

3rd International Workshop on Cooperative Robots and Sensor Networks (Robosense-2014)

A Distributed Market-Based Algorithm for the Multi-Robot Assignment Problem

Sahar Trigui^{a,b,*}, Anis Koubaa^{b,c,d}, Omar Cheikhrouhou^e, Habib Youssef^f, Hachemi Bennaceur^g, Mohamed-Foued Sriti^g, Yasir Javed^g

^aUniversity of Manouba, ENSI, Tunisia.

^bCooperative Intelligent Networked Systems (COINS) Research Group, Riyadh, Saudi Arabia.

^cPrince Sultan University, College of Computer and Information Sciences, Riyadh, Saudi Arabia.

^dCISTER/INESC-TEC, ISEP, Polytechnic Institute of Porto, Porto, Portugal.

^eUniversity of Sfax, Computer and Embedded Systems Lab, Tunisia.

^fUniversity of Sousse, PRINCE Research Unit, Tunisia.

^gAl Imam Mohammad Ibn Saud Islamic University, Computer Science Department, Saudi Arabia.

Abstract

Assigning tasks to a set of robots is a fundamental problem in robotics. It consists in finding the best task assignment to the available robots. In this paper, we present two distributed market-based algorithms to solve the assignment problem where n robots compete for n tasks with the assumption that each robot can be assigned to only one task. The first algorithm, called DMB, represents a Distributed Market-Based algorithm where each robot bids for every task. The second algorithm is an extension of the DMB. It consists in swapping tasks between robots in order to improve the efficiency of the whole assignment. We conducted both real-world experimental testing, and MATLAB simulations to evaluate performance of the proposed algorithms and compare them against the centralized Hungarian algorithm in terms of traveled distance. Simulation results show that the IDMB algorithm produces near optimal solutions and in several cases it gives the optimal solution. In addition, we demonstrated the feasibility of our algorithms through real-world experimentation on robots.

© 2014 The Authors. Published by Elsevier B.V.

Selection and peer-review under responsibility of Elhadi M. Shakshuki.

Keywords: Task Assignment, Market-based Algorithms, Multi-Robot Task Allocation, ROS.

1. Introduction

In our days, robotics is a key technology that is integrated into human being's life. In fact, it can be applied to several domains and contexts in order to perform complex missions. The use of such systems will ease several tasks and may lead to a higher accuracy. Robotic applications become more and more complex, difficult and in some cases surpass a single robot capabilities (e.g. pushing a weighty object). For this reason, researchers tend to develop multi-robot systems (MRSs) where a set of robots communicates and collaborates together. MRSs are generally

* Sahar Trigui. Tel.: +216-96-708-395.

E-mail address: sahar.trigui@coins-lab.org

employed in order to improve the performance of the overall application. MRSs face several challenges, but the most typical problem is the Multi Robot Task Allocation (MRTA) problem. It consists in finding the efficient allocation mechanism in order to assign different tasks to the set of available robots. Towards this objective, robots will work as cooperative agents. The MRTA problem can be formulated as follows: given a set of constraints, n robots and m tasks, the objective consists in ensuring an efficient assignment of tasks under consideration and thus minimizes the overall system cost. The MRTA can be considered as assignment problem¹ that can be formulated as follows: Consider a set of n individuals $i = 1, 2, \dots, n$ available to perform a set of n jobs $j = 1, 2, \dots, n$. The assignment consists in giving one job j to the individual i in such a way no job is assigned to two different individuals². The assignment problem consists in allocating jobs to workers in order to maximize the performance.

In the literature, there are three typical approaches devised for ensuring multi-robot coordination^{3,4}, namely: (i.) the centralized approach: it assumes the knowledge of global information by a central agent (e.g. control station), which is able to calculate an optimal (or near-optimal) solution to the allocation problem, (ii.) the distributed approach: decisions (or local solutions) are based on local information for each agent performing the task (e.g. robot) (iii.) the market-based approach: it assumes that solutions are built based on a bidding-auctioning procedure between the working agent (e.g. robot) and a coordinator for allocating tasks for low-cost bidders. Three categories exist to classify the MRTA problems⁵: (i.) Single-task robots (ST) vs. multi-task robots (MT): In the ST approach, robots can execute one task at most. However, in the MT approach, robots can simultaneously execute multiple tasks. (ii.) Single-robot tasks (SR) vs. multi-robot tasks (MR): Using the SR approach, each task requires exactly one robot to be completed. However, in the MR approach, some tasks need multiple robots. (iii.) Instantaneous assignment (IA) vs. time-extended assignment (TA): In the IA approach, robots are only concerned with the task they are handling at the moment and they cannot plan for future allocations. However, in the TA approach, robots have more global information and they are able to build a more efficient plan for future allocations by defining a schedule. The task assignment problem presented in this work can be coined as single-task robots, single-robot tasks, time-extended assignment problem (ST-SR-TA). We assume that the robots have decentralized decision making.

