Multi-Robot Perception and Action: World Modeling and Task Allocation

Francesco Riccio, Maria T. Lázaro, Guglielmo Gemignani, Daniele Nardi
Sapienza University of Rome
Email: {riccio, mtlazaro, gemignani, nardi}@dis.uniroma1.it

***Abstract*—Multi-Robot Systems (MRS) is one research area where Artificial Intelligence and robotic techniques can be efficiently integrated. In this paper, we attempt to highlight some common aspects of the robotic and AI literature on multi-robot systems by surveying the most recent works in such an area. In particular, we focus on works that deal with the problem of coordinating a team of autonomous robots perceiving the world and acting in it to carry out a common task. By surveying these works, we attempt to give a new perspective on the problem of multi-robot world modeling and distributed multi-robot coordination.**

I. INTRODUCTION

In recent years, intelligent robots started to enter our lives moving out from research labs. Indeed, their employment reaches a broad range of fields, finding several applications in human-dangerous environments exploration, surveillance, health care assistance or domestic environments and entertainment. The Multi-Robot Systems (MRSs) research area of study constitutes one of the cross-points between AI and Robotics. In fact, it aims at developing systems of multiple robots that are capable of coexisting in various scenarios and of achieving common goals.

Generally, Multi-robot systems can be categorized in three main groups based on their structure. In particular, they can be designed as centralized, decentralized or distributed. Centralized systems are based on a central unit that allows for communication, coordination and information storage for the robots. This approach to MRSs is an efficient choice in terms of consistency and cohesion of the system, but it penalizes real applications due to the common drawbacks of single-point-of-failure frameworks. Conversely, in a decentralized setting each robot has a complete and redundant representation of the world, built through the exchange of broadcasted messages among the robots. Finally, in a distributed approach, robots take their decisions based on their partial internal representation of the world updated with the information obtained from their teammates.

In this survey, we will focus on distributed multi-robot systems. In particular, we are interested in the problem of allowing a group of autonomous robots to perceive and act in an environment that might or might not be known a-priori. As an example, let's consider the typical problem faced in disaster scenarios. In order to perform tasks in a coordinated manner in such scenarios, a suitable representation of the operational environment needs to be given to the robots. However, the robots must build a new model of the environment since it often changes during disasters. To this end, both position estimation and mapping problems must be solved concurrently, leading to the problem of Simultaneous Localization and Mapping (SLAM), extensively studied in mobile robotics. Once the world has been modeled, the robots have to coordinate in order to effectively carry out the given assignments. To this end, several approaches to multi-robot coordination have been proposed in literature.

Abstracting away from the example scenario, the problem can be generalized as shown in Fig. 1. This figure sketches the data flow of a multi-robot system mainly composed by two fundamental parts: the distributed world model (DWM) and the task allocation modules (Φ). The DWM module is in charge of representing and sharing information about the perceived world, and therefore, to build a reliable world representation for reasoning about the operational environment. The module labeled Φ is instead responsible of coordinating the robots within such an environment by assigning each task t_j to the best robot r_i.

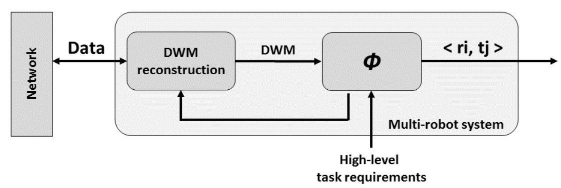

Fig. 1. Data flow sketch for a general multi-robot system able to perceive an act in a given environment. The distributed world model module is in charge of representing the operational environment. Instead, the module Φ assigns to each task t_j a robot r_i based on the current representation of the world.

In the next section we will survey the most recent approaches to multi-robot world modeling and coordination. Then, we will conclude the paper discussing the common aspects of the surveyed related works to give a new perspective on the two areas analyzed.

II. MULTI-ROBOT SYSTEMS

As previously stated, in order to allow a team of robots to carry out a task in a coordinated manner, two fundamental elements are necessary: a suitable representation of the world and a reliable coordination technique. There are approaches that do not require an explicit representation of the world but they are outside the scope of this work. Because of this fact, in this section we will first review several world modeling approaches proposed in the literature. Next, we will discuss multiple recent works in multi-robot coordination, referring the reader to previous surveys for a more detailed analysis of prior works.

