

Efficient Multi-AUV Cooperation using Semantic Knowledge Representation for Underwater Archaeology Missions

Nikolaos Tsiogkas*, Georgios Papadimitriou*, Zeyn Saigol*, David Lane*

*Ocean Systems Laboratory

Heriot Watt University, EH14 4AS Edinburgh, Scotland, UK

Email: nt95;d.m.lane@hw.ac.uk

Abstract—Advances in the fields of communication technology and software, electrical and mechanical engineering enable the replacement of a single robot by cooperative robotic team in highly demanding applications, such as search and rescue. A robotic team could perform better than a single robot, if certain challenges, such as action planning, coordination, and decision making, are successfully tackled. One key factor for the successful performance of a robotic team is the multi-robot task allocation. Specifically, the challenge is to define which robot executes which task, considering an efficient solution for the successful completion of the complex mission. The task allocation could be even more challenging when real-world communication constraints and uncertainties are presented, such as limited bandwidth, high latency and high packet loss. In the current study, we attempt to resolve the issue of a cooperative robotic team under communication constraints. To reach this goal, the use of a distributed world model for multi-robot task allocation is proposed. This ontology based distributed world model is capable of successfully handling to a great extent the aforementioned communications limitations, thus allowing successful mission execution even under harsh communication conditions. An efficient centralised task allocation mechanism, using k-means clustering, is described, and its performance is compared to a greedy centralised task allocation method. Experimental simulation results indicate that the efficient method performs better on average than the greedy one, without extra time requirements.

I. INTRODUCTION

Recent advances in robotics have enabled the use of teams of robots to solve complex real-life tasks. Using robots means that tasks can be executed in a more safe and cost effective manner than the past, where humans had to perform these tasks. One particularly interesting and challenging area for autonomous robotics is the underwater domain. Working underwater is dangerous and challenging for humans, and the use of autonomous underwater vehicles (AUVs) can lower both the risk and the cost of underwater operations. Moreover it can enable operations that were not possible for humans.

This paper focuses on underwater archaeological exploration, as described in [1]. The mission examined is one where a heterogeneous team of AUVs is required to search a specified area for archaeological artefacts, inspect them and classify them. At first AUVs with mapping and searching capabilities (Search-AUVs or SAUVs) have to map the area and find specific points of interest that have to be inspected. Then another type of AUVs with inspection capabilities (Inspection-AUVs or IAUVs) moves to these points of interest and inspects

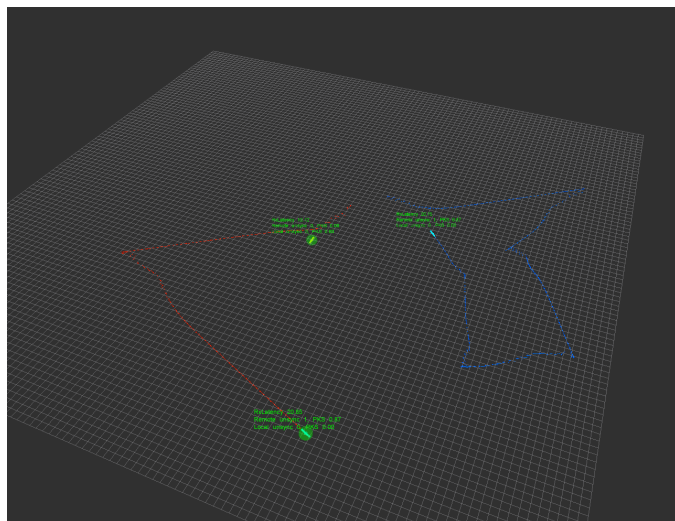


Fig. 1. Mission execution instance.

the targets. The targets are then classified as archaeological artefacts or not.

