussed in Section III-A, although the algorithm can be implemented in a distributed manner, it requires certain common knowledge among agents and some central management to ensure that all agents agree on the resulting assignment.

Another approach that efficiently solves the assignment problem in a distributed manner was proposed by Cheng et. al. [7]. The algorithm is based on max-product message passing update rules and does not explicitly require common knowledge among agents or central management. However, it requires that each agent can communicate with all the other agents. This effectively implies that global knowledge can be obtained by each agent. In addition, convergence is only guaranteed when the optimal solution is unique.

In this paper, we extend the auction algorithm [6] by utilizing the consensus-based scheme to eliminate the need for common knowledge and central management. A similar approach has been proposed by Choi et al. [8] which guarantees convergence to an assignment whose benefit is at least half of the benefit of an optimal assignment. In our approach, however, an arbitrary degree of optimality can be achieved as it can be regarded as a design parameter that also affects the convergence rate. Specifically, for any given $\epsilon > 0$, our approach converges to a solution that is within $n\epsilon$ of being optimal where n is the number of agents and the number of operations sufficient to guarantee convergence is inversely proportional to ϵ. We then apply the proposed technique to solve the sensor dispatch problem in a distributed manner. To the authors' knowledge, existing approaches either solve the sensor dispatch problem in a centralized manner [3] or require that each of the mobile sensors can directly talk to all the other mobile sensors in the network [4]. In other words, reference [4] requires that the communication graph is complete. In contrast, our approach only requires that the communication graph is strongly connected.

The remainder of the paper is organized as follows: In Section II, we formally describe the assignment problem for multi-agent systems. In Section III, we briefly present the auction algorithm and results on consensus problems. In

Sections IV and V, we present a consensus approach to solve the assignment problem based on the auction algorithm and provide the proof of optimality and convergence. Finally, in Section VI, we apply the proposed auction-based consensus protocol to the sensor dispatch problem.

II. THE ASSIGNMENT PROBLEM

Consider a problem of matching n agents, $a_1, \ldots, a_n$, with n objects, $b_1, \ldots, b_n$. Each agent a_i has an associated set $\mathcal{N}_i \subseteq \{1, \ldots, n\}$ of indices of its neighbors with whom it may communicate. Let $\mathcal{G}$ be a directed graph that represents the interconnection of the agents. Formally, $\mathcal{G}$ consists of a vertex set $\mathcal{V}(\mathcal{G}) = \{v_1, \ldots, v_n\}$ and an edge set $\mathcal{E}(\mathcal{G})$ where $(v_i, v_j) \in \mathcal{E}(\mathcal{G})$ if $j \in \mathcal{N}_i$.

Let $Q \in \mathbb{R}^{n \times n}$ be an objective matrix whose (i,j) entry, $Q_{i,j}$, represents a benefit of associating agent a_i with object b_j. An assignment $x \in \mathbb{N}^n$ is such that the ith element of x, x_i, represents the index of agent associated with object b_i. For an unassigned object b_i, we let $x_i = 0$. We consider the case where each agent can only be associated with at most one object. Formally, $x_i \neq x_j$ for any $i \neq j$ and $x_i \neq 0$. Let X be the set of all the one-to-one assignments so X contains all the permutations of $\{1, \ldots, n\}$. An optimal assignment x^* maximizes the total benefit. Formally,

$$L^* = \sum_{i=1}^{n} Q_{x_i^*,i} \geq \sum_{i=1}^{n} Q_{x_i,i}, \forall x \in X. \tag{1}$$

Assuming that each agent can only exchange information with its neighbors, we want to find a protocol for each agent to determine its associated object such that the resulting assignment is one-to-one and the total benefit is near optimal.

III. PRELIMINARIES

A. The Auction Algorithm

The assignment problem can be formulated as the following integer linear programming (ILP) problem:

$$\begin{array}{lll} \max\limits_{F \in \mathbb{N}^{n \times n}} & \sum\limits_{i=1}^{n}\sum\limits_{j=1}^{n} Q_{i,j}F_{i,j} & \\ \text{s.t.} & \sum\limits_{j=1}^{n} F_{i,j} = 1, \ \sum\limits_{j=1}^{n} F_{j,i} = 1 & \forall i \in \{1, \ldots, n\} \\ & F_{i,j} \geq 0, & \forall i,j \in \{1, \ldots, n\} \end{array} \tag{2}$$

where F is an $n \times n$ permutation matrix whose (i,j) entry, $F_{i,j}$, is an indicator variable that is equal to 1 if agent a_i is associated with object b_j and is equal to 0 otherwise.