A. Multi-robot World Modeling

World modelling is a broad research problem that can be addressed from a variety of perspectives. In general, robots use their own perception to acquire a representation of the environment suitable for the task it has to accomplish. For example, in a robot soccer domain, the model of the environment has to include static and dynamic tracked objects where the cooperation of the team is crucial due to the robots' narrow field of view [1]. Cooperative localization approaches [2], [3], [4] have also interesting advantages in the absence of an *a priori* model of the environment where, by relying upon inter-robot measurements and the correlations established between them, an update on the position estimation of one robot improves the estimations of the rest of the team.

In the literature, different kind of map representations have been proposed, usually grouped as metric, semantic [5] and topological maps [6]. In this section, we focus on how an accurate metric representation can be obtained by a team of robots, referring the reader to other surveys for more details on semantic and topological mapping [7], [8], where approaches to topological map merging [9], [10], [11] could be also suitable for multi-robot map reconstruction.

Due to the different sources of errors present in the sensors used during the map acquisition process, the SLAM problem is usually approached using a probabilistic formulation. Over the years, different techniques have been proposed in the framework of Bayesian filtering, such as Extended Kalman Filters (EKFs), Extended Information Filters (EIFs) and Particle Filters leading to different kind of maps, from feature-based stochastic maps to dense maps like occupancy grid maps. We refer the reader to [8], [12] for a complete review of these approaches applied to single-robot systems. More recent approaches to the SLAM problem have their origin in the early work by Lu and Milios [13] where the SLAM problem is represented as a graph of spatial relationships between poses (i.e., pose-graph) of the complete robot trajectory, optimally estimated using a maximum likelihood criterion. Current successful implementations like iSAM [14] or g^2o [15] allow to solve the SLAM problem in real time.

The extension of the aforemetioned single-robot approaches to multi-robot systems has recently interested the SLAM community. In fact, these approaches present multiple advantages, such as an increased accuracy and robustness or the possibility of cooperation to cover larger environments in less time. However, by increasing the number of robots, the complexity of the problem increases leading to new important issues related to data fusion, inter-robot data association (i.e., map matching) or communication constraints.

During the data fusion process it is important to take into account the problem of *double-counting* or *data incest*. This is a common problem in decentralized or distributed networks and is exemplified in the following way: consider a robot A which computes its position $\mathbf{x}_A$ using its own observations $\mathbf{z}_A$. Then, it communicates its position to a robot B, which uses this received information and its measurements $\mathbf{z}_B$ to improve its position $\mathbf{x}_B$ which sends later to robot A. At this point, robot A can not treat $\mathbf{x}_B$ directly because it already contains information from A and it would be integrating its own information twice, producing inconsistent estimates. To solve this problem, several techniques were proposed: [16] uses a Channel Filter to substract the information that has already sent through the network, [17] keeps track of the origins of measurements sent over the network and [18] presents the Covariance Intersection technique which provides consistent solutions, when correlations between estimates are unknown.

Data fusion can be easily achieved when global or relative positions between robots are known; however this is a strong assumption only satisfied in an idealistic scenario. Therefore, a considerable amount of SLAM approaches have considered robots that lack any prior knowledge of their initial locations or their relative positions with respect to their teammates. In this context, Thrun and Liu [19] exploit the inherent decoupling property of EIFs to propose a decentralized SLAM approach. This technique allows to identify overlapping maps acquired by different robots and recover the joint map, even with unknown initial robot positions or landmark correspondences between maps.

Zhou [20] presents a multi-robot EKF-based SLAM approach, where robots acquire relative pose measurements to determine the transformations between pairs of maps. Common landmarks are also identified through an improved Nearest Neighbor method and by exploiting the Mahalanobis distance to validate their correspondences. This process reduces the overall map complexity and increases its accuracy. Once two maps from different robots have been merged they keep working using this fused map.

Howards [21] proposes a multi-robot SLAM approach with particle filters in occupancy grid maps, where inter-robot relations are determined through mutual detection using retro-reflective markers mounted on the robots. The maps acquired by different robots are fused into a single centralized map once the encounter between them takes place. An improved distributed version using a Rao-Blackwellized particle filter is implemented in [22]. Each time the robots meet they share all the measurements gathered since the last communication, which is incorporated into the other robots' map accounting also for the uncertainty of their relative transformations.

