On the Communication Range in Auction-Based Multi-Agent Target Assignment

Marin Lujak and Stefano Giordani

Dip. Ingegneria dell'Impresa, University of Rome "Tor Vergata", Italy

Abstract. In this paper, we consider a decentralized approach to the multi-agent target assignment problem and explore the deterioration of the quality of assignment solution in respect to the decrease of the quantity of the information exchanged among communicating agents and their communication range when the latter is not sufficient to maintain the connected communication graph. The assignment is achieved through a dynamic iterative auction algorithm in which agents (robots) assign the targets and move towards them in each period. In the case of static targets and connected communication graph, the algorithm results in an optimal assignment solution. The assignment results are compared with two benchmark cases: a centralized one in which all the global information is known and therefore, the optimal assignment can be found, and the greedy one in which each agent moves towards the target with the highest benefit without communication with any other agent.

Keywords: Multi-agent, auction algorithm, target assignment, random geometric graph, network optimization.

1 Introduction

Cooperation of a team of robotic agents is a crucial factor of success of any multi-robot mission and in the target assignment it requires shared information, either via a priori information or communication of information gathered during a mission. In this work, we address a multi-agent (robot) target assignment problem with a team of mobile homogeneous cooperating robotic agents and a set of spatially distributed target locations. Each agent has knowledge of all target positions and only a local information about other agents due to a limited communication range $\rho > 0$; furthermore, it must solve the problem of its target assignment individually through the local interaction with the environment, the communication with connected agents exchanging only its partial, possibly outdated information, and the motion towards the assigned target.

A multi-agent system in which agents have limited communication capabilities achieves an optimal target assignment solution if the global information is up- to-date and at the disposal of all decision makers (agents) at least in a multi-hop fashion; sufficient condition for optimal solution is that the communication graph among agents is connected [20,16]. Recall that a graph is said to be connected, if for every pair of nodes there exists a path between them; otherwise

C. Bettstetter and C. Gershenson (Eds.): IWSOS 2011, LNCS 6557, pp. 32–43, 2011.

it is disconnected and formed of more than one connected component. In general, a connected communication graph is very difficult to obtain and dynamic breaking and (re)-establishment of communication links can be due to many communication imperfections, e.g., link bandwidth, delays, need for encryption, message mis-ordering, as well as range constraints due to physical setup of the network or to the power limitations on communications [4]. The frequency of the communication graph's disconnections depends on the choice of the agents' transmitting range: the larger the range, the less likely it is that the network becomes disconnected. The critical transmitting range (CTR) for connectivity is a minimum common value of the agents' transmitting range that produces a connected communication graph. At CTR, all global data is available to each agent. Under this value, starts the creation of detached connected components and when the value of communication range is zero, the agents communicate only over the target states (free or occupied). If CTR is not high enough and, therefore, the information at the disposal of agents (decision makers) is not up-to-date and complete then we naturally expect to achieve a poor group decision making and lower levels of cooperation and hence assignment performance.

Multi-agent target assignment problem is equivalent to the (linear) assignment problem from the operations research field, where, more generally, the agents compete for targets based on some cost function. Hungarian method or Kuhn-Munkres algorithm is one of the most known polynomial time complexity algorithms that solves the problem in the case of a single decision maker. From the other side, there are also distributed algorithms for the assignment problem, one of which is Bertsekas' auction algorithm [1]. In this algorithm, the targets' prices are regulated by the bids which agents give based on the current prices and potential benefit. The bidding process increases the prices of targets which makes them gradually less attractive to agents. Eventually each agent is assigned to a target and the bidding stops. These as well as many other parallel and distributed algorithms for solving the assignment problem, e.g., [3,5,8,17,20], all assume the complete inter-agent communication graph, the latter with possibly substantial but bounded delays was a subject of [3,8,20].