For the successful mission accomplishment using a multi-robot team, a structured approach to coordinate the execution is required. The mission is decomposed to smaller tasks, able to be executed by a single robot. Then the robots are required to distribute and execute these tasks. The mission is considered successful when all the defined tasks are successfully executed. To make a decision upon tasks the robots must have a certain level of knowledge about the world and a way to communicate with other robots so as to coordinate, both in terms of motion and action. There has been extensive literature regarding the task allocation problem. However, it is usually assumed that the communication among the members of the robotic team is error-free. Real-world underwater communication can be hindered because of high latency, high error rates and low bandwidth, as discussed in [2], [3] and [4]. This implies that with imperfect communication the information sharing and task allocation among robots can be a quite challenging and demanding problem.

To overcome the communication challenges, the use of a distributed world model, as extensively described in [5] and [6], is proposed. The distributed world model has the responsibility to handle all the information sharing needs

among members of the robot team. Moreover, to increase the planning capabilities of the robots, the ontology based world model can be utilised as described in [7]. Additionally, the distributed world model can be used for the multi-robot team coordination, since the task allocation can be viewed as extra information that has to be shared among the robots.

Regarding the task allocation an efficient centralised task allocation method is utilised, implemented and compared with a greedy centralised allocation method. In the centralised task allocation one robot acts as a leader of the team and decides for the other robots. In the greedy task allocation scheme, tasks are allocated based on a greedy nearest neighbour approach. The proposed method involves an efficient approach that utilizes k-means clustering and a travelling salesman algorithm which try to minimize the distance travelled by the robots.

The rest of the paper is organized as follows: In Section II, a review of previous work on task allocation is presented. In Section III background information regarding methods used to solve the problem is presented. In Section IV one challenging application of the multi-robot team is described. Section V includes the efficient task allocation scheme that is implemented to tackle the described challenges. In Section VI the evaluation of the methods is performed. In Section VII useful conclusions are drawn. Finally, Section VIII details the future research directions.

II. PREVIOUS WORK

Research in the field of multi-robot task allocation has been quite active in recent years. A key study formalizing the multi-robot task allocation problem is presented in [8]. A newer study regarding the taxonomy of the multi-robot task allocation is presented in [9]. In [10] the multi-robot problem is viewed as a sub-problem of the general distributed intelligence problem. Additionally, in the same work a categorisation based on the interactions among the team members is attempted.

In the literature, there are many solution approaches to the multi-robot task allocation problem. In [11] a centralised approach for the multi-robot exploration problem is presented. There frontier search is used to explore an unknown environment based on the maximisation of a utility function. This approach only considers robots that are able to communicate with each other. An extension of this approach is presented in [12] where the task execution utility is calculated based on the information gain of the whole path the robot executes and not merely from the destination. In [13] a centralised behaviour-based architecture is presented. In this approach, the robots have basic capabilities and they are matched with certain tasks based on their expertise. This architecture can derive an optimal allocation given enough time. In [14] a k-means clustering method is used to balance the load on a multi-robot team.

In robotics another important aspect is planning the task execution. Ontologies have been proved to be able to facilitate the planning process. In [15] a reactive planning system is described that is able to re-plan the mission to accommodate a component failure, thanks to semantic knowledge encoded in the ontology. In [16] semantic maps are used to increase the planning capabilities of the robot by using semantic information and improve the planning efficiency. In [7] ontologies are

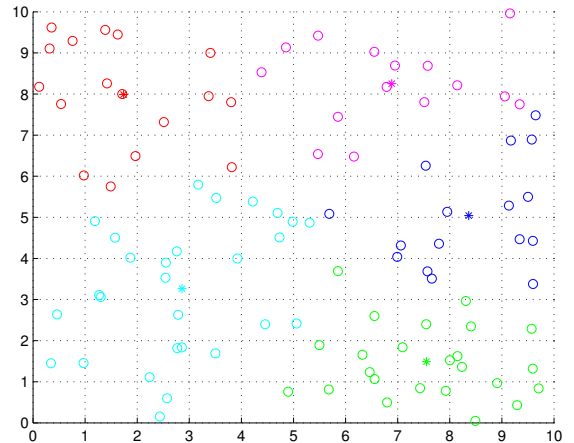


Fig. 2. Clustering of 100 2-d observations in five clusters using the k-means algorithm.

used to increase the robustness of plan execution and to aid the planning effort in cases of plan failures.