Addressing the multi-robot SLAM problem with the more recent nonlinear optimization approaches has one additional interesting advantage: the complete individual robot trajectories are also recomputed during the optimization process to minimize the joint overall map error. In this context, Andersson and Nygårds [23] present a Collaborative Smoothing and Mapping (C-SAM) algorithm to build a joint graph-based map with features using a team of robots. It is a centralized solution to the problem in which graphs from different robots are related through *base nodes* and rendezvous measurements, which allow to represent them in a common reference frame.

An extension of iSAM for multi-robot pose-graph SLAM is presented in [24]. Robots build incrementally their own maps which are related through relative measurements obtained during robot encounters. The introduction of *anchor nodes* allows to express all maps in a common global frame while each robot keeps working in its own coordinate system.

All these approaches provide solutions to the problem of modeling the environment with a team of robots, however they consider few communications aspects. Normally, two robots can share information only when they are within a certain communication range, using a connection with a limited bandwidth. Such limitations are often ignored, often assuming that the robots can either share their complete knowledge of the map or the complete history of measurements acquired.

To this end, in [25] the authors propose an extended Smoothing and Mapping approach called Decentralized Data Fusion (DDF-SAM) using factor graphs with landmarks. In this approach, each robot maintains its own local map whose landmarks are marginalized to create a summarized graph which is sent to other robots. This compressed map together with those received from other robots are fused in a *neighborhood* map which is updated as new information is received. However this approach does not allow for a feedback from the neighborhood map to their local maps to prevent the double-counting problem. This issue was addressed a subsequent work [26] where only one single structure is needed to manage both local and neighbor maps and the *anti-factor* is introduced to subtract replicated information. The data association is based on a triangulation algorithm that provides matching between maps.

In [27], the authors propose a multi-robot graph-based SLAM approach which makes use of *condensed measurements* [28] and a communication protocol to efficiently exchange information among the robots. These condensed measurements are a summarized version of each robot's graph containing only the *relevant* information for other robots, reducing the communication overload and the size of the optimization problem each robot has to solve. Furthermore, these condensed measurements do not require any special treatment to be included in the optimization back-end (g^2o in this case) and are computed only with each robot's own information to avoid the multiple integration of information. Relations between pair of maps are determined by a robust data association algorithm each time robots are in the same communication range.

B. Distributed Task Allocation

Distributed Task Allocation (DTA) is the problem of coordinating multiple cooperative units with a common goal in various scenarios. Such a problem introduces different issues to be tackled and different constraints to be satisfied. In a MRS scenario, in fact, we can have different configurations. For instance, not only the number of robots operating within the same environment is important, but also if the robots share the same goal or if the team is divided into sub-teams with different goals. In this latter setting, the teams have to consider conflicting behaviors and have to model other teams as an active part of the environment or even as adversarial agents.

Additionally, apart from the kind of scenario that we are facing, an MRS can have different features depending on the kind of coordination criterion itself. In particular, we can have *explicit* or *implicit* coordination. The main difference in these two configuration is the way in which the robots exchange local information. The former case implies an explicit sharing of information, and agreement of the robots about the best mapping between robots and task, while the latter allows the team of robot to coordinate even if the world model of each robot is not known by all the units. For example, a DTA approach based on *utility estimations* belongs to this class.

Despite the kind of coordination requirements and the operating scenario the problem of multiple robot coordination and task allocation can be generally expressed as the problem of assign a given set of tasks $T = \{\tau_1, \ldots, \tau_M\}$ to a set of acting robots $R = \{r_1, \ldots, r_N\}$, or in other words, the formalization of a function

$$\Phi : R \longrightarrow T \quad (1)$$

which maps the robots into a set of available tasks that generally is called the *task-space*. In characterizing this function, manifold issues need to be addressed, typically depending on the task to be satisfied, the environment, the context in which the robots are acting, and most importantly, the characteristics of the robots itself.

The mapping function Φ has to be expressed by considering the surrounding world. Generally, the environment needs to be properly modeled in a suitable representation in order to enable the robots to take informed decisions. Therefore, in most of the recent works, that try to solve the problem of of distributed coordination and task assignment, the world model is usually taken for granted.