The main issue we explore in this paper is the fluctuation of the assignment quality in respect to the reduction of the quantity of the information exchanged among agents, the maximum step size towards the targets between two consecutive auction runs, and the contraction of the communication range under the CTR value. We present a modified dynamic version of Bertsekas' auction algorithm [1] in which agents without any a priori assignment information negotiate for and dynamically get assigned to targets based on their local environment observations and the information exchanged only with the agents in their connected component (if any). In this way, agents move from any initial configuration to distinct target positions. The final assignment is optimal if the communication graph described previously is connected. The main idea behind our approach to this problem is to let every agent maximally benefit the presence of the network of communicating agents and mutually exchange the information regarding the targets' assignments so as to increase the assignment quality. The size of the communication graph among agents is dependent on each agent's communication range and maximum distance of movement between two consecutive communication rounds (maximum step size); the latter is inversely proportional to the frequency of data exchange among communicating agents.

The experimentation of the algorithm is performed in simulated distinct environment cases with the number of agents ranging from 1 to 30. We show through the simulation results that the quality of the assignment solution in the case of the disconnected communication graph is not a monotonic function of agents' communication range and their maximum step size, but is so on the average.

This paper is organized as follows. In Section 2 we present related work. In Section 3, problem formulation and definitions are presented. In Section 4, we recall the minimum weight perfect matching model for the assignment problem and the background of the Bertsekas' distributed auction algorithm. In Section 5, the modified distributed auction algorithm for the case of disconnected communication graph is presented. Section 6 contains simulation results demonstrating the performance of the modified algorithm compared to the greedy and centralized algorithm. We close the work with the conclusions in Section 7.

2 Related Work

Critical transmitting range (CTR) $\rho > 0$ of n randomly employed agents in environment $[0, l]^2$ was a topic of diverse works, e.g., [11,13,14,15]; in sparse ad-hoc networks, it was analyzed in [13,14,15] using the occupancy theory where nodes (agents) are connected with high probability if $\rho = l\sqrt{(c \log l)/n}$ for some constant $c > 0$ [15].

The conditions evaluated in [13] are how many nodes are required and what transmitting ranges must they have in order to establish a mobile wireless ad-hoc network with connectedness in an obstacle free area of limited size; CTR was proven to be $\rho_M = c\sqrt{\ln n/(\pi n)}$ for some constant $c \geq 1$, where n is the number of network nodes (agents) and M is some kind of node mobility [13].

Connected communication graph among agents whether with or without delays in the target assignment was assumed in some recent works, e.g. [8,9,18]. An assignment of a team of vehicles to distinct targets when the assignment cost function is discounted proportionally to the time it takes an agent to complete the task was a topic of [8] where agents communicate over a connected communication network with substantial delays and perform online assignment simultaneously with movements towards their assigned targets. Moore and Passino in [9] modify the Bertsekas' distributed auction algorithm [1] to minimize the loss of benefit that may occur in the target assignment problem when mobile agents cannot remain stationary (e.g., if they are autonomous air vehicles) and may travel away from their assigned task (target) during the assignment interval, which delays and degrades the collective benefit received. [18] addresses dynamical assignment of tasks to multiple agents using distributed multi-destination potential fields that are able to drive every agent to any available destination, and nearest neighbor coordination protocols to ensure that each agent will be assigned to distinct destination.

On the contrary to the assumption of the complete communication graph, [19] deals with the case without communication among the agents; the authors' approach of dynamic assignment problem is to let every agent explore a sequence of targets and eventually be assigned to the first one that is available.

In wireless networks, a frequently used method for inter-agent communication is broadcasting. It is defined as a dissemination mode of data for which the group of receivers is completely open [6]. When a message is disseminated to locations beyond the transmission range, the mobility of the nodes and multi-hopping is used. A simple broadcast scheme assumes a spanning tree of the network nodes which any node can utilize to send a message. Source node broadcasts a copy of the message to each of its neighbors, which broadcast the message to their neighbors, and so on until the leaf nodes (agents) are reached [10]. How well this service performs depends on its availability and the movements of the nodes that participate in the content dissemination [7].

The correlation between the performance of the solution of the distributed multi-agent target assignment problem and the increase of the number of connected components within the communication graph has not been explored so far to the best of our knowledge.