The dual problem to the assignment problem (2) is given by (see, for example, [6])

$$\begin{array}{ll} \min\limits_{p_1,\ldots,p_n,\pi_1,\ldots,\pi_n} & \sum\limits_{j=1}^{n} p_j + \sum\limits_{i=1}^{n} \pi_i \\ \text{s.t.} & p_j + \pi_i \geq Q_{i,j}, \forall i,j \in \{1, \ldots, n\} \end{array} \tag{3}$$

where the dual variables p_i and π_i can be regarded as the price of object b_i and the profit margin of agent a_i, respectively. Since at the optimum, we have $\pi_i = \max_{j \in \{1,\ldots,n\}}\{Q_{i,j} - p_j\}$, π_i can be eliminated and the problem in (3) can be reduced to the following equivalent unconstrained problem:

$$\min_{p_1,\ldots,p_n} \sum_{j=1}^{n} p_j + \sum_{i=1}^{n} \max_{j \in \{1,\ldots,n\}} \{Q_{i,j} - p_j\}. \tag{4}$$

The auction algorithm uses the notion of the ϵ-CS condition to solve the dual assignment problem in (4). The ϵ-CS condition asserts that for any $j \in \{1, \ldots, n\}$ such that $x_j \neq 0$,

$$Q_{x_j,j} - p_j \geq \max_k \{Q_{x_j,k} - p_k\} - \epsilon. \tag{5}$$

The algorithm maintains an assignment x and a price vector p. Each iteration consists of two phases—the bidding phase and the assignment phase. In the bidding phase, each unassigned agent a_i computes the bid for the object with index $j^* \triangleq \arg\max_j \{Q_{i,j} - p_j\}$ given by

$$\begin{array}{lll} bid_{i,j^*} & = & p_{j^*} + (Q_{i,j^*} - p_{j^*}) - w_{i,j^*} + \epsilon \\ & = & Q_{i,j^*} - w_{i,j^*} + \epsilon \end{array} \tag{6}$$

where $w_{i,j^*} = \max_{j \neq j^*}\{Q_{i,j} - p_j\}$. That is, agent a_i increases the price of the most profitable object b_{j^*} by ϵ plus the difference between its highest profit $Q_{i,j^*} - p_{j^*}$ and its second highest profit w_{i,j^*}. Hence, if a_i wins the bid, it will get the object b_{j^*} which is within ϵ of begin optimal.

Next, in the assignment phase, each object b_j that received a bid in the bidding phase increases its price p_j to the highest bid and sets $x_j = i^*$ where i^* is the index of some highest bidder of b_j. The ϵ-CS condition is preserved throughout the execution.

It can be shown [6] that the algorithm terminates with a solution within $n\epsilon$ of being optimal in $O\left(\frac{n^2 \max\{|Q_{i,j}|\}}{\epsilon}\right)$ operations. Thus, if all the $Q_{i,j}$'s are integer and ϵ is chosen to be less than $1/n$, the auction algorithm terminates with an optimal solution. Although the algorithm can be straightforwardly implemented in a distributed manner, it requires common knowledge of the price vector p and assignment x among agents. Moreover, communication is not among agents as it requires some "central" management associated with each object to update its price and assignment based on the received bids.

B. Solving Consensus Problems

The problems of average-consensus, max-consensus and min-consensus have been studied extensively in the field of control and distributed computing. A recent review can be found, for example, in [9]. The consensus protocols and the necessary and sufficient conditions on the communication graph $\mathcal{G}$ that guarantee convergence in the presence of communication time-delays, packet drops, channel noises, link failures and quantization errors have been studied by many researchers [10], [11], [12], [13], [14], [15], [16], [17]. However, the problems of interest are typically limited to the classical average-consensus or max-consensus. In particular, the following max-consensus protocol has been used for determining the max-leader (i.e., the agent with the maximum initial value) [10]:

$$\begin{array}{lll} x^i[t+1] & = & \max_{j \in \mathcal{N}_i \cup \{i\}} x^j[t] \\ f^i[t+1] & = & \begin{cases} f^i[t] & \text{if } x^i[t+1] > x^i[t] \\ \overline{f}^i[t] & \text{otherwise} \end{cases} \end{array} \tag{7}$$

where x^i represents the value of agent a_i, the Boolean variable $f^i \in \{0,1\}$ is the max-flag of agent a_i that indicates whether agent a_i thinks it is the max-leader and $\overline{f}^i$ is the negation of f^i. It was shown in [10] that starting with $f^i[0]=1, \forall i$ and using protocol (7), the value of each agent converges to the value of the max-leader and the max-flag of all the agents except the max-leader converges to zero in $O(n)$ time, assuming that the communication graph $\mathcal{G}$ is strongly connected.

IV. AUCTION-BASED CONSENSUS PROTOCOL

Suppose each agent a_i maintains its local knowledge of the price vector $p^i \in \mathbb{R}^n$ and assignment $x^i \in \mathbb{N}^n$, both initialized to zero. For each $i,j \in \{1,\ldots,n\}$, we let p^i_j and x^i_j be the jth element of p^i and x^i, respectively. We want to develop a consensus protocol to ensure that eventually, all the agents agree on a near optimal, one-to-one assignment. Formally, there exist an $x \in X$ and a time $\hat{t} \in \mathbb{N}$ such that $x = x^i[t], \forall i \in \{1,\ldots,n\}, t \geq \hat{t}$ and $x_i \neq x_j \neq 0$ for any $i \neq j$. In addition, for a given $\delta > 0$, $\sum_{j=1}^n Q_{x_j,j} \geq L^* - \delta$ where L^* is the optimal benefit given in (1).