III. BACKGROUND

In this section some background information regarding the methods used in Section V are presented.

A. k-means clustering

The k-means clustering algorithm, presented in [17], is a method of partitioning a set of observations into clusters using some distance metric. Given a set of n observations $X \in \mathbb{R}^d$, the algorithm tries to partition these observations into k sets with $k \leq n$. It is achieved by choosing k centres C that minimize the potential function,

$$\phi = \sum_{x \in X} \min_{c \in C} \|x - c\|^2$$

The algorithm is simple and can be seen in the following listing:

- 1) Randomly choose k centres $C = (c_1, \dots, c_k)$.
- 2) For each $i \in (1, \dots, k)$, calculate the points of X that are closer to c_i than c_j for all $i \neq j$, and assign them to cluster C_i .
- 3) For each $i \in (1, \dots, k)$, set c_i as the center of mass of points in C_i .
- 4) Repeat steps 2 and 3 until C not changes.

In this paper the k-means++ algorithm is used as it is presented in [18]. This algorithm proposes a method to choose the initial k centres and is shown to improve the speed and the accuracy of the k-means algorithm. A clustering of 100 2-d observations using the aforementioned k-means algorithm can be seen in figure 2.

B. Travelling Salesman problem

The travelling salesman problem is a famous problem which tries to find the shortest path required to traverse a set

of n cities. It can be formulated as an integer programming problem as shown below [19].

$$\begin{aligned}
\min \quad & \sum_{i=0}^n \sum_{j \neq i, j=0}^n c_{ij} x_{ij} \\
0 \leq x_{ij} \leq 1 \quad & i, j = 0, \dots, n \\
u_i \in \mathbf{Z} \quad & i = 0, \dots, n \\
\sum_{i=0, i \neq j}^n x_{ij} = 1 \quad & j = 0, \dots, n \\
\sum_{j=0, j \neq i}^n x_{ij} = 1 \quad & i = 0, \dots, n \\
u_i - u_j + n x_{ij} \leq n - 1 \quad & 1 \leq i \neq j \leq n
\end{aligned}$$

Where x_{ij} denotes a path that goes from city i to city j , c_{ij} is the cost of travelling from city i to city j and n is the number of the cities. u_i and u_j are artificial variables used to ensure that only a single tour covers all the cities.

In [20] are provided methods for solving hard and large scale TSP problems.

C. Distributed world model

Knowledge sharing among members of a robotic team is crucial for the successful completion of missions where adaptive behaviour is required. The distributed world model, presented in [5] and in [6], provides a fully decentralised knowledge-base system that utilises acoustic communications to share knowledge among the members of the robotic team. The distributed world model is capable of operating in challenging communications conditions (high latency, low bandwidth, high error rates) which are typical in the underwater domain [4]. Knowledge is represented using OWL ontology language. The ontology based representation allows for a structured representation of data and provides with methods to inference and reason over the data.

In more detail, the world model runs on each robot as a standalone server process. The other subsystems of the robot can access the world model server through a client-side library. The world model server also decides which information should be forwarded to the other robots, based on their information needs, and shares these updates in a push-based manner. This method allows each robot to focus on mission execution and planning, as the world model service handles the information exchange among the team members. The higher level planning uses a C++ client to query the ontology for information that it needs. Additionally, the decentralized nature enables the seamless operation of a vehicle even if it is temporarily isolated from the other team members. The distributed world model architecture is depicted in figure 3.

The ontology form that is applied is inspired from [21]. In this work it is proposed that each variable as an instance of the `Attribute` concept. The `Attribute` can have one or more timestamped `PropertyValues`, each representing the value of the variable at a specific time. This form allows the synchronisation of data among the vehicles, especially in cases where the communication is hindered by high latency and high packet loss. Additionally, this ontology form allows the storage