Originally, multi-robot coordination has been studied for engineering reasons, where multi-robot systems were designed to improve the effectiveness of a robotic system. In a later moment, a significant development of this research topic has stemmed from studies on biological systems or complex models arising in cognitive science and economics. This large amount of work has been the subject of multiple survey papers that provide interesting characterizations of this research area. For example, Cao *et al.* [29] propose multiple dimensions for characterizing multi-robot coordination. Instead, Dudeck *et al.* [30] give a classification that is more focused on the communication and computation aspects. Stone and Veloso [31] present an introduction to the field of Multi-Agent Systems and MRS, along with a conceptual framework to organize the possible systems, while the research topics in the MRS field are discussed by Parker [32]. Finally, Farinelli *et al.* [33] classify multiple approaches proposed in literature, specifically focusing on the coordination aspects of MRS.

Building upon these surveys, in this section we focus on recent works that have approached the problem of task allocation in a decentralized fashion. For example, Wicke *et al.* [34] have recently proposed an approach to multi-agent task allocation inspired by bounty hunters and bail bondsmen. In this approach, while a task remains uncompleted, its bounty gradually rises, making it more and more desirable to pursue. Unlike auctions, this model does not assume rationality in agents bids. An alternative approach has been recently presented by Johnson *et al.* [35]. In this work the authors present a decentralized assignment approach for communicating agents by means of an asynchronous channel. The authors modify the Asynchronous Consensus-Based Bundle Algorithm (ACBBA) in order to take into account more real-time implementations. More specifically, first a new metric of convergence for ACBBA is identified and then, the system is enabled to asynchronously re-plan its task execution.

Okamoto *et al.* [36] focus instead on a large-scale assignment problem. They address the problem of task allocation by enhancing the Distributed Constraint Optimization

(DCOP) [37] approach for large-scale problem, by modifying the task acceptance rules for each robot. Their solution allows the robot to exploit the available local knowledge and adapt their acceptance criteria. DCOP has also been recently considered by Wu and Jennings [38]. In this work, the authors extend standard DCOP models to consider uncertain task rewards, where the outcome of completing a task depends on its current state, which is randomly drawn from unknown distributions. In order to do so, the authors propose a decentralized algorithm that incorporates Max-Sum with iterative constraint generation to solve the problem. Max-Sum is exploited also by Cerquides *et al.* [39]. In this work, the authors consider inter-team coordination in order to achieve a desirable cooperation between multiple specialized teams. With this approach the authors show how a team of firefighters and police forces can coordinate more effectively in polynomial time by using Tractable High Order Potential. Analogously, Corrêa [40] considers the problem of coordinating a team of robots to form groups of agents to solve disaster tasks. In this work, the author discusses an algorithm that creates partitions of agents in a tree-structure factor graph. Also Pennisi *et al.* [41] have recently faced the problem of distributed task assignment to address *multi-robot surveillance*. In this work, the authors perform a greedy task allocation routine based on utility estimation computed according to Euclidean distances between areas to be covered and the actual positions of the robots.

Alternatively, a distributed coordination problem can be seen as a *Decentralized Partially Observable Markov Decision Process (Dec-POMDP)*. In this framework, a team of robots has the goal of *learning* the best collective policy in order to maximize a common reward. For example, Eker *et al.* [42] use evolution strategies to generate policies and learn joint actions for two robots in a grid-world. Matignon *et al.* [43], instead, use *distributed value function (DVF)* to formalize a Dec-POMDP as several MPDs, one for each robot, which are locally solved in order to succeed in multi-robot exploration. Capitan *et al.* [44] use decentralized POMDP to multi-robot surveillance and tracking using UAVs. The authors employ auctions in order to flexibly coordinate the agents' individual policies formalized as a POMDPs.

Another commonly used approach for multi-robot coordination is Dynamic Task Assignment. This standard approach however is not able to handle unpredictable tasks in dynamic environments. To tackle this issue, recently Luo *et al.* [45] have described an iterative greedy auction algorithm in order to approximate an optimal task allocation solution. However, the tasks to be assigned are not known beforehand, but they may be discovered *online*. For this reason, the authors run an auction episode for each arising task, assigning it to the robot with the maximum payoff. Additionally, to allow a fully decentralized approach, the auctioneer is substituted with a maximum consensus technique among the entire network of robots. Luo *et al.* [46] have subsequently extended their previous work by proposing an auction-based task allocation algorithm. In this work, the authors first propose a centralized approach that is able to compute the best mapping between tasks and robots, and then a decentralized one which guarantees provable-suboptimal solutions. In their distributed solution, the authors locally maximize the objective function for each robot. One of the key contributions of this paper consists of dropping the assumption of independence among a given group of tasks, optimizing them altogether.