For each $i,j \in \{1,\ldots,n\}$, we define auxiliary variables β^i_j and α^i_j where β^i_j is the highest bid and α^i_j is the smallest index of the highest bidders of object b_j according to the local knowledge of agent a_i and its neighbors. Formally, $\beta^i_j \triangleq \max_{k \in \mathcal{N}_i \cup \{i\}} p^k_j$ and $\alpha^i_j \triangleq \min_{k \in \mathcal{N}_i \cup \{i\}} \{x^k_j \mid p^k_j = \beta^i_j\}$. Let $\alpha^i = [\alpha^i_1,\ldots,\alpha^i_n]^T$ and $\beta^i = [\beta^i_1,\ldots,\beta^i_n]^T$. We also define the following shorthand notations for any vector α and scalar m: (a) $\alpha\{j \leftarrow m\}$ represents the vector α whose jth element is replaced by m, and (b) $m \in \alpha$ if there exists an element of α whose value equals m.

Fix $\epsilon > 0$. We propose the following auction-based consensus protocol:

$$\begin{aligned} p^i[t+1] &= \begin{cases} \beta^i[t] & \text{if } i \in \alpha^i[t] \\ \beta^i[t]\{j^* \leftarrow Q_{i,j^*} - w_{i,j^*} + \epsilon\} & \text{otherwise} \end{cases} \\ x^i[t+1] &= \begin{cases} \alpha^i[t] & \text{if } i \in \alpha^i[t] \\ \alpha^i[t]\{j^* \leftarrow i\} & \text{otherwise} \end{cases} \end{aligned} \tag{8}$$

where $j^* = \min\left\{ \arg\max_j \{Q_{i,j} - \beta^i_j[t]\}\right\}$ and $w_{i,j^*} = \max_{j \neq j^*}\{Q_{i,j} - \beta^i_j[t]\}$.

Note that $Q_{i,j^*} - w_{i,j^*} + \epsilon = \beta^i_{j^*} + (Q_{i,j^*} - \beta^i_{j^*}) - w_{i,j^*} + \epsilon$. Thus, essentially, each agent first updates its knowledge of the highest bid of each object (i.e., the price vector) and the smallest index of the highest bidders of each object (i.e., the assignment) based on the local knowledge of itself and its neighbors. Then, based on the updated price vector and assignment, if it is not associated with any object, it will increase the bid price of its most profitable object b_{j^*} by ϵ plus the difference between its highest profit $Q_{i,j^*} - \beta^i_{j^*}$ and its second highest profit (w_{i,j^*}) in order to make itself the highest bidder of object b_{j^*} while still keeping b_{j^*} an object that is within ϵ of being most profitable.

Remark 1: For each $i,j \in \{1,\ldots,n\}$, we define α^i_j to be the smallest index of the highest bidders to resolve ties when there are multiple highest bidders for object b_j.

Remark 2: Based on protocol (8), the objective matrix Q only needs to be known "locally". Specifically, each agent a_i only needs to know $Q_{i,j}, \forall j \in \{1,\ldots,n\}$.

Remark 3: From the definitions of j^* and w_{i,j^*}, we see that $Q_{i,j^*} - \beta^i_{j^*}[t] \geq w_{i,j^*}$. Hence, $Q_{i,j^*} - w_{i,j^*} = \beta^i_{j^*}[t] + (Q_{i,j^*} - \beta^i_{j^*}[t]) - w_{i,j^*} \geq \beta^i_{j^*}[t]$. Protocol (8) thus increases $p^i_{j^*}$ by at least ϵ when $i \notin \alpha^i$.

V. ANALYSIS

In this section, we show that protocol (8) converges in finite time to a near optimal assignment. First, we use the notion of ϵ-CS to prove the near-optimality of the resulting assignment assuming that protocol (8) converges. Then, we show that protocol (8) actually converges.

The following lemma shows that the ϵ-CS condition is satisfied throughout the execution.

Lemma 1: For any $i \in \{1,\ldots,n\}$ and $t \in \mathbb{N}$, the assignment-price pair $(x^i[t], p^i[t])$ of agent a_i satisfies the ϵ-CS condition (5).

Proof: Since we assume that each agent starts with an empty assignment (i.e., $x^i[0]$ is a zero vector), it is obvious that the ϵ-CS condition is satisfied at the initial state. We want to show that protocol (8) preserves ϵ-CS throughout the execution. Consider an arbitrary agent a_i and assume that for all $k \in \mathcal{N}_i \cup \{i\}$, the assignment-price pair at time t, $(x^k[t], p^k[t])$, satisfies the ϵ-CS condition. We want to show that $(x^i[t+1], p^i[t+1])$ also satisfies the ϵ-CS condition.