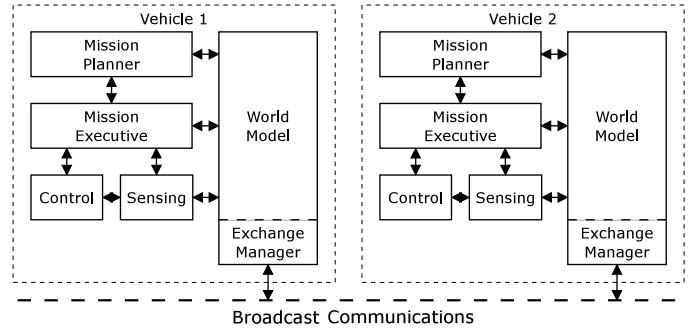


Fig. 3. The distributed world model architecture.

of the history of a variable. This aspect is significant as, the robot can, for example, reason over the past and reach some conclusion regarding the current mission status. In addition the data logging allows for a full offline mission repeatability, which makes possible a post-mission analysis.

IV. MISSION DESCRIPTION

In this section a mission relevant for the multi-robot team performance is described.

The mission involves a target detection and classification. In this type of missions the robots discover targets in an area and then they classify them as a certain class of a real-world object. In the current work the robotic team is assembled with two types of robots: (a) robots capable of searching for targets (SAUVs) and (b) robots capable of inspecting and classifying the detected targets (IAUVs). The mission goal is to search a specific area for targets to be inspected and then inspect and classify those targets. The mission is completed when all of the discovered targets are classified and the classification information has propagated to the whole team.

To complete the mission a centralised task allocation scheme is implemented. In the centralised task allocation there is one SAUV which acts as a leader and assigns tasks to the other AUVs of the robotic team. As the mission progresses and new targets are discovered the leader AUV creates new tasks and assigns them to the other robots. Whenever a new task appears, the leader AUV selects the appropriate sub-team of robots, based on their capabilities and their fitness to perform the task, and finds the lowest cost robot and the task is then assigned to that robot. The robots then find out the tasks they have to execute from the world model and whenever a new task is received, the robot tries to insert it into its current execution plan by adapting its original plan.

The detection and classification mission was selected to be investigated as it can be applied to various real-world scenarios. This type of mission can be adapted to a search and rescue mission, or to an archaeological search and recovery mission or to a mine counter-measures mission. All the above scenarios share a common action strategy: The robots should locate certain targets, then classify them and finally act appropriately according to some plan. To sum up, the selected mission is believed to be quite challenging, with rich research opportunity and highly potential, as it could be used for a wide range of applications.

V. TASK ALLOCATION

In this section the efficient task allocation scheme is going to be described. A detailed description of the proposed method is given. For sake of completeness the greedy task allocation method will also be described.

A. Efficient task allocation method

For the efficient task allocation and execution the techniques described in section III are utilised. As targets are discovered the leader AUV creates tasks for the inspection of those targets. Then the leader AUV periodically tries to assign these tasks to inspection AUVs in an efficient manner. To do that it takes advantage of the spatial distribution of targets by trying to cluster the targets using the k-means algorithm. In each assignment iteration the leader AUV finds which are the remaining targets to be classified. It then creates clusters of targets, and then assigns one cluster to each robot by calculating the minimum cost assignment based on the distances of the robots from the cluster centres.

The periodical task assignment allows to cope with the dynamic discovery of targets. As the mission progresses and new targets are discovered the previously calculated clusters will be invalid. By creating new clusters we ensure that targets which are spatially close are visited by the same robot thus minimizing the distance that has to be travelled.

The inspection AUVs on the other hand periodically check for tasks that they have to execute. They then order the targets they have to visit by solving the equivalent travelling salesman problem. To solve the travelling salesman problem two strategies are followed depending on the number of targets to be visited. If the number of targets is low (i.e. ≤ 8) an exact solution is calculated by calculating all the possible solutions and getting the one that minimizes the distance to be travelled. If the number of targets is high then a simulated annealing technique is used as described in [22]. The simulated annealing method can provide close to optimal solutions with a small computational cost.

Task execution is performed by always performing the first task on the ordered list (FIFO scheduling). It should be noted that as the tasks to be executed are updated periodically, the task execution may be interrupted in order to execute a more favourable task based on the new data received. This can help the efficiency of task execution as the targets are dynamically discovered.