Finally, a few years ago Stone *et al.* [47] have approached the task assignment problem in a completely different but still affine perspective. In this work, the authors address the multi-robot cooperation problem without relying on *a priori* protocols. Their final goal consists of building a working team of robots that only shares the knowledge of the operational environment. In other words, their goal consists of building a single agent able to cooperate with other unknown agents that are not necessarily programmed by the same team. These other robots might also feature different capabilities and may also not share the same knowledge of the current state of the world. Drawing from this initial study, other similar works have been presented. In particular, Liemhetcharat and Veloso have represented the synergy among robots as a connected weighted graph. This graph has been then used to evaluate the ability of each pair of robots to work together [48]. Instead, Barret and Stone have adopted pre-learned policies to quickly react to teammate actions [49].

III. Comparison and Conclusion

When coordinating multiple autonomous robots, it is often necessary to rely onto a distributed system. Specifically, this is a common choice when the task to be accomplished requires robustness, scalability and performance. Depending on the task, the coordination system can significantly change, and the used approaches may vary accordingly. In the previous section, we surveyed and highlighted two different aspects of a Multi-Robot System: *World modeling*, where the team of robots are goal-driven and the aim consists of efficiently reconstructing the surrounding world while perceiving it, and *Task assignment*, where the main problem consists of deciding how to decompose tasks and allocate resources.

As we presented in the previous section, there are three main issues one has to solve in multi-robot *world modeling*. Firstly, the robots must establish relations between pair of maps obtained from different robots. Some approaches [20], [21] obtain these relations through direct robot-to-robot measurements, although sometimes this constrains robot movements as their identifiers must be visible to the others. Inter-robot measurements can also be obtained indirectly by comparing the robots local maps which involves solving a data association problem, as done in [19], [25], [27].

Once the transformations between maps have been obtained, they can be merged into a centralized map which allows to have a common global vision of the environment as done in [20], [21], [23]. These systems may require constant communications between robots or a communication infrastructure to connect to a central server. On the other hand, in distributed approaches [22], [26], [27] robots work with their own maps and only incorporate information from others, if they eventually meet.

Conversely, if we consider *distributed task allocation* approaches, the key issue is the formalization of the function mapping a set of robots into a set of tasks. In fact, the mapping function Φ must take into account the goal of the team, the tasks that are available to the team of robots and, most importantly, it has to encode the features and capabilities of each unit in the team.

Indeed, depending on the robots capabilities the mapping function changes significantly. For instance, if we consider a team of *equally skilled* robots, and the problem to solve is well formalized, then we can employ DCOP based approaches as in [36], [38]. In a DCOP formulation each robot contributes to the optimization of a common objective function and governs a set of variables (usually one) constituting the objective function. Obviously, a direct application of a DCOP is the optimization of a team of robots, but in general it can be adapted to every decentralized problem solving formulation.

Conversely, if the team of robots has heterogeneous units with different skills and capabilities, then the mapping function Φ needs to be reshaped accordingly. In these terms, marked-based approaches, or more in general *distributed task assignment (DTA) approaches* present promising results. In this setting, each robot performs a *self evaluation* with respect to the available tasks and propose a possible assignment which maximizes its payoff. For example, these techniques may rely on auctions [45] or utility estimations [41]. Utility estimation, for example, represent a possibility of expressing how good a robot is with respect to a given task. Thus, if there are M tasks, each robot computes a vector of M dimensions of utility values, one for each task. According to this formulation, the evaluation is locally performed by the robot itself, and then, shared among the entire team. This allows to easily implement a fully distributed system in which robots heterogeneity is encrypted into the utility estimations and can be collectively evaluated.

Multi-robot World Modeling and *Distributed Task Allocation* are generally addressed separately. This splits the MRS problem into two "easier" parts, however we foresee future works into complete operational systems where both parts are integrated would benefit the whole system. For example, the robots could build a better map of the environment and more efficiently if, by using some of the above mentioned techniques for task allocation, the team can decide which area each robot is going to cover and establish strategic meeting points where the data association would be more easily solved.