First, we will show that the pair $(\alpha^i[t], \beta^i[t])$ satisfies the ϵ-CS condition. Pick an arbitrary object b_j and fix j for the remainder of the paragraph. If $\alpha^i_j[t] = 0$, then the ϵ-CS condition is automatically satisfied. Thus, we only have to consider the case where $\alpha^i_j[t] \neq 0$. Let κ be an element in the set $\{k \in \mathcal{N}_i \cup \{i\} \mid p^k_j[t] = \beta^i_j[t] \text{ and } x^k_j[t] = \alpha^i_j[t]\}$ and let $\gamma = \alpha^i_j[t] = x^\kappa_j[t]$. By assumption, the assignment-price pair $(x^\kappa[t], p^\kappa[t])$ satisfies the ϵ-CS condition which implies that

$$Q_{\gamma,j} - p^\kappa_j[t] \geq \max_l \{Q_{\gamma,l} - p^\kappa_l[t]\} - \epsilon. \tag{9}$$

In addition, from the definition of β^i, it is obvious that $\beta^i_l[t] \geq p^\kappa_l[t], \forall l \in \{1,\ldots,n\}$ and from the definition of κ, $\beta^i_j[t] = p^\kappa_j[t]$. Thus, $\max_l\{Q_{\gamma,l} - p^\kappa_l[t]\} \geq \max_l\{Q_{\gamma,l} - \beta^i_l[t]\}$. Combining this with (9), we get $Q_{\gamma,j} - \beta^i_j[t] = Q_{\gamma,j} - p^\kappa_j[t] \geq \max_l\{Q_{\gamma,l} - \beta^i_l[t]\} - \epsilon$.

Now we will show that protocol (8) preserves ϵ-CS. Consider an arbitrary agent a_i. The case where $i \in \alpha^i[t]$ is obvious since by protocol (8), $x^i[t+1] = \alpha^i[t]$ and $p^i[t+1] = \beta^i[t]$ and as previously shown, the pair $(\alpha^i[t], \beta^i[t])$ satisfies the ϵ-CS condition. So we only have to consider the case where $i \notin \alpha^i[t]$. Let $j^* = \min\left\{\arg\max_j \{Q_{i,j} - \beta^i_j[t]\}\right\}$ and $w_{i,j^*} = \max_{j \neq j^*}\{Q_{i,j} - \beta^i_j[t]\}$. From (8), we get

$$\begin{aligned} &\max_{j \in \{1,\ldots,n\}} \{Q_{i,j} - p^i_j[t+1]\} \\ &\quad = \max\{Q_{i,j^*} - p^i_{j^*}[t+1], \max_{j \neq j^*}\{Q_{i,j} - p^i_j[t+1]\}\} \\ &\quad \leq \max\{w_{i,j^*} - \epsilon, w_{i,j^*}\} = w_{i,j^*}. \end{aligned}$$

Thus, we obtain

$$Q_{i,j^*} - p^i_{j^*}[t+1] = w_{i,j^*} - \epsilon \geq \max_{j \in \{1,\ldots,n\}} \{Q_{i,j} - p^i_j[t+1]\} - \epsilon.$$

From protocol (8) and Remark 3, it is obvious that $p^i_j[t+1] = \beta^i_j[t], \forall j \neq j^*$ and $p^i_{j^*}[t+1] > \beta^i_{j^*}[t]$. Consider an arbitrary object b_j where $j \neq j^*$ and $\alpha^i_j[t] \neq 0$. Let $\gamma =$

$\alpha^i_j[t] = x^i_j[t+1]$. As proved earlier, the assignment-price pair $(\alpha^i[t], \beta^i[t])$ satisfies the ϵ-CS condition. Thus, we get

$$\begin{aligned}\max_l\{Q_{\gamma,l} - p^i_l[t+1]\} - \epsilon &\leq \max_l\{Q_{\gamma,l} - \beta^i_l[t]\} - \epsilon \\ &\leq Q_{\gamma,j} - \beta^i_j[t] = Q_{\gamma,j} - p^i_j[t+1].\end{aligned}$$

■

From the ϵ-CS condition, the following lemma shows that protocol (8) ensures a near-optimal assignment, assuming that it converges.

Lemma 2: Suppose that protocol (8) converges to a one-to-one assignment x. Then the final assignment is within $n\epsilon$ of being optimal. Formally, $\sum_{j=1}^n Q_{x_j,j} \geq L^* - n\epsilon$ where L^* is the optimal benefit given in (1).

Proof: Let x and p be the final assignment and object price vector, respectively. Then by summing the inequality (5) over j, we get

$$\begin{aligned}\sum_{j=1}^n Q_{x_j,j} &\geq \sum_{j=1}^n (p_j) + \sum_{j=1}^n (\max_k\{Q_{x_j,k} - p_k\}) - n\epsilon \\ &= \sum_{j=1}^n (p_j) + \sum_{i=1}^n \big(\max_j\{Q_{i,j} - p_j\}\big) - n\epsilon\end{aligned}$$

The desired result follows from the duality of (2) and (4). ■

Next, we will prove that protocol (8) actually converges to a one-to-one assignment. For the purposes of the analysis, we define the global price vector $p = [p_1, \ldots, p_n]^T$ and the global assignment $x = [x_1, \ldots, x_n]^T$ where for each $j \in \{1, \ldots, n\}$, p_j and x_j are the globally highest bid and the smallest index of the globally highest bidders, respectively, of object b_j. Formally, p_j and x_j are defined as $p_j \triangleq \max_k p^k_j$ and $x_j \triangleq \min_k\{x^k_j \mid p^k_j = p_j\}$.