B. Greedy task allocation method

In the greedy task allocation a simpler strategy is followed. As new targets are discovered the leader robot creates tasks and tries to assign them to the other robots. The task assignment is more straightforward. The leader robot, knowing the targets that are already assigned to the other robots, tries to find which robot could insert this target to its existing tour with the minimum cost. For the assignment, firstly the tour of each robot is computed by using a greedy nearest neighbour approach based on the already assigned targets. Then new tours are computed including the newly discovered target. The differences in cost are calculated by subtracting the cost of the old tour from the new tour. Finally the target is assigned in

the robot that has the minimum cost increase for inspecting the extra target.

The tasks in turn are executed by the inspection AUVs by greedily selecting the nearest target that is assigned to them and proceeding with inspection and classification of the object.

VI. EXPERIMENTAL RESULTS

To evaluate the performance of the two aforementioned methods simulation experiments were conducted using the communications and navigation simulator developed by the Ocean Systems Laboratory. The simulation scenario is based on the archaeological survey mission as defined in [1]. In this scenario AUVs search, inspect and classify objects of archaeological interest in a predefined area. The mission requires the vehicles to allocate tasks to search, inspect and classify a certain number of targets. For this scenario there was two AUV types: (a) a single search AUV that detected randomly generated targets and (b) a number of inspection AUVs that inspected and classified the targets. The mission was considered to be successful when all the targets were classified. Experiments take into account the uncertainties and restrictions in underwater communications (latency, low bandwidth, errors). For the evaluation of both methods, a packet size of 512 bytes and a packet travel time of 2 seconds were chosen as they are typical values for underwater communications. The search AUV detected targets at a rate of one minute per target. Both methods are compared in the view of mission execution time and energy efficiency, where energy efficiency is measured by the total path length that has to be travelled by the AUVs. Relative performance of the two techniques will be compared as the quality of the communications link (packet error rate) and the number of inspection targets are varied.

The first experiment was designed to test the time and cost efficiency of each method in different values of packet error rates. The vehicles had to inspect ten different sets of ten random targets each. The packet error rate was incremented from 0 to 0.8 with a step of 0.2. In figure 4 the average mission completion time is depicted. It can be observed that in terms of time efficiency both methods perform equally. As it is expected the time to complete a mission rises as the packet error rate becomes higher, and thus the communication becomes harder. As shown in the figure the time grows with at least a quadratic rate to the packet error rate.

In figure 5 the average total distance travelled by the two vehicles is shown. The cost efficiency, which is measured by the average total distance the inspection vehicles had to travel, is generally in favour of the efficient method. As it can be seen from the standard deviation analysis the efficient method produces more consistent results.

The second experiment was designed to test the scaling of the two methods in larger target numbers. The vehicles had to inspect ten different sets of twenty random targets each. The packet error rate was set to 0. The results, shown in tables I and II are in accordance with the previous results. The efficient method still outperforms on average the greedy one while performing in almost the same time.

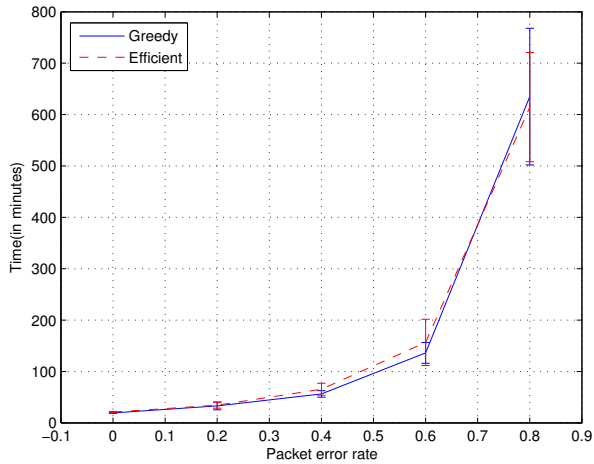


Fig. 4. Average mission completion time for different values of packet error rate.