3 Problem Formulation and Definitions

Considering a time horizon made of T time periods, given are set $A = \{1, \ldots, n\}$ of n collaborative mobile robot agents, represented by points in the plane positioned, w.l.o.g., in a square environment $E = [0, l]^2 \subset \mathbf{R}^2$ of side length $l > 0$, with $p_a(t) \in E$ being the position of robot $a \in A$ at time $t = 1, \ldots, T$, and a set $\Theta = \{1, \ldots, n\}$ of n targets (tasks), with $q_\theta \in E$ being the static position of target $\theta \in \Theta$; the latter is seen as mere object without any processing, memory or communication capacities which are all sustained by robotic agents.

Each agent, $a \in A$, is described by the touple

$$a = \{p_a(t),\ \rho,\ d_{max}^{[a]}\}\ , \tag{1}$$

where $\rho \in \mathrm{R}_{>0}$ is a fixed transmitting (communication) range of agent a's wireless transceiver for limited range communication, and $d_{max}^{[a]}$ is its maximum movement distance (maximum step size) in each assignment interval. At any time t, each agent a knows its $p_a(t)$ and the position q_θ of each target $\theta \in \Theta$. Let $c_{a\theta}(t)$ be the (Euclidean) distance between the position of agent a and target θ.

Examples of such an agent-target setup can be found on the production shop floor where targets are production machines, the position of which is globally known to all the agents (mobile robots or automated guided vehicles). This can be also the case in the supply chain management where trucks or any mobile resource of some distribution center has to optimize its group behavior in the distribution of homogeneous goods to fixed locations of clients.

As the agents move through the environment, their sets of connected agents dynamically change due to a limited communication range. In each period t, each agent a is able to communicate to a set of agents $C_a(t) \subseteq A$ (belonging to the same connected component) reachable in a multi-hop fashion within the communication graph; at any time t, the latter is the random geometric graph (RGG) [2], that is the undirected graph $G(t) = (A, E(t))$ with vertex set A randomly distributed in some subset of R^2, and edge set $E(t)$ with edge $(i, j) \in E(t)$ if and only if

$$\|p_i(t) - p_j(t)\|_2 \leq \rho \ . \tag{2}$$

In this way, two agents which are not within the communication range of each other can communicate over a third agent (communication relay point) in a multi-hop fashion as long as the latter is placed within the communication range of the both. Therefore, agent a together with the set of agents communicating with the same induce a connected subgraph (connected component) of $G(t)$.

We consider the problem of dynamic assignment of set A of robot agents to set Θ of target locations where each agent has to be assigned to at most one target. Total traveled distance by all agents moving towards their targets has to be minimized. We assume that no a priori global assignment information is available and that agents are collaborative and only receive information through their local interaction with the environment and with the connected agents in the communication graph. T is the upper time bound in which all the agents reach their distinct assigned target locations.

When a communication range among the agents participating in the target assignment is lower than CTR, the communication graph is disconnected or even formed of only isolated vertices; in this case the updated information in each period doesn't necessarily arrive to all the agents in the system. The classical Bertsekas' auction algorithm will not guarantee an optimal or even feasible assignment in this context, and certain modifications are necessary. To respond to this open issue, we present a distributed and dynamic version of Bertsekas' Auction algorithm [1]. The latter finds minimum total length routes from agents' initial to distinct final positions (targets) in the multi-agent target assignment problem with a disconnected communication graph. Through this algorithm, we explore the connection between the assignment solution and the sparsity of communication graph varying from a connected communication graph to the case with only isolated vertices.

4 Minimum Weight Perfect Matching Problem

Assuming single decision maker, the problem of finding a feasible assignment of minimum total cost (distance) between agents and targets, can be modeled as the minimum weight perfect matching problem on the weighted bipartite graph $G = (A \bigcup \Theta, E)$ with weight $c_{a\theta}$ on edge $(a, \theta) \in E$ where each target should be assigned to exactly one agent and each agent to exactly one target.