The next five lemmas provide the important properties of the global price vector p and the global assignment x.

Lemma 3: For any given time $t \in \mathbb{N}$ and any $i, j, k \in \{1, \ldots, n\}$, if $x^k_j[t] = i$, then $p^k_j[t] \leq p^i_j[t]$.

Proof: Consider arbitrary agents a_i and a_k, object b_j and time $t \in \mathbb{N}$. Suppose $x^k_j[t] = i$. Then according to protocol (8), there exists some time $\hat{t} \leq t$ such that $x^i_j[\hat{t}] = i$ and $p^i_j[\hat{t}] = p^k_j[t]$. ($x^k_j[t] = i$ means that according to the knowledge of agent a_k, a_i is a bidder of object b_j with the highest bid price $p^k_j[t]$. Thus, at some time $\hat{t} \leq t$, agent a_i actually put a bid for object b_j with the bid price $p^i_j[\hat{t}] = p^k_j[t]$.) From (8) and Remark 3, it is obvious that p^i_j cannot decrease with time. Thus, we obtain $p^i_j[t] \geq p^i_j[\hat{t}] = p^k_j[t]$. ■

Lemma 4: For any given time $t \in \mathbb{N}$ and any $j \in \{1, \ldots, n\}$, if $x_j[t] = i \neq 0$, then $x^i_j[t] = \alpha^i_j[t] = i$ and $p^i_j[t] \geq p^k_j[t], \forall k \in \{1, \ldots, n\}$.

Proof: Consider arbitrary $t \in \mathbb{N}$ and $j \in \{1, \ldots, n\}$. Suppose $x_j[t] = i \neq 0$. From the definition of x_j, there exists $k \in \{1, \ldots, n\}$ such that $x^k_j[t] = i$ and $p^k_j[t] = p_j[t] = \max_l p^l_j[t]$. Applying Lemma 3, we get $p^i_j[t] \geq p^k_j[t] = \max_l p^l_j[t]$. But since $p^i_j[t] \leq \max_l p^l_j[t]$, it must be the case that $p^i_j[t] = p^k_j[t] = \max_l p^l_j[t]$. Furthermore, since $x^k_j[t] = i$, according to protocol (8), there exists some time $\hat{t} \leq t$ such that $x^i_j[\hat{t}] = i$ and $p^i_j[\hat{t}] = p^k_j[t] = p^i_j[t]$. From Remark 3, this implies that for any $\tilde{t} \in [\hat{t}, t-1]$, $p^i_j[\tilde{t}+1] = p^i_j[\tilde{t}]$ and therefore $x^i_j[\tilde{t}+1] \leq x^i_j[\tilde{t}]$. (Because the highest bid stays the same, the smallest index of the highest bidders cannot increase.) In particular,

$$x^i_j[t] \leq x^i_j[\hat{t}]. \tag{10}$$

But from the definition of x_j, $i = \min\{x^l_j[t] \mid p^l_j[t] = \max_m p^m_j[t] = p^i_j[t]\}$. Combining this with (10), we get

$$\begin{aligned}i &= x^i_j[\hat{t}] \geq x^i_j[t] \\ &\geq \min\{x^l_j[t] \mid l \in \mathcal{N}_i \cup \{i\}, p^l_j[t] = p^i_j[t] = \max_m p^m_j[t]\} \\ &= \alpha^i_j[t] \geq \min\{x^l_j[t] \mid p^l_j[t] = p^i_j[t]\} = i.\end{aligned}$$

Since the chain of inequalities above can be satisfied only if all the inequalities are equalities, we obtain the desired result $x^i_j[t] = \alpha^i_j[t] = i$. ■

Lemma 5: At any given time $t \in \mathbb{N}$, each agent is associated with at most one object according to the global assignment $x[t]$. Formally, $x_j[t] \neq x_k[t]$ for any $j \neq k$ such that $x_j[t] \neq 0$.

Proof: Suppose, for the sake of contradiction, that at some time $t \in \mathbb{N}$, there exist j and k such that $x_j[t] = x_k[t] = i \neq 0$, i.e., agent a_i is assigned to both objects b_j and b_k. From protocol (8), there must exist two different times $t_1, t_2 \leq t$ when agent a_i assigned itself to object b_j and object b_k, respectively, This effectively implies that $i \notin \alpha^i[t_1]$, $j = \min\{\arg\max_l Q_{i,l} - \beta^i_l[t_1]\}$, $i \notin \alpha^i[t_2]$ and $k = \min\{\arg\max_l Q_{i,l} - \beta^i_l[t_2]\}$. Without loss of generality, we assume $t_1 < t_2$. Following the proof of Lemma 4, we can show that $x^i_j[\tilde{t}] = i, \forall \tilde{t} \in [t_1, t]$ and $i \in \alpha^i_j[t_2]$. However, this contradicts the previous argument that $i \notin \alpha^i[t_2]$. ■

Lemma 6: Once a global object price becomes positive, the object is assigned (in the global sense) and remains assigned throughout the remainder of the execution. Formally, for any $j \in \{1, \ldots, n\}$, if there exists $t \in \mathbb{N}$ such that $p_j[t] > 0$, then $x_j[\hat{t}] \neq 0, \forall \hat{t} \geq t$.