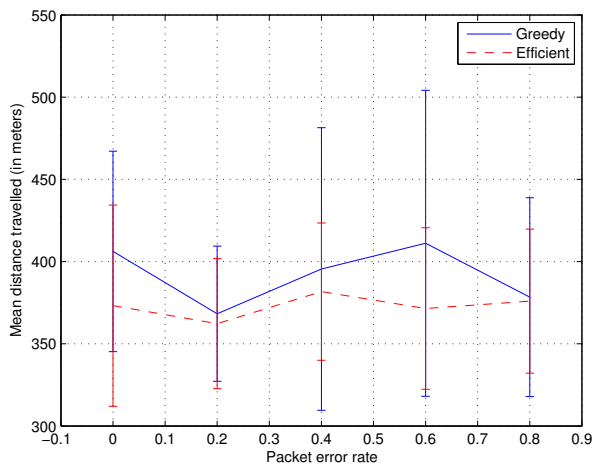


Fig. 5. Average total distance travelled for different values of packet error rate.

VII. CONCLUSIONS

In this work the problem of multi-robot task allocation was studied under high latency and unreliable communications. Two task allocation methods were implemented and experimentally tested. On average the efficient method performed better in terms of cost, while not requiring more time to complete a mission. The use of the distributed world model has enabled the successful mission execution even under harsh communication conditions with packet error rates reaching 80%.

VIII. FUTURE WORK

In this work significant effort was paid to successfully implementing and testing the above efficient centralised task allocation for the investigated mission. The conclusions from the current work indicate that there is space for further improvements. It has been noted that the performance regarding the mission execution time greatly decreases as the packet

TABLE I. AVERAGE MISSION COMPLETION TIME FOR 20 TARGETS.

	Average mission completion time	Standard deviation
Efficient Method	35.54	1.13
Greedy method	33.39	1.44

TABLE II. AVERAGE TOTAL DISTANCE TRAVELLED FOR 20 TARGETS.

	Average total distance travelled	Standard deviation
Efficient Method	457.75	70.52
Greedy method	490.02	109.04

error rate goes up. This is probably caused due to the high communication demands of the centralised solution. It would be interesting to investigate if the communication aspect of the task allocation would drop by implementing a decentralised task allocation scheme, where all robots act as equal peers, and each robot decides which targets to inspect by itself only by knowing the position of other robots and the position of targets. Moreover, the decentralised task allocation scheme would provide robustness to the mission execution, as the failure of a single robot would not prohibit the mission execution, compared to the failure of the leader robot of a centralised team.

ACKNOWLEDGMENT

The authors would like to thank all the members of the Ocean Systems Laboratory at Heriot-Watt University. The research leading to these results has received funding by the European Union Seventh Framework Programme FP7 ARROWS - under grant agreement No. 308724.

REFERENCES

- [1] ARROWS. Online - <http://www.arrowsproject.eu/>, 2012. ARchaeological Robot systems for the Worlds Seas.
- [2] M. Roth, D. Vail, and M. Veloso, "A real-time world model for multi-robot teams with high-latency communication," in *Intelligent Robots and Systems, 2003. (IROS 2003). Proceedings. 2003 IEEE/RSJ International Conference on*, vol. 3, pp. 2494–2499 vol.3, Oct 2003.
- [3] I. Vasilescu, K. Kotay, D. Rus, M. Dunbabin, and P. Corke, "Data collection, storage, and retrieval with an underwater sensor network," in *Proceedings of the 3rd International Conference on Embedded Networked Sensor Systems, SenSys '05*, (New York, NY, USA), pp. 154–165, ACM, 2005.
- [4] Z. Saigol, G. Frost, N. Tsiogkas, F. Maurelli, D. Lane, A. Bourque, and B. Nguyen, "Facilitating cooperative AUV missions: experimental results with an acoustic knowledge-sharing framework," in *Proceedings of IEEE-MTS Oceans'13, San Diego, USA*, 2013.
- [5] J. J. Cartwright and D. M. Lane, "Design and analysis of a decentralised world model service for collaborating autonomous underwater vehicles," (*submitted in RAS*), 2012.
- [6] F. Maurelli, Z. Saigol, J. Cartwright, D. Lane, A. Bourque, and B. Nguyen, "Tdma-based exchange policies for multi-robot communication of world information," in *Proceedings of IFAC MCMC 2012 - Manoeuvring and Control of Marine Craft*, 2012.
- [7] G. Papadimitriou and D. Lane, "Semantic based knowledge representation and adaptive mission planning for MCM missions using AUVs," in *Proceedings of IEEE-MTS Oceans'14, Taipei, Republic of China*, 2014.
- [8] B. P. Gerkey and M. J. Mataric, "A formal analysis and taxonomy of task allocation in multi-robot systems," *The International Journal of Robotics Research*, vol. 23, no. 9, pp. 939–954, 2004.
- [9] G. A. Korsah, A. Stentz, and M. B. Dias, "A comprehensive taxonomy for multi-robot task allocation," *The International Journal of Robotics Research*, vol. 32, no. 12, pp. 1495–1512, 2013.
- [10] L. E. Parker, "Distributed intelligence: Overview of the field and its application in multi-robot systems," *Journal of Physical Agents*, vol. 2, no. 1, pp. 5–14, 2008.