In fact, we could implement complete systems that guarantee reliable performance in every type of multi-robot system. In this perspective, different approaches are arising where the robots first *perceive* the world and update their distributed model, and then, effectively *act* in it in order to achieve a given goal. For example, De Hoog *et al.* [50] implement a dynamic task assignment to select the best explorer in a a frontier-based exploration task. Instead, in [41], the authors select the best robot to cover a particular area thought utility estimations in order to succeed in a surveillance task. Similarly, in [51] the authors use utility functions to select the best task association depending on the reconstructed world and the current context of the environment.

To summarize, in this paper we have attempted to highlight common aspects of the works on multi-robot systems both in the robotic and AI literature. In particular, we have focussed on works that deal with the problem of coordinating a team of autonomous robots that need to perceive the world and act in it to achieve a common goal. By surveying these works, we have underlined common aspects of both categories of works, sketching a general multi-robot system able to carry out such a task. We hope that this contribution will help stimulating more works that at the same time try to integrate techniques for distributed world modeling with approaches for multi-robot coordination.

References

[1] B. Coltin, S. Liemhetcharat, C. Mericli, J. Tay, and M. Veloso, "Multi-humanoid world modeling in standard platform robot soccer," in *IEEE-RAS Int. Conf. on Humanoid Robots*, Dec 2010, pp. 424–429.

[2] S. Roumeliotis and G. Bekey, "Distributed multi-robot localization," *IEEE Transactions on Robotics and Automation*, vol. 18, no. 5, pp. 781–795, Oct. 2002.

[3] E. Nerurkar, S. Roumeliotis, and A. Martinelli, "Distributed maximum a posteriori estimation for multi-robot cooperative localization," in *IEEE Int. Conf. on Robotics and Automation*, Kobe, Japan, May 12-17 2009, pp. 1402–1409.

[4] T. Bailey, M. Bryson, H. Mu, J. Vial, L. McCalman, and H. Durrant-Whyte, "Decentralised cooperative localisation for heterogeneous teams of mobile robots," in *IEEE Int. Conf. on Robotics and Automation*, May 9-13 2011, pp. 2859–2865.

[5] A. Nüchter and J. Hertzberg, "Towards semantic maps for mobile robots," *Robotics and Autonomous Systems*, vol. 56, no. 11, pp. 915–926, 2008.

[6] S. Thrun, "Learning metric-topological maps for indoor mobile robot navigation," *Artificial Intelligence*, vol. 99, no. 1, pp. 21–71, 1998.

[7] I. Kostavelis and A. Gasteratos, "Semantic mapping for mobile robotics tasks: A survey," *Robotics and Autonomous Systems*, vol. 66, pp. 86–103, 2015.

[8] S. Thrun, "Exploring artificial intelligence in the new millennium," G. Lakemeyer and B. Nebel, Eds. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 2003, ch. Robotic Mapping: A Survey, pp. 1–35.

[9] W. H. Huang and K. R. Beevers, "Topological map merging," *The International Journal of Robotics Research*, vol. 24, no. 8, pp. 601–613, 2005.

[10] S. Saeedi, L. Paull, M. Trentini, M. Seto, and H. Li, "Group mapping: A topological approach to map merging for multiple robots," *Robotics Automation Magazine, IEEE*, vol. 21, no. 2, pp. 60–72, June 2014.

[11] T. Bonanni, G. Grisetti, and L. Iocchi, "Merging partially consistent maps," in *Simulation, Modeling, and Programming for Autonomous Robots*, ser. Lecture Notes in Computer Science, D. Brugali, J. Broenink, T. Kroeger, and B. MacDonald, Eds. Springer International Publishing, 2014, vol. 8810, pp. 352–363.

[12] J. Aulinas, Y. Petillot, J. Salvi, and X. Lladó, "The slam problem: a survey," vol. 184, pp. 363–371, Oct. 22-24 2008.

[13] F. Lu and E. Milios, "Globally consistent range scan alignment for environment mapping," *Autonomous robots*, vol. 4, no. 4, pp. 333–349, 1997.

[14] M. Kaess, A. Ranganathan, and F. Dellaert, "iSAM: Incremental smoothing and mapping," *IEEE Transactions on Robotics*, vol. 24, no. 6, pp. 1365–1378, Dec. 2008.