Proof: From protocol (8), it can be easily shown that for any $i, j \in \{1, \ldots, n\}$, if $p^i_j > 0$, then $x^i_j \neq 0$. The result can therefore be straightforwardly obtained from the definition of p_j and x_j. ■

Lemma 7: $p_j[t]$ is bounded for all time $t \in \mathbb{N}$.

Proof: From protocol (8) and the definition of the global price vector p, it is obvious that for the global price p_j of object b_j to increase, there must exist an agent a_i where $i \notin \alpha^i$ and $j = \min\big\{\arg\max_k\{Q_{i,k} - \beta^i_k\}\big\}$. Suppose, to reach a contradiction, that there exists a subset $J^\infty \subseteq \{1, \ldots, n\}$ such that for each $j \in J^\infty$, $p_j \to \infty$. Since each agent only updates the price of the most "profitable" object, this implies that $J^\infty = \{1, \ldots, n\}$ and for each $i \in \{1, \ldots, n\}$, $\max_k\{Q_{i,k} - \beta^i_k\} \to -\infty$. Furthermore, since for any $j \in J^\infty$, $p_j \to \infty$, there exists a time t where $p_j[t] > 0$. Thus, from Lemma 6, $x_j[\hat{t}] \neq 0, \forall j \in J^\infty, \hat{t} \geq t$. That is, after time t, each of the objects will be assigned to some agent. However, since some agent a_i must update the price p_j infinitely often, this means that at infinitely many instances of time, $i \notin \alpha^i$ and by Lemma 4, $i \notin x$. Combining these arguments and using Lemma 5, we can conclude that the number of agents is strictly greater than the cardinality of J^∞ but this contradicts the assumption that there is an equal number of agents and objects. ■

Using the boundedness of the global price vector p and the result on max-consensus problem, the following lemma proves the convergence of protocol (8).

Lemma 8: If $\mathcal{G}$ is a strongly connected digraph, then protocol (8) converges in finite time to a one-to-one assignment.

Proof: Since the global price vector p is bounded (cf. Lemma 7), from Remark 3, eventually, $i \in \alpha^i$ for all $i \in \{1, \ldots, n\}$. Thus, each agent a_i essentially updates its local price vector p^i and assignment x^i using the max-consensus protocol (7). The result thus follows from the result on max-consensus problem [10] and Lemma 5. (Each element of x^i can be regarded as the index of max-leader in [10].) ■

Combining Lemma 2 and Lemma 8, we obtain the desired convergence and near-optimality of protocol (8) as formally stated in the following theorem.

Theorem 1: If $\mathcal{G}$ is a strongly connected digraph, then protocol (8) converges in finite time to a one-to-one assignment x that is within $n\epsilon$ of being optimal. Formally, $\sum_{j=1}^{n} Q_{x_j,j} \geq L^* - n\epsilon$ where L^* is the optimal benefit given in (1).

Remark 4: It can be easily shown that Lemmas 1–7 hold even in the presence of arbitrary bounded communication time delays. Since the max-consensus protocol is robust with respect to arbitrary bounded communication time delays, it can be shown that protocol (8) is also robust with respect to arbitrary bounded communication time delays.

Remark 5: Since the proof of the boundedness of the global price vector (Lemma 7) does not include any assumption on graph topology, we can apply the result on max-consensus problem [10] to show that protocol (8) is robust with respect to changing communication topology.

Remark 6: The proof of Lemma 7 can be modified to handle a more general problem where each agent can be assigned to only a subset of objects (see, e.g., [6] for this more general proof). Thus, protocol (8) can be applied to the problem where each agent can be assigned to only a subset of objects, provided that a valid complete assignment exists.

VI. EXAMPLE

In this section, we show how the proposed auction-based consensus protocol (8) can be applied to sensor dispatch in order to optimize deployment cost without any central control or common knowledge among mobile sensors.

Consider a sensing field with static sensors collecting information from the environment. Some of these static sensors may break down; thus, they need to be recovered by the redundant mobile sensors to maintain the coverage. We assume that each of the mobile sensors can communicate with some subset of other mobile sensors. Furthermore, we consider the case where the number of mobile sensors is exactly the same as the number of failed static sensors. This is not a very restrictive assumption since if the number of failed static sensors is less than the number of mobile sensors, we can add some dummy failed sensors at arbitrary positions very far from each of the mobile sensors to make the numbers equal without affecting the result. If the number of failed static sensors is more than the number of mobile sensors, it can be shown that protocol (8) still works but some of the failed static sensors will not be recovered by any of the mobile sensors. We will apply protocol (8) to determine which failed static sensor should be covered by each of the mobile sensors so that the sum of the distances travelled by the mobile sensors is minimized.