In this paper, we propose a distributed market-based assignment algorithm where robots are able to bid for tasks with the assumption that each robot can perform only one task and each task is assigned to only one robot at a time. The market-based approach provides a good trade-off between centralized² and distributed solutions. It eliminates the need for global information maintenance at a control station, while it provides more efficient solutions than other distributed approaches where only local information is provided. Also, a market-based approach can work even without knowing exactly the number of available robots in communication range, and this number could change dynamically without compromising the operation of whole system.

The remainder of the paper is organized as follows: In Section 2, we present the most important works dealing with the task assignment problem. In Section 3, both a basic Distributed Market-Based algorithm (DMB) and an Improved DMB (IDMB) will be presented. Section 4 presents a real-world experimental testing, and a simulation performance evaluation study. Finally, Section 5 concludes the paper and discusses future works.

2. Related Works

In the last years, different approaches were used to solve the task assignment problem. Although, this problem was solved optimally in a centralized manner using the Hungarian method², these solutions have the drawbacks of centralized systems such as the slow response to dynamic changes.

In the literature, several multi-robot assignment problems have been studied^{3,5}. In⁶, the authors addressed the linear assignment problem in the context of networked systems. The challenge is the lack of information due to the limited communication between agents. For that purpose, a distributed auction algorithm was proposed. Each agent bids for the task to which it wishes to be assigned. The algorithm considers local information thus each agent is required to keep pricing information and update it if necessary. The first step of the auction algorithm is to update the prices and highest bidders for all tasks and the second step is to check whether the price of the current assignment of each agent has been increased by other agents in the network or whether a larger-indexed agent has placed an equal bid. Simulation results show that the auction algorithm converges to an assignment that maximizes the total assignment benefit within a linear approximation of the optimal.

In⁷, the authors tackle the problem of task assignment for multiple agents. The problem is formulated as follows: Given n identical agents and m tasks where $m < n$, determine distributed control laws that split the agents into m groups of size n_k associated with each task $k = 1, \dots, m$. The solution is based on the market process where there is no centralized auctioneer and each agent can make bids. It is assumed that the agents have the knowledge of all tasks and the maximum number of agents that can be assigned to every task. The agents communicate together in order to compare bids and the information propagated on the network considering the availability of the requested tasks.

In⁸, the authors tackle the initial formation problem. It aims to decide which robot should go to each of the positions of the formation : in order to minimize a certain objective function. They proposed a new algorithm called the Robot and Task Mean Allocation Algorithm (RTMA). In the RTMA algorithm, the cost is considered as the difference between the distance from the robot to the task and the mean of distances from all the robots to that task. Therefore, the robot will win the task that is best for the team, not just for itself. Thus, the robots try to minimize the global cost. The auctioneer is responsible of calculating the mean of the distances for a certain task and sent it to the robot that won the task.

In⁹, the authors considered the multi-robot task assignment problem with set precedence constraints (SPC-MAP). The tasks are divided into a set of groups linked by precedence constraints. Each robot can exactly perform only one task from each group. After executing its tasks, the robot takes benefits. The SPC-MAP is formulated as follows: Given n_r robots, n_t tasks with the tasks divided into n_s disjoint subsets, maximize the total benefits of robot-task assignment with the set precedence constraints for tasks, such that, each task is performed by one robot, and each robot r_i performs exactly N_i tasks and at most one task from each subset. To solve the problem, the authors proposed an auction mechanism to bid for tasks. First of all, the robot must update its assignment information from the previous iteration. In the bidding part, the robot bids for tasks that are not already assigned to it. In the bidding process, the following constraints should be satisfied: (a) robot r_i is assigned to exactly N_i tasks; (b) robot r_i is assigned to at most one task in each subset. It is noted that each task is assigned to only one robot.