[15] R. Kümmerle, G. Grisetti, H. Strasdat, K. Konolige, and W. Burgard, "g2o: A general framework for graph optimization," in *IEEE Int. Conf. on Robotics and Automation*, May 9-13 2011, pp. 3607–3613.

[16] S. Grime and H. Durrant-Whyte, "Data fusion in decentralized sensor networks," *Control Engineering Practice*, vol. 2, no. 5, pp. 849–863, 1994.

[17] A. Bahr, M. Walter, and J. Leonard, "Consistent cooperative localization," in *IEEE Int. Conf. on Robotics and Automation*, Kobe, Japan, May 12-17 2009.

[18] S. J. Julier and J. K. Uhlmann, "Using covariance intersection for {SLAM}," *Robotics and Autonomous Systems*, vol. 55, no. 1, pp. 3 – 20, 2007.

[19] S. Thrun and Y. Liu, "Multi-robot SLAM with sparse extended information filers," in *The 11th International Symposium Robotics Research.* Springer, 2005, pp. 254–266.

[20] X. S. Zhou and S. Roumeliotis, "Multi-robot slam with unknown initial correspondence: The robot rendezvous case," in *2006 IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, 2006, pp. 1785–1792.

[21] A. Howard, "Multi-robot simultaneous localization and mapping using particle filters," *The Int. Journal of Robotics Research*, vol. 25, no. 12, pp. 1243–1256, 2006.

[22] L. Carlone, M. Ng, J. Du, B. Bona, and M. Indri, "Rao-blackwellized particle filters multi robot SLAM with unknown initial correspondences and limited communication," in *IEEE Int. Conf. on Robotics and Automation*, Anchorage, Alaska, May 3-8 2010, pp. 243–249.

[23] L. Andersson and J. Nygards, "C-SAM: Multi-robot SLAM using square root information smoothing," in *IEEE Int. Conf. on Robotics and Automation*, Pasadena, CA, USA, May 19-23 2008, pp. 2798–2805.

[24] B. Kim, M. Kaess, L. Fletcher, J. Leonard, A. Bachrach, N. Roy, and S. Teller, "Multiple relative pose graphs for robust cooperative mapping," in *IEEE Int. Conf. on Robotics and Automation*, Anchorage, Alaska, May 3-8 2010, pp. 3185–3192.

[25] A. Cunningham, K. Wurm, W. Burgard, and F. Dellaert, "Fully distributed scalable smoothing and mapping with robust multi-robot data association," in *IEEE Int. Conf. on Robotics and Automation*, St. Paul, Minnesota, USA, May 14-18 2012, pp. 1093–1100.

[26] A. Cunningham, V. Indelman, and F. Dellaert, "DDF-SAM 2.0: Consistent distributed smoothing and mapping," in *IEEE Int. Conf. on Robotics and Automation*, 2013, pp. 5220–5227.

[27] M. T. Lázaro, L. M. Paz, P. Piniés, J. A. Castellanos, and G. Grisetti, "Multi-robot SLAM using condensed measurements," in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, Nov 3-8 2013.

[28] G. Grisetti, R. Kümmerle, and K. Ni, "Robust optimization of factor graphs by using condensed measurements," in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, Vilamoura, Portugal, Oct. 7-12 2012, pp. 581–588.

[29] Y. U. Cao, A. S. Fukunaga, and A. Kahng, "Cooperative mobile robotics: Antecedents and directions," *Autonomous robots*, vol. 4, no. 1, pp. 7–27, 1997.

[30] G. Dudek, M. R. Jenkin, E. Milios, and D. Wilkes, "A taxonomy for multi-agent robotics," *Autonomous Robots*, vol. 3, no. 4, pp. 375–397, 1996.

[31] P. Stone and M. Veloso, "Robot teams: From diversity to polymorphism," *A Survey of Multiagent and Multirobot System*, 2002.

[32] L. E. Parker, "Current state of the art in distributed autonomous mobile robotics," in *Distributed Autonomous Robotic Systems 4*. Springer, 2000, pp. 3–12.

[33] A. Farinelli, L. Iocchi, and D. Nardi, "Multirobot systems: a classification focused on coordination," *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics*, vol. 34, no. 5, pp. 2015–2028, 2004.