The mobile sensors and the failed static sensors can be regarded as the agents and the objects, respectively, in the assignment problem. Let n be the number of failed static sensors (which is the same as the number of mobile sensors). For each $i, j \in \{1, \ldots, n\}$, we let the benefit $Q_{i,j}$ of associating a mobile sensor a_i with a failed static sensor b_j be the negative of the distance from a_i to b_j. Note that based on Remark 2, each mobile sensor only needs to know its distance to each of the failed static sensors.

We consider the case where $n = 15$ and the initial positions of the mobile and static sensors are as shown in Figure VI. It can be shown that the minimum total distance that the mobile sensors need to travel to cover all the failed static sensors is 1348.90 meters.

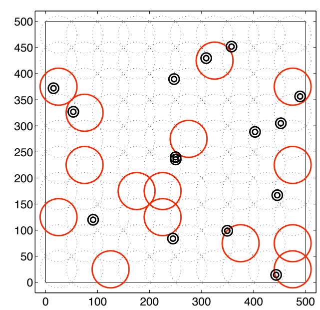

Fig. 1. Sensor dispatch problem. The double circles are the mobile sensors. The dotted (black) circles are the coverage areas of the static sensors. The solid (red) circles are the coverage areas of the failed static sensors.

$\mathcal{G}$	ϵ	# operations required for convergence	Sum of distances traveled by mobile sensors	Difference from optimal distance
complete	1	102	1348.90	0
complete	5	60	1348.90	0
complete	50	22	1444.93	7.12%
complete	100	14	1438.84	6.67%
cycle	1	682	1348.90	0
cycle	5	442	1350.25	0.10%
cycle	50	92	1367.70	1.39%
cycle	100	94	1437.06	6.54%

TABLE I

CONVERGENCE RATES AND TOTAL DISTANCES TRAVELLED BY THE MOBILE SENSORS FOR DIFFERENT TYPES OF COMMUNICATION GRAPH $\mathcal{G}$ AND DIFFERENT VALUES OF ϵ.

Table VI summarizes the results. The assignments between mobile sensors and failed static sensors for the case where $\mathcal{G}$ is a complete graph and for the case where $\mathcal{G}$ is a cycle graph (i.e., each agent can communicate with only one other agent in such a way that the communication graph is complete) are shown in Figure 2 and Figure 3, respectively, with different values of ϵ. For both cases, the optimal solution is obtained with $\epsilon = 1$. In most cases, both the number of operations required for convergence and the degree of optimality decrease as ϵ increases. With the same ϵ, a complete communication graph converges faster than a cycle graph. For the case where $\epsilon = 50$ and $\epsilon = 100$, a cycle communication graph converges to a better solution although it takes longer to converge. This result is not counter-intuitive as protocol (8) only ensures convergence to an assignment that is within $n\epsilon$ of being optimal regardless of the type of communication graph. However, from the results on max-consensus, it is expected that a communication graph with less distance between each pair of agents gives a better convergence rate where the distance between a pair of agents in the graph is defined as the length of the shortest path in the graph that connects those agents.

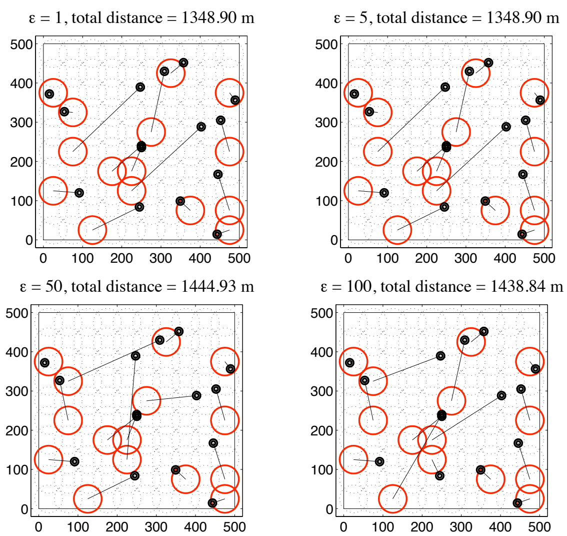

Fig. 2. Simulation results for complete communication graph $\mathcal{G}$. Each line indicates an associated pair of mobile sensor and failed static sensor.

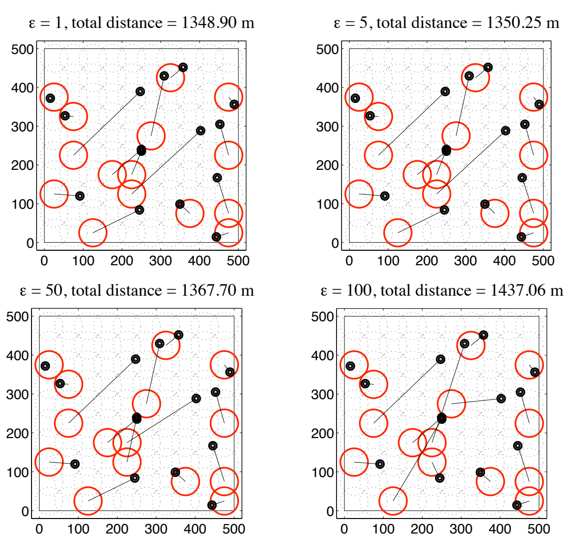

Fig. 3. Simulation results for cycle communication graph $\mathcal{G}$.