In¹⁰, the authors studied the multi-robot task assignment problem with task deadline constraints (DiMAP) which is an extension of the SPC-MAP problem⁹. As each robot has a limited battery life, it can only execute a limited number of tasks. The objective is to assign tasks to robots in a distributed manner in such a way the total cost is minimized and the task deadline constraints are achieved. Task assignment with deadlines is necessary for many application scenarios such as disaster recovery scenarios where robots must clear paths to find victims within some time. The problem is formulated as follows: given a set of tasks T , with each task $t_j \in T$ having a deadline d_j . Each task has to be done by only one robot and each robot can do one task at a time. The maximum number of tasks that robot r_i can do is N_i . Each robot r_i obtains a payoff a_{ij} for doing task t_j . The overall payoff is the sum of the individual robot payoffs. To solve the DiMAP problem, the authors proposed an auction-based distributed algorithm. The bidding algorithm for a robot r_i at an iteration is described as follows: firstly, the robot r_i receives the local task price from its neighbor robots, then updates its own price and calculates its task value. Secondly, the robot r_i selects a task set with task indices J . This step satisfies these conditions: (a) robot r_i is assigned to at most N_i tasks; (b) r_i is assigned to at most k tasks of all tasks with deadline no more than k . Thirdly, the robot r_i is assigned to task set, and updates the task price. At the end of the algorithm, each robot is assigned to its task. The algorithm was analyzed in terms of soundness (does the final solution satisfy all constraints?), completeness (Will the solution be found in a finite number of iterations and be feasible?) and optimality (How good is the solution?). The simulation results of a scenario of 20 robots, 100 tasks, and 5 tasks per robots show that there is a tradeoff between the solution quality and the convergence time. In¹¹, the multi-robot patrolling task was studied. An approach for monitoring robot performance in a patrolling task and dynamically reassigning tasks from those team members that perform poorly was considered. In the beginning, each robot is assigned to a partition of the environment and charged to visit every node in the partition. If the refresh time of any robot's assigned nodes exceeds a threshold, a central monitor process performs a task reassignment using a market-based auction algorithm. After announcing the candidate node, each robot add this node to its patrol partition and calculates the new bid. The node will be assigned to the robot that has the smallest refresh time. For simulation, three scenarios that differ from the number of robots (3, 5 and 8 robots) were run for 2 hours. The experiments were performed using a team of three TurtleBots in the same office environment. The authors compared their approach to a naive approach which does not consider individual robot performance. The results shows that the auction strategy produced better performance in terms of refresh time.

Paper¹² addressed the following Multi Robot Task Allocation problem. Consider a set n collaborative mobile robot, a set of n target locations and an $n * n$ matrix of the distances c_{ij} between the robots j and the task i . Each robot can perform only one task at a time and each task requires one robot to be accomplished. The problem consists in finding a feasible assignment of minimum total cost (distance) between the robots and the tasks. To solve this problem, authors proposed a distributed allocation algorithm based on the Hungarian Method. Each robot is assumed to make decision based on the information it has (e.g. distance to the target) and the information received from the other robots. The communication between robots is performed over a connected dynamic communication network. It is shown that the computational time performed by each robot (in the order of $O(n^2)$) results minor than the time required by the standard (centralized) Hungarian algorithm.

3. The Market-Based Algorithms for Task Assignment Problem

The task assignment problem can be defined as follows: Suppose that there are n identical robots, $\{r_1, r_2, \dots, r_n\}$ and n tasks to be accomplished by the robots, $\{t_1, t_2, \dots, t_n\}$, find the best and the most efficient policy to assign the n tasks to the n robots while maximizing the overall expected performance and minimizing the global cost defined as $\sum_{i=1}^n C(t_i, r_j)$ where $C(t_i, r_j)$ is the cost of executing task t_i by robot r_j . In this work, we consider the task of visiting target positions using a group of autonomous mobile robots.