Denoting the above as the primal problem, it is possible to define the dual problem as follows. Let v_θ be a dual variable representing the profit that each agent a will get if it gets matched with the target θ, i.e., v_θ is the value of θ, and u_a the dual variable representing utility of agent a for being matched with certain target related to agent a. Both variables are real and unrestricted in sign. The objective of the dual problem is

$$\max(\sum_a u_a + \sum_\theta v_\theta) \ , \tag{3}$$

subject to

$$u_a + v_\theta \le c_{a\theta} \quad , \forall\ a \ . \tag{4}$$

The dual objective function (3) to be maximized is the sum of the utilities of the agents and the sum of the values of the targets, i.e., the total net profit of the agent-target system. Constraints (4) state that the utility u_a of agent a cannot be greater than the total net cost $(c_{a\theta} - v_\theta)$ of allocation (matching) to target θ, and each agent a would logically want to be assigned to the target θ_a with minimal value of the total net cost,

$$c_{a\theta_a} - v_{\theta_a} = \min_\theta(c_{a\theta} - v_\theta) \ , \ \forall a \ . \tag{5}$$

Agent a is assumed happy if this condition holds and an assignment and a set of prices are at equilibrium when all agents are assumed happy. Equilibrium assignment offers minimal total cost and thus solves the assignment problem while the corresponding set of prices solves an associated dual optimization problem. On the base of the duality in linear programming, $\sum_\theta v_\theta + \sum_a u_a \le \sum_{a\theta} c_{a\theta} \cdot x_{a\theta}$, where $x_{a\theta}$ represents the assignment variable equal to 1 if agent a is assigned to target θ, and 0 otherwise. Therefore, the total net profit of the system cannot be greater than the total assignment cost that agents have to pay for moving to targets, and only at the optimum, those two are equal [1].

Agent a is almost happy with an assignment to target θ_a and a set of target values if the total net cost of being assigned to target θ_a is within ϵ (a positive scalar) of being minimal, that is,

$$c_{a\theta_a} - v_{\theta_a} \le \min_\theta(c_{a\theta} - v_\theta) + \epsilon \ , \forall a \ . \tag{6}$$

The condition (6) is known as ϵ-complementary slackness. An assignment is almost at equilibrium when all agents are almost happy.

In the forward auction algorithm, which we treat in this paper, v_θ can only increase in each iteration and period. For a detailed description of the Bertsekas' forward auction algorithm we invite the reader to refer to the work [1].

5 Dynamic Online Auction with Mobility

In the following, we propose a partially asynchronous algorithm to solve the distributed target assignment problem. We integrate n copies of the modified auction algorithm which is run online by every agent as it moves toward its assigned target in each period. Each agent a keeps in its memory the value $v_{\theta a}$, that is, its most recent knowledge about the actual value of target θ, for each $\theta \in$

Θ, and the set S_a of its most recent knowledge about all the agents' assignments. Both $v_{\theta a}$ and S_a do not have to coincide with the actual target value and agents' assignments, respectively; they may also differ from one agent to another due to the dynamics of their previous communication and local interaction with the environment.

- Initially, i.e., at time $t = 1$ and auction iteration $h = 0$, for each robot agent $a \in A$, set $S_a(t = 1, h = 0)$ of assignments is assumed empty and all target values $v_{\theta,a}(1, 0)$ are set to zero.

At each time $t \in [1, \ldots, T]$, a new round of iterative auction is performed, starting from the assignments and target values locally available to the agents from the end of the round before. In more detail, during iteration h of round t:

- each agent a broadcasts and receives target values $v_{\theta,a}(t, h-1)$ and assignments $S_a(t, h-1)$ to/from all the agents within the connected component the agent a belongs to within the communication graph.
- Each agent updates its local list of the assignments $S_a(t, h)$ and target values $v_{\theta,a}(t, h)$ by adopting the largest value $v_\theta(t, h-1)$ among the set of agents $C_a(t)$ within the connected component for each target of interest $\theta \in \Theta$ and the assignment resulting from this value. However, if the target θ_a is assigned to more than one agent, its assignment is canceled making it unassigned and eligible for bidding in the ongoing round unless if it has already come to its target position q_{θ_a}; in that case, the agent on the target position remains assigned to the target while other agents within the connected component update the value of the target θ_a and cancel the assignment to the same.
- If agent a is unassigned, it finds its best target, calculates the bid value and bids for that target using the following bidding and assignment procedure.