[34] D. Wicke, D. Freelan, and S. Luke, "Bounty hunters and multiagent task allocation," in *Int. Conf. on Autonomous Agents and Multiagent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 2015, pp. 387–394.

[35] L. B. Johnson, S. S. Ponda, H.-L. Choi, and J. P. How, "Asynchronous decentralized task allocation for dynamic environments," in *Proc. of the AIAA Infotech@ Aerospace Conference*, vol. 2, no. 1, 2011, pp. 2–2.

[36] S. Okamoto, N. Brooks, S. Owens, K. Sycara, and P. Scerri, "Allocating spatially distributed tasks in large, dynamic robot teams," in *Int. Conf. on Autonomous Agents and Multiagent Systems*. Int. Foundation for Autonomous Agents and Multiagent Systems, 2011, pp. 1245–1246.

[37] P. Scerri, A. Farinelli, S. Okamoto, and M. Tambe, "Allocating tasks in extreme teams," in *Int. Joint Conference on Autonomous Agents and Multiagent Systems*. ACM, 2005, pp. 727–734.

[38] F. Wu and N. R. Jennings, "Regret-based multi-agent coordination with uncertain task rewards," *arXiv preprint arXiv:1309.1973*, 2013.

[39] M. Pujol-Gonzalez, J. Cerquides, A. Farinelli, P. Meseguer, and J. A. Rodriguez-Aguilar, "Efficient inter-team task allocation in robocup rescue," in *Int. Conf. on Autonomous Agents and Multiagent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 2015, pp. 413–421.

[40] A. Corrêa, "Distributed team formation in urban disaster environments," in *2014 IEEE Symposium on Intelligent Agents (IA)*. IEEE, 2014, pp. 57–64.

[41] A. Pennisi, F. Previtali, C. Gennari, D. D. Bloisi, L. Iocchi, F. Ficarola, A. Vitaletti, and D. Nardi, "Multi-robot surveillance through a distributed sensor network," in *Cooperative Robots and Sensor Networks 2015*. Springer, 2015, pp. 77–98.

[42] B. Eker, E. Özkucur, C. Meriçli, T. Meriçli, and H. L. Akin, "A finite horizon dec-pomdp approach to multi-robot task learning," in *Application of Information and Communication Technologies (AICT), 2011 5th International Conference on*. IEEE, 2011, pp. 1–5.

[43] L. Matignon, L. Jeanpierre, and A.-I. Mouaddib, "Coordinated multi-robot exploration under communication constraints using decentralized markov decision processes." in *AAAI*, 2012.

[44] J. Capitan, M. T. Spaan, L. Merino, and A. Ollero, "Decentralized multi-robot cooperation with auctioned pomdps," *The International Journal of Robotics Research*, vol. 32, no. 6, pp. 650–671, 2013.

[45] L. Luo, N. Chakraborty, and K. Sycara, "Competitive analysis of repeated greedy auction algorithm for online multi-robot task assignment," in *IEEE Int. Conf. on Robotics and Automation*. IEEE, 2012, pp. 4792–4799.

[46] ——, "Provably-good distributed algorithm for constrained multi-robot task assignment for grouped tasks," *IEEE Transactions on Robotics*, vol. 31, no. 1, pp. 19–30, 2015.

[47] P. Stone, G. A. Kaminka, S. Kraus, J. S. Rosenschein, *et al.*, "Ad hoc autonomous agent teams: Collaboration without pre-coordination." in *AAAI*, 2010.

[48] S. Liemhetcharat and M. Veloso, "Weighted synergy graphs for effective team formation with heterogeneous ad hoc agents," *Artificial Intelligence*, vol. 208, pp. 41–65, 2014.

[49] S. Barrett and P. Stone, "Cooperating with unknown teammates in complex domains: A robot soccer case study of ad hoc teamwork," in *Proc. of the 29th AAAI Conference on Artificial Intelligence*, January 2015.

[50] J. De Hoog, S. Cameron, and A. Visser, "Autonomous multi-robot exploration in communication-limited environments," in *Proc. of the 11th Conference Towards Autonomous Robotic Systems*. University of Plymouth, School of Computing and Mathematics, 2010.

[51] F. Riccio, E. Borzi, G. Gemignani, and D. Nardi, "Context-based coordination for a multi-robot soccer team," *RoboCup Symposium*, 2015.