The solution to the multi-robot assignment problem is based on a market process³. We define two main entities: the agents (i.e. robots) and the tasks. Each agent can play two roles: bidder (buyer) and auctioneer (seller). The auctioneer is the agent that announces offers of several items during an announcement phase. In our case, the items are the tasks. Then the bidders submit bids to the offers to the auctioneer. We assume that all robots are able to announce and bid for tasks. Every robot is characterized by a unique identifier, an initial position, velocity, field of view, and the task to which it is assigned. We assume that each robot has the knowledge about its own cost of the task, called robot cost vector, from the beginning. A robot does not know about the cost of other robots as it does not know about their positions.

3.1. The Basic Distributed Market-Based Algorithm (DMB)

Since the number of tasks is equal to the number of robots, we initially assign each robot to the task with the same ID . For example, the robots r_1 and r_2 are assigned respectively to the tasks t_1 and t_2 . Each robot agent plays the role of the auctioneer in order to announce its task and the rest of robots team are the bidders. The auction process starts with an announcement phase, where each robot sends an announcement message to provide its bid on the announced task. For example, the robot r_1 will announce the task t_1 for the other robots. We assume that all robots are in the same communication range. In the announcement phase, the robot sends its task ID , its own ID , and its cost of the task ($Mess_{ann}(t_i, r_j, C(t_i, r_j))$). Assuming that robots are in the same communication range, all announcement messages will be received by all robots. Possible collisions can be overcome by using contention resolution protocols. Each robot receives $(n - 1)$ announcement messages at a time, so each robot bids only for the tasks in which it has a cost lower than it was announced. Note that each robot sends bidding messages containing task ID , its own ID , its cost of the task. After receiving all bids, the auctioneer allocates the task to the robot with the best bid (lowest cost). After the assignment, a robot may win more than one task. In this case, it keeps the task with the best bid and sells the others. We assume that only the robots which didn't win tasks can bid for the new announced tasks. At the end, each task will be assigned to a robot. In this algorithm, each robot considers only its own benefit without taking into account the benefit of the whole system. Fig. 1.a and Fig. 1.b show the difference in cost between the Hungarian method and the DMB algorithm in the case where the number of robots and tasks is equal to 10. In most cases, the DMB algorithm produces non optimal solutions due to the fact that each robot works independently of the others.

3.2. The Improved Distributed Market-Based algorithm (IDMB)

The DMB consists of three main phases: announcement, bidding and assignment. In order to make improvement, we proposed a new phase called swap-Tasks. This step consists in swapping tasks between robots in order to minimize the total cost of the whole assignment. In this step, communication between robots is required. Each robot




Fig. 1. Difference in cost between the solutions obtained with the Hungarian algorithm the DMB algorithm and the IDMB algorithm

communicates with the other robots and asks for swapping.

With regard to communication complexity, each robot needs to send $(n - 1)$ messages where n is the total number of robots. Thus, the total number of messages sent by n robots is equal to $n * (n - 1)$. The swap function is repeated n times. It follows that the communication complexity is $O(n^3)$.

Assuming that the robot r_i tries to swap its task t_i with the robot r_j which is assigned to the task t_j . The robot r_i should compare the two following expressions:

$$C(t_i, r_i) + C(t_j, r_j) \quad (1)$$

$$C(t_i, r_j) + C(t_j, r_i) \quad (2)$$

where $C(t_i, r_i)$ is the cost to execute task t_i by robot r_i . If (1) < (2) then the robot r_i keeps its task t_i else robots r_i and r_j exchange their tasks. Fig. 1.b and Fig. 1.c represent an example to assign 10 robot to 10 tasks. It is observed that the IDMB algorithm gives better results than the DMB.

4. Experimentation and Simulations

4.1. Experimentation

We implemented the proposed distributed market-based approach on real robots to demonstrate the feasibility of the algorithm in real-world deployment. We used the Turtlebot¹ V2 robots as robotic platform and the open source Robot Operating System (ROS²) to implement the robot messaging, low-level control, and task assignment. Each robot relies on the default ROS navigation stack for navigation, localization and obstacle avoidance purposes.

The implementation mainly consists of three main components: (i.) a TurtleBot controller component, which consists of a set of modular classes that provide all functionalities needed to monitor sensor data of the robot and control its motors (e.g. sending the robot a goal). (ii.) a MRTA server, which is a UPD server that receives bidding and auctioning messages from neighbor robots and publish the receive message as a ROS topic, *mrtatopic*, for other ROS subscribers processing the messages, (iii.) a MRTA processing node, which is a ROS node that subscribes to the *mrtatopic* topic published by the MRTA server, and processes the received messages depending on their types, which are target announcement, target auctioning, target bidding, and target assignment.