5.1 Bidding

To submit a bid, each agent a unassigned in its partial assignment $S_a(t, h)$:

- finds target θ_a which offers the best possible value $\theta_a = \arg\min_{\theta\in\Theta}\{c_{a\theta}(t) - v_{\theta a}(t, h)\}$, and calculates bid for target θ_a as follows: $b_{a\theta_a}(t, h) = v_{\theta_a a}(t, h) + u_a(t, h) - w_a(t, h) + \epsilon = c_{a\theta_a}(t) - w_a(t, h) + \epsilon$, where $u_a(t, h) = \min_{\theta\in\Theta}\{c_{a\theta}(t) - v_{\theta a}(t, h)\}$ and $w_a(t, h) = \min_{k\neq\theta_a\in\Theta}\{c_{ak}(t) - v_{ka}(t, h)\}$ is the second best utility that is, the best value over targets other than θ_a.
- raises the value of its preferred target by the bidding increment γ_a so that it is indifferent between θ_a and the second best target, that is, it sets $v_{\theta a}(t, h)$ to $v_{\theta a}(t, h) + \gamma_a(t, h)$, where $\gamma_a(t, h) = v_{\theta a}(t, h) - w_a(t, h) + \epsilon$.

The bidding phase is over when all the unassigned agents calculate their bid.

5.2 Assignment

Let $P(\theta_a)(t, h) \subseteq C_a(t)$ be the set of agents with bids pending for target θ_a. Only one agent $a_{coord} \in P(\theta_a)(t, h)$ is responsible for the assignment of target θ_a and if there is more than one agent in $P(\theta_a)(t,h)$, the first one in the lexicographic ordering coordinates the auction for that target.

Each agent $a \neq a_{coord} \in P(\theta_a)(t,h)$, broadcasts its bid $b_{a\theta_a}(t,h)$. Agent a_{coord} receives the bids $b_{k\theta_a}(t,h)$ of all other agents $k \in P(\theta_a)(t,h)$, $k \neq a_{coord}$, regarding θ_a. Following steps are performed to resolve the assignment:

- Agent a_{coord} selects agent $a_{\theta_a} = \arg\max_{a \in P(\theta_a)(t,h)} b_{a\theta_a}$ with the highest bid $b_{a\theta_{a\max}} = \max_{a \in P(\theta)(t,h)} b_{a\theta_a}$.
- If $b_{a\theta_{a\max}} \geq v_{\theta_a a}(t,h) + \epsilon$ then $v_{\theta_a a}(t,h) := b_{a\theta_{a\max}}$, the updated assignment information is broadcasted to all the agents $k \neq a_{coord} \in C_a$ which update their sets of assignments S_a by replacing the current agent assigned to it (if any), with agent a_{θ_a}.
- If $b_{a\theta_{a\max}} < v_{\theta_a a}(t,h) + \epsilon$ then all bids for target θ_a are cleared, no reassignment or target value change is made.

If there are any unassigned agents left within C_a, the assignment algorithm starts again from the bidding phase within iteration $h+1$. This process terminates when each agent $a \in A$ has a target assignment.

5.3 Agent Movement

When all agents are assigned to their respective targets in the final auction iteration, say $h_{fin}(t)$, at round t of the iterative auction, the agent movement phase takes place:

- if agent $a \in A$ is not at its target position p_{θ_a}, it moves one step toward θ_a, covering at most a maximum distance $d_{max}^{[a]}$.
- Once when agent a comes to the position of its assigned and available target, it sets its target value $v_{\theta_a a}(t, h_{fin}(t))$ to $-\inf$, and broadcasts the latter in C_a; this will disable assigning the rest of agents to the same target.

If the initial targets' values $v_\theta(1,0)$ are identically zero, the ϵ-CS condition of (6) will be satisfied because the value of any assigned target is strictly increasing during the auction process and must be at least ϵ. To assure the convergence of the algorithm and to exclude the possibility of multiple exchanges of agents on the target positions, and therefore of infinite loops, we model the target value in the way that it decreases infinitely when an agent arrives on its assigned available target position so that once arrived, an agent remains on the same. Since there is a limited number of bids for any target while there still exist targets that have not yet received any bids, the dynamic auction algorithm will finish when all targets have been given at least one bid and the last agent comes to its target.

The presented algorithm works also when the number of agents differs from the number of targets. In fact, if the number of targets is not less than the number of agents, the assignment process terminates when each agent reaches its assigned target; otherwise, the assignment process terminates when all the unassigned agents realize that all the targets are occupied.