VII. CONCLUSIONS AND FUTURE WORK

We presented an auction-based consensus protocol to solve the assignment problem. The protocol allows agents to agree on a near optimal assignment in a completely distributed manner without any central management or global knowledge. Furthermore, it is robust with respect to arbitrary bounded communication time delays and changing communication graph topology. We then apply our approach to the mobile sensor dispatch problem to let each mobile sensor determines, in a completely distributed manner, the failed static sensor it should recover so that the total distance moved by all the mobile sensors is minimized.

Future work includes computing the convergence rate of the proposed protocol with respect to communication graph topology. In addition, we want to investigate the effects of quantization errors on the convergence and optimality results.

VIII. ACKNOWLEDGMENTS

The authors gratefully acknowledge Professor Richard Murray for his generous support and for giving us the opportunity to collaborate on this work. We also acknowledge Ufuk Topcu for thoughtful comments on the paper.

REFERENCES

[1] W. J. Cook, W. H. Cunningham, W. R. Pulleyblank, and A. Schrijver, *Combinatorial Optimization*. Wiley-Interscience, 1997.

[2] B. P. Gerkey and M. J. Matarič, "A formal analysis and taxonomy of task allocation in multi-robot systems," *International Journal of Robotics Research*, vol. 23, no. 9, pp. 939–954, 2004.

[3] Y.-C. Wang, C.-C. Hu, and Y.-C. Tseng, "Efficient placement and dispatch of sensors in a wireless sensor network," *IEEE Transactions on Mobile Computing*, vol. 7, no. 2, pp. 262–274, 2008.

[4] N. H. Viet, N. A. Vien, S. Lee, and T. Chung, "Efficient distributed sensor dispatch in mobile sensor network," *Advanced Information Networking and Applications Workshops, International Conference on*, vol. 0, pp. 1561–1566, 2008.

[5] H. W. Kuhn, "The Hungarian method for the assignment problem," *Naval Research Logistic Quarterly*, vol. 2, pp. 83–97, 1955.

[6] D. P. Bertsekas, "The auction algorithm: a distributed relaxation method for the assignment problem," *Annals of Operations Research*, vol. 14, no. 1-4, pp. 105–123, 1988.

[7] Y. Cheng, M. Neely, and K. M. Chugg, "Iterative message passing algorithm for bipartite maximum weighted matching," in *IEEE International Symposium on Information Theory*, 2006, pp. 1934–1938.

[8] H.-L. Choi, L. Brunet, and J. P. How, "Consensus-based decentralized auctions for robust task allocation," *IEEE Transactions on Robotics*, vol. 25, no. 4, pp. 912–926, 2009.

[9] R. Olfati-Saber, J. A. Fax, and R. M. Murray, "Consensus and cooperation in networked multi-agent systems," *Proceedings of the IEEE*, vol. 95, no. 1, pp. 215–233, 2007.

[10] R. Olfati-Saber and R. M. Murray, "Consensus problems in networks of agents with switching topology and time-delays," *IEEE Transactions on Automatic Control*, vol. 49, no. 9, pp. 1520–1533, 2004.

[11] F. Fagnani and S. Zampieri, "Average consensus with packet drop communication," *SIAM Journal on Control and Optimization*, vol. 48, no. 1, pp. 102–133, 2009.

[12] M. Huang and J. H. Manton, "Coordination and consensus of networked agents with noisy measurements: stochastic algorithms and asymptotic behavior," *SIAM Journal on Control and Optimization*, vol. 48, no. 1, pp. 134–161, 2009.

[13] T. Li and J. Zhang, "Mean square average-consensus under measurement noises and fixed topologies: Necessary and sufficient conditions," *Automatica*, vol. 45, no. 8, pp. 1929–1936, 2009.

[14] S. Kar and J. M. F. Moura, "Distributed consensus algorithms in sensor networks with imperfect communication: Link failures and channel noise," *IEEE Transactions on Signal Processing*, vol. 57, no. 1, pp. 355–369, 2009.

[15] T. Li, M. Fu, L. Xie, and J.-F. Zhang, "Distributed consensus with limited communication data rate," *IEEE Transactions on Automatic Control*, 2009, accepted.

[16] R. Carli, F. Bullo, and S. Zampieri, "Quantized average consensus via dynamic coding/decoding schemes," *International Journal on Robust and Nonlinear Control*, 2009, (Submitted May 2008) to appear.

[17] T. C. Aysal, M. Coates, and M. G. Rabbat, "Distributed average consensus with dithered quantization," *IEEE Transactions on Signal Processing*, vol. 56, no. 10 Part 1, pp. 4905–4918, 2008.