The experiments were performed in the corridors of the College of Computer and Information Sciences of Prince Sultan University (refer to map in Fig. 2) with two Turtlebot robots. In the experiments, we tested a simple scenario with two robots $R1$ and $R2$ that select two targets locations $T1$ (location 26 in Fig. 2) and $T2$ (location 4 in Fig. 2), respectively, such that the optimal assignment is opposite to the initial assignment. After the execution of the DBM distributed algorithm, both robots reached the consensus about the optimal assignment such that $R2$ assigns its initial

¹ <http://www.turtlebot.com/>

² <http://www.ros.org/>



Fig. 2. ROS Map used for Experiments in Prince Sultan University

T2 to R1 and R1 assigns its initial target T1 to R2. A video demonstration of the experiments is available on this iroboapp project website¹³. We used the Euclidean distance for cost estimation. More accurate costs can be obtained using the total length of the path planned by the ROS global planner.

We also tested other simulated scenarios with five robots using our ROS implementation and confirmed the feasibility and correctness of assignment.

4.2. Simulation Study

In this section, we present a simulation study to evaluate the performance of the IDMB algorithm in scenarios without obstacles where the task cost is calculated as the Euclidean distance. The DMB and IDMB algorithms have been tested using a simple simulator implemented using MATLAB. The robots and tasks are scattered in an environment of 1000*1000 meters without obstacles. We tested the algorithms on different cases (i.e. 2, 5, 8, 10, 12, 15, 20 and 30 robots and tasks). We considered 30 different scenarios for each case. Each scenario is specified by the coordinates of a randomly chosen initial position of both robots and tasks. Each scenario is repeated 30 times (i.e. 30 runs for each scenario) and we recorded the global estimated cost. The optimal solution has been calculated using the Hungarian method². We also implemented the Robot and Task Mean Allocation algorithm (RTMA) proposed by Viguria and al.⁸. In order to demonstrate the efficiency of the IDMB algorithm, we considered the error in percentage in comparison with the optimal solution obtained with the Hungarian algorithm. The results are shown in Fig. 3. It can be observed that the solutions obtained with the IDMB are very close to the optimal (Fig. 4). For less than 10 robots and tasks, the algorithms (DMB, IDMB and RTMA) give good results (the error is less than 7.8% for DMB, 0.5% for IDMB and 6.8% for RTMA). Increasing the number of robots and tasks, the IDMB algorithm still produces near optimal solutions as shown in Fig. 3 and Fig. 4. The maximum error obtained with the IDMB algorithm in all cases did not exceed 2%. This means that the IDMB algorithm can generate the optimal solution in several simulations. Notice that in the case where the number of robots and tasks is 2, the IDMB algorithm always produces the optimal solution (error is 0%).

5. Conclusions

In this paper, we considered the task assignment problem in multi-robot systems. We presented two distributed market-based algorithms called DMB and IDMB. The IDMB algorithm consists in swapping tasks between robots in order to reduce the cost of the assignment. In our strategy, each robot can perform only one task. The IDMB algorithm produces near optimal solutions and in several cases it gives the optimal solution. We are currently working on extending the DMB approach to consider more complex assignment problems with heterogeneous robots. In addition, we aim at reducing the complexity of the communication overhead of the swap task by considering more intelligent heuristics instead of a pre-fixed number of message exchanges. Furthermore, we aim at extending the

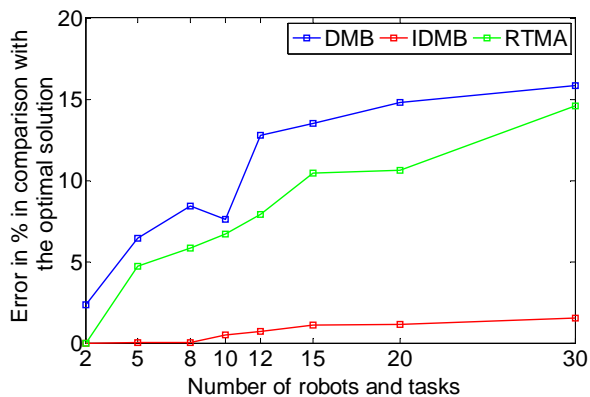


Fig. 3. Error in percentage in comparison with the optimal solution for the DMB, the IDMB and the RTMA algorithms.