6 Simulation Setup and Results

We simulate a Multi-Agent (robot) System with the forward auction algorithm implemented in MatLab. The dynamic modified auction algorithm was experimented with 4 different kinds of information exchange among agents within each $C_a(t)$, $\forall a \in A$, $t \in \{1, \ldots, T\}$, and $h \in \{0, \ldots, h_{fin}(t)\}$:

- exchange of all assignment data in S_a and target values $v_{\theta,a}(t,h)$ for targets $\theta \in \Theta$ among all agents $a \in C_a$;
- partial exchange of target values $v_{\theta_a,a}(t,h)$ and assignments in S_a regarding only the targets assigned to connected agents $a \in C_a$;
- exchange only of the target value $v_{\theta_a,a}(t,h)$ among the agents assigned to the same target θ_a;
- no information exchange at all. The result is a greedy algorithm where each agent behaves in a selfish and greedy way.

W.l.o.g., and for simplicity, we model the agents as points in plane which move on a straight line towards their assigned targets. Experiments were performed for up to 30 agents in $[0, 100]^2 \subset \mathbf{R^2}$ where the initial robot agent and target positions were generated uniformly randomly. The value of communication range ρ is set from 0 to 50 since a critical communication range for 30 agents calculated by the CTR formula in [13] is 24; furthermore, maximum step size d_{max} varies from 1 to 70 since above the latter value, the number of exchanged messages, crossed distance, and the number of algorithm runs remain unchanged. For each number of robots n varying from 1 to 30, we considered 10 different instances.

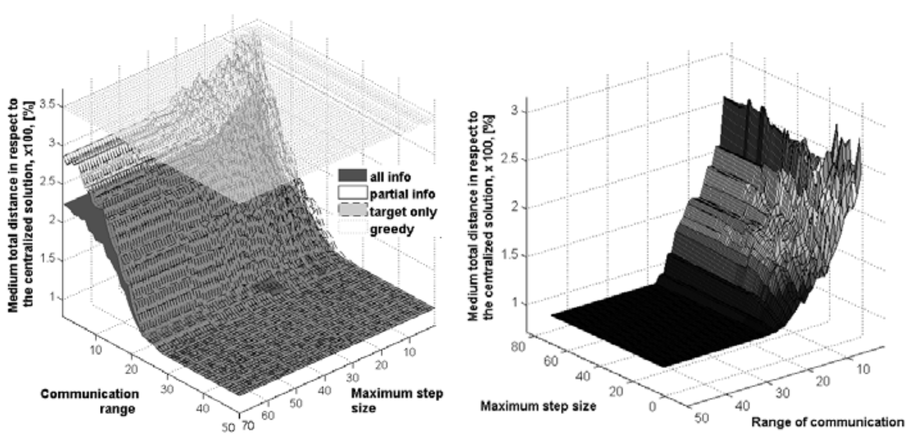

Fig. 1. Left: medium total distance (relative to optimal solution) crossed by 30 agents in $[0, 100]^2$ environment in respect to communication range ρ and maximum step size $d^{[a]}_{max}$; right: assignment dynamics of one representative instance with 30 agents and all information exchange in respect to communication range ρ and max step size $d^{[a]}_{max}$ in the range $[1, 70]$

The average assignment solution (over 10 instances) for the problem with 30 agents is presented in Figure 1, left. The latter represents the variation of the assignment solution in respect to the optimal case with varying maximum step size $d_{max}^{[a]}$ and the communication range ρ.

The presented auction algorithm with included mobility is stable, i.e., it always produces a feasible solution. The experiments show that the average total crossed distance for all the 3 cases with information exchange give the same optimal solution as Bertsekas' algorithm in the case of connected communication graph among all the agents; furthermore, in all of these cases, the solution degrades gracefully with the decrease of the communication range ρ and the increase of the maximum step size $d_{max}^{[a]}$. In the case when the communication range is lower than CTR, the previously mentioned cases give very unexpected results. The number of messages exchanged among the robots in all of the communication models is limited by $O(n^3)$ (Figure 2, right).