- [11] W. Burgard, M. Moors, and C. Stachniss, "Coordinated multi-robot exploration," *IEEE Transactions on Robotics*, vol. 21, no. 3, pp. 376–386, 2005.
- [12] J. Butzke and M. Likhachev, "Planning for multi-robot exploration with multiple objective utility functions," in *Intelligent Robots and Systems (IROS), 2011 IEEE/RSJ International Conference on*, pp. 3254–3259, IEEE, 2011.
- [13] F. Tang and L. E. Parker, "Asymtre: Automated synthesis of multi-robot task solutions through software reconfiguration," in *Robotics and Automation, 2005. ICRA 2005. Proceedings of the 2005 IEEE International Conference on*, pp. 1501–1508, IEEE, 2005.
- [14] M. Elango, S. Nachiappan, and M. K. Tiwari, "Balancing task allocation in multi-robot systems using k -means clustering and auction based mechanisms," *Expert Systems With Applications*, vol. 38, no. 6, pp. 6486–6491, 2011.
- [15] E. Miguelanez, P. Patron, K. E. Brown, Y. R. Petillot, and D. M. Lane, "Semantic knowledge-based framework to improve the situation awareness of autonomous underwater vehicles," *IEEE Transactions on Knowledge and Data Engineering*, vol. 23, pp. 759–773, May 2011.
- [16] C. Galindo, J.-A. Fernández-Madrigo, J. González, and A. Saffiotti, "Robot task planning using semantic maps," *Robotics and Autonomous Systems*, vol. 56, no. 11, pp. 955 – 966, 2008. Semantic Knowledge in Robotics.
- [17] J. MacQueen, "Some methods for classification and analysis of multivariate observations," 1967.
- [18] D. Arthur and S. Vassilvitskii, "k-means++: The advantages of careful seeding," in *Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*, pp. 1027–1035, Society for Industrial and Applied Mathematics, 2007.
- [19] C. E. Miller, A. W. Tucker, and R. A. Zemlin, "Integer programming formulation of traveling salesman problems," *Journal of the ACM (JACM)*, vol. 7, no. 4, pp. 326–329, 1960.
- [20] C. Rego, D. Gamboa, F. Glover, and C. Osterman, "Traveling salesman problem heuristics: leading methods, implementations and latest advances," *European Journal of Operational Research*, vol. 211, no. 3, pp. 427–441, 2011.
- [21] C. J. Mathews, M. M. Kokar, and K. Baclawski, "A core ontology for situation awareness," in *Proceedings of the Sixth International Conference of Information Fusion, 2003*, vol. 1, pp. 545–552, 2003.
- [22] C. C. Skiścim and B. L. Golden, "Optimization by simulated annealing: A preliminary computational study for the tsp," in *Proceedings of the 15th conference on Winter Simulation-Volume 2*, pp. 523–535, IEEE Press, 1983.