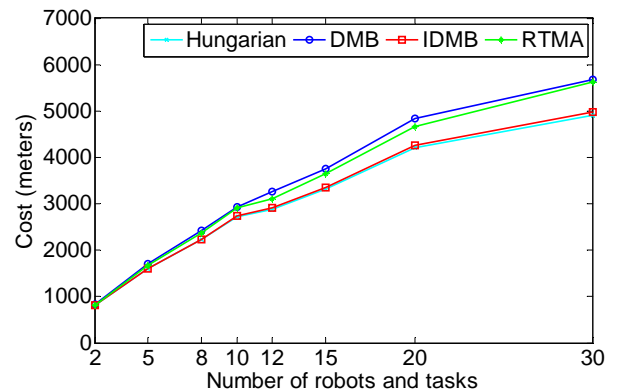


Fig. 4. Results of the estimated cost of the Hungarian, DMB, IDMB and RTMA algorithms over 30 simulations per case.

implementation and perform experiments with more robots to evaluate the algorithm performance in large-scale real deployment.

Acknowledgements

This work is supported by the iroboapp project “Design and Analysis of Intelligent Algorithms for Robotic Problems and Applications”¹³ under the grant of the National Plan for Sciences, Technology and Innovation (NPSTI), managed by the Science and Technology Unit of Al-Imam Mohamed bin Saud University and by King AbdulAziz Center for Science and Technology (KACST). This work is partially supported by Prince Sultan University.

References

1. Z. Yan, N. Jouandeau, A. A. Cherif, A survey and analysis of multi-robot coordination, *International Journal of Advanced Robotic Systems* 10.
2. H. W. Kuhn, The hungarian method for the assignment problem, *Naval Research Logistics Quarterly* 2 (1955) 83–97.
3. B. M. Dias, R. Zlot, N. Kalra, A. Stentz, Market-based multirobot coordination: A survey and analysis, *Proceedings of the IEEE* 94 (7) (2006) 1257–1270.
4. S. Trigui, A. Koubaa, M. Ben Jamaa, I. Chaari, K. Al-Shalfan, Coordination in a multi-robot surveillance application using wireless sensor networks, in: *16th IEEE Mediterranean Electrotechnical Conference (MELECON)*, 2012, 2012, pp. 989–992.
5. B. Gerkey, M. Mataric, A formal analysis and taxonomy of task allocation in multi-robot systems, *International Journal of Robotics Research* 23 (9) (2004) 939–954.
6. M. M. Zavlanos, L. Spesivtsev, G. J. Pappas, A distributed auction algorithm for the assignment problem, in: *47th IEEE Conference on Decision and Control*, 2008, p. 1212–1217.
7. N. Michael, M. M. Zavlanos, V. Kumar, G. J. Pappas, Distributed multi-robot task assignment and formation control, in: *IEEE International Conference on Robotics and Automation*, 2008, pp. 128–133.
8. A. Viguria, A. M. Howard, An integrated approach for achieving multirobot task formations, *Mechatronics, IEEE/ASME Transactions on* 14 (2) (2009) 176–186.
9. L. Luo, N. Chakraborty, K. Sycara, Multi-robot assignment algorithm for tasks with set precedence constraints, in: *IEEE International Conference on Robotics and Automation (ICRA)*, 2011, pp. 2526–2533.
10. L. Luo, N. Chakraborty, K. Sycara, Distributed algorithm design for multi-robot task assignment with deadlines for tasks, in: *IEEE International Conference on Robotics and Automation (ICRA)*, 2013, pp. 3007–3013.
11. C. Pippin, H. Christensen, L. Weiss, Performance based task assignment in multi-robot patrolling, in: *Proceedings of the 28th Annual ACM Symposium on Applied Computing*, ACM, 2013, pp. 70–76.
12. S. Giordani, M. Lujak, F. Martinelli, A distributed algorithm for the multi-robot task allocation problem, in: *Trends in Applied Intelligent Systems*, Springer, 2010, pp. 721–730.
13. iroboapp: Design and analysis of intelligent algorithms for robotic problems and applications, video demos, <http://www.iroboapp.org>. URL <http://www.iroboapp.org>