In the model with all the assignment information exchange, the assignment solution is, as expected, the closest to the optimal solution, and in the worst case, when the communication range is close to zero and the maximum step size is higher than 60, the average total distance crossed by agents comes up to 260% of the optimal solution.

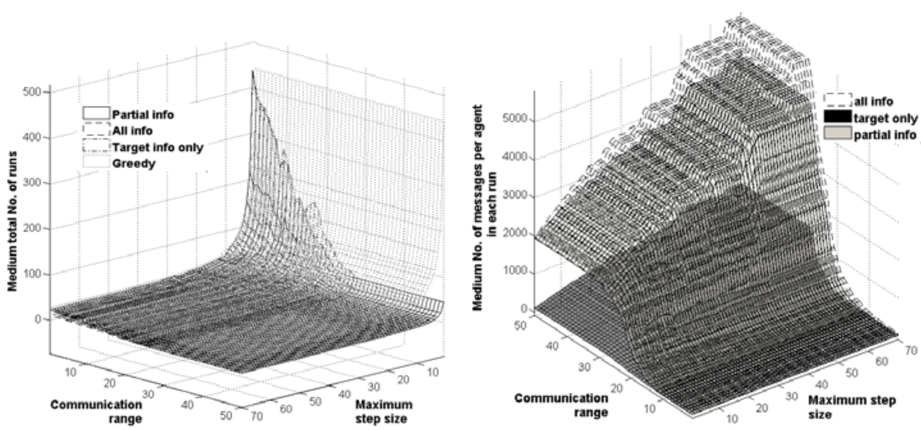

Fig. 2. Left: Medium total number of runs of the algorithms with exchange of all of the information (dashed) vs. partial exchange of information (solid), target only (dash-dot), and greedy (dots); right: medium number of messages per agent in each run for the algorithm with all information exchange (dashed) vs. the one with partial (solid), and target only (black) information exchange

The highest number of exchanged messages per agent and per run, (5733), however, is achieved by the model with all the information exchange. This large amount of information that must be passed over the agents' communication network is a potential problem if the latter is subject to imperfections such as delays, noise, and changing topology.

Unexpectedly, the model with partial information exchange and the model with exchange of information regarding only own target show very close crossed distance results with the worst case approaching 339% and 350% respectively, while in the terms of the medium number of exchanged messages per agent in each run, the model with partial information exchange is up to 4 times more demanding than the target only model.

It is important to notice on the Figure 1, right, that the increase of the communication range when the communication graph is not connected, doesn't necessarily give a better assignment solution; this was the case in all the simulated problem instances.

Considering the communication imperfections, it might be useful to forgo optimality of the assignment for the stability by settling for suboptimal solution by only exchanging the own target information with the agents assigned to the same. The reason not to use a greedy approach is that all of the presented models show, however, better performance than the greedy algorithm which is an upper bound on the average total distance crossed with the value of 351% of the optimal solution, and on the average number of runs, with the worst case of 443 (Figures 1, left, and 2, left). Even if the decentralized method with only target information exchange is cumulatively heavier at the level of exchanged messages (Figure 2, right), the running time of each robot (in the order of $O(n^2)$) results minor in any case than the time required by the greedy algorithm (Figure 2, left).

7 Conclusions

A decentralized implementation of the Bertsekas' auction algorithm has been presented to solve a Multi-Agent (robot) Target Assignment problem for the case when there is no centralized controller or a shared memory and a connected communication graph among agents is not available due to low inter-agent communication range. The agents must operate autonomously and the decentralized implementation is possible on the basis of an exchange of messages. We have assumed homogeneous agents, but the case where agents can be assigned to a specific subset of targets can be handled with a proper elimination of targets within their local assignments. Furthermore, dynamic appearance of targets is supported by the proposed algorithm as long as the positions of targets are known to each agent in each period. The cumulative execution time of the decentralized algorithm and the total number of messages exchanged are in the order of $O(n^2)$, and $O(n^3)$ respectively, resulting in an average computational payload $O(n^2)$ for each agent. A decentralized implementation is intrinsically robust: how robot failures can be handled will be extended in future research.

References

1. Bertsekas, D.P.: Auction algorithms for network flow problems: A tutorial introduction. Comput. Opt. and Applications 1(1), 7–66 (1992)
2. Díaz, J., Mitsche, D., Pérez-Giménez, X.: On the connectivity of dynamic random geometric graphs. In: Proc. of 19th ACM-SIAM Symposium on Discrete Algorithms, pp. 601–610 (2008)

3. Finke, J., Passino, K.M., Sparks, A.: Cooperative control via task load balancing for networked uninhabited autonomous vehicles. In: Proc. of 42nd IEEE Conference on Decision and Control, vol. 1, pp. 31–36 (2003)
4. Gil, A.E., Passino, K.M., Ganapathy, S., Sparks, A.: Cooperative scheduling of tasks for networked uninhabited autonomous vehicles. In: Proc. of 42nd IEEE Conference on Decision and Control, pp. 522–527 (2003)
5. Giordani, S., Lujak, M., Martinelli, F.: A Distributed Algorithm for the Multi-Robot Task Allocation Problem. In: García-Pedrajas, N., Herrera, F., Fyfe, C., Benítez, J.M., Ali, M. (eds.) IEA/AIE 2010. LNCS, vol. 6096, pp. 721–730. Springer, Heidelberg (2010)
6. Keshavarz, H.A., Ribeiro, V., Riedi, R.: Broadcast capacity in multihop wireless networks. In: Proc. of 12th Int. Conf. on Mob. Comp. & Netw., pp. 239–250 (2006)
7. Korkmaz, G., Ekici, E., Özgüner, F., Özgüner, Ü.: Urban multi-hop broadcast protocol for inter-vehicle communication systems. In: Proc. of 1st ACM Int. Workshop on Vehicular Ad-hoc Networks, pp. 76–85 (2004)
8. Moore, B.J., Passino, K.M.: Coping with information delays in the assignment of mobile agents to stationary tasks. In: IEEE Conference on Decision and Control, Institute of Electrical and Electronics Engineers, UK, pp. 3339–3344 (2004)
9. Moore, B.J., Passino, K.M.: Distributed task assignment for mobile agents. IEEE Transactions on Automatic Control 52(4), 749–753 (2007)
10. Murphy, A.L., Picco, G.P.: Reliable communication for highly mobile agents. Aut. Agents and Multi-Ag. Sys. 5(1), 81–100 (2002)
11. Penrose, M.: Random geometric graphs. Oxford University Press, USA (2003)
12. Ren, W., Beard, R.W., McLain, T.W.: Coordination Variables and Consensus Building in Multiple Vehicle Systems. In: Coop. Control: a Post-Workshop: Block Island Workshop on Coop. Control, pp. 171–188 (2004)
13. Santi, P.: The critical transmitting range for connectivity in mobile ad hoc networks. IEEE Trans. on Mob. Comp. 4(3), 310–317 (2005)
14. Santi, P., Blough, D.M.: An evaluation of connectivity in mobile wireless ad hoc networks. In: Proc. of Int. Conf. on Depend. Sys. and Netw., pp. 89–98 (2002)
15. Santi, P., Blough, D.M., Vainstein, F.: A probabilistic analysis for the range assignment problem in ad-hoc networks. In: Proc. of ACM Mobihoc, pp. 212–220 (2001)
16. Savkin, A.V.: The problem of coordination and consensus achievement in groups of autonomous mobile robots with limited communication. Nonlinear Analysis 65(5), 1094–1102 (2006)
17. Smith, S., Bullo, F.: A Geometric Assignment Problem for Robotic Networks. In: Modeling, Estimation and Control, pp. 271–284. Springer, Heidelberg (2008)
18. Zavlanos, M.M., Pappas, G.J.: Dynamic assignment in distributed motion planning with local coordin. IEEE Trans. on Rob. 24(1), 232–242 (2008)
19. Zavlanos, M.M., Pappas, G.J.: Sensor-based dynamic assignment in distributed motion planning. In: IEEE Int. Conf. on Rob. and Autom., pp. 3333–3338 (2007)
20. Zavlanos, M.M., Spesivtsev, L., Pappas, G.J.: A distributed auction algorithm for the assignment problem. In: Proc. of 47th IEEE Conf. on Decision and Control, pp. 1212–1217 (2008)