

Multi-Robot Systems: A classification focused on coordination

Alessandro Farinelli, Luca Iocchi, Daniele Nardi

Dipartimento di Informatica e Sistemistica Università “La Sapienza”

Via Salaria 113 00198, Roma, Italy.

E-mail: {farinelli,iocchi,nardi}@dis.uniroma1.it

Abstract—Multi-Robot Systems (MRS) are, nowadays, an important research area within Robotics and Artificial Intelligence and a growing number of systems has been recently presented in the literature. Since application domains and tasks that are faced by MRS are of increasing complexity, the ability of the robots to cooperate can be regarded as a fundamental feature. In this paper, we present a survey of the recent work in the area by specifically examining the forms of cooperation and coordination realized in the MRS. In particular, we propose a new taxonomy for classification of the approaches to coordination in MRS and we describe some systems, which we consider representative in our taxonomy. We finally discuss the outcomes of our analysis and try to highlight future trends of the research on MRS.

I. INTRODUCTION

Multi-Robot Systems (MRS) have been proposed in the last decade in a variety of settings and frameworks, pursuing different research goals, and successfully applied in many application domains. Special attention has been given to MRS developed to operate in dynamic environments, where uncertainty and unforeseen changes can happen due to the presence of robots and other agents that are external to the MRS itself.

Generally speaking, a Multi-Robot System can be characterized as a set of robots operating in the same environment. However, robotic systems may range from simple sensors, acquiring and processing data, to complex human-like machines, able to interact with the environment in fairly complex ways. Moreover, it is not easy to give a definition of the level of autonomy that is required for a robot in order to be considered an entity acting in the environment, as opposed to a simple machine that provides services to the operator (a printer or a even a light switch). While we discuss several different settings of Multi-Robot Systems, we primarily focus on fairly complex mobile platforms, equipped with sophisticated sensors and actuators, able to execute complex tasks. We can further characterize the subset of MRS, that is addressed in the present work, by considering three main aspects: (i) the rationale for the design of the MRS, (ii) the basic functionalities and technologies (both hardware and software) used in the MRS development and (iii) the tasks that the robots should perform and the intended application domains. In the sequel we explain and discuss each of these characteristics in further details.

A significant body of work on MRS has been originated from motivations that are essentially of engineering nature, where MRS are designed and realized in order to improve

the effectiveness of a robotic system. From an engineering stand point, the MRS can improve the effectiveness of a robotic system either from the viewpoint of the performance in accomplishing certain tasks, or in the robustness and reliability of the system, which can be increased by modularization [4]. In fact, MRS are useful not only when the robots can accomplish different functions, but also when they have the same capabilities [31]. Moreover, even when a single robot can achieve the given task, the possibility of deploying a team of robots can improve the performance of the overall system. Another significant development of MRS stems from the studies on biological systems or complex models arising in cognitive science and economics (see for example [19]). In this work we take an engineering perspective, although we also look at a few biologically inspired approaches.

Technological improvements both in the hardware and in the associated software are two of the key reasons beyond the growing interest in MRS [8], [81]. The increased availability of complex sensor devices and robotic platforms in the research laboratories favored their development and customization, resulting in robots equipped with reliable and effective hardware that improves their basic capabilities (laser range finders, cameras, infrared sensors, robotic arms, gripping devices etc.). In addition, the software techniques developed for the robotic applications take advantage of the hardware improvements and provide complex and reliable solutions for the basic tasks that a robot should be able to perform, while acting in real world environments: localization, path planning, object transportation, object recognition and tracking, etc. Although several problems faced in single robot applications are not solved in a general and effective way, under specific assumptions, some of them can be tackled reliably. Moreover, the effectiveness of a solution to a single robot task could be, in some cases, improved using coordination among several robotic agents [86], [37], [83]. Therefore, the study and development of MRS applications is particularly relevant and significant at this stage.

MRS are well suited for several application domains, which require high complexity coordinated tasks to be performed. Examples of these applications can be found both in industrial environments (see for example [97]), and in hostile and dangerous environments, such as space [47], military [82], [63], and rescue operations [77]. In order to be used in such applications, the MRS should fulfill several requirements: long term autonomy and robustness are two of the main features

required to a MRS. In fact, space and security applications may require robots to work without human interaction for a long time. Therefore, the system should be able to handle unpredictable situations and execute complex tasks while maintaining the power level of its own members. Moreover, in order to reduce the development cost of the overall application, the proposed solutions often involve heterogeneous team members. Heterogeneity in the team composition often entails complex coordination strategies, in order to exploit the characteristics of all the team members and to obtain efficient solutions. A significant boost to the work on MRS has recently been given also by Robotics competitions, such as AAI robotic contests¹ and RoboCup² [61]. In fact, the design and the realization of MRS is regarded as one of the major scientific challenges and robotic contests are extremely useful for comparing and analyzing different strategies and techniques by providing a common test-bed for experiments. Moreover, these competitions offer new challenges in the design of MRS: for example in the RoboCup soccer domain, as compared with other domains for MRS, the environment is highly dynamic and includes an opponent team.

Even considering only the subset of MRS we are addressing, a common framework for the technical solutions that are being developed for MRS is difficult to identify. Moreover, a MRS cannot be simply regarded as a generalization of the single robot case and the proposed approaches need to be precisely characterized in terms of assumptions about the environment and in terms of the internal system organization [69], [70]. Neither a MRS can be simply considered as a special case of a Multi-Agent System (MAS), because of the issues arising when dealing with a physical environment, such as uncertainty and incompleteness on information acquired from the environment. In fact, the need to cope with the acquisition of knowledge from a real environment, makes the experimental evaluation of MRS much more challenging. In addition, the forms of cooperation used in MRS need to take into account the uncertainty, the limitations, and the mistakes arising from the processing of sensor information. The large amount of work in MRS has been the subject of a few survey papers that provide interesting characterizations and perspectives of the research in the area. In [19] several dimensions for characterizing a MRS are proposed, while in [31] a classification of MRS, that is more focused on the communication and computation aspects, is presented. In [110] an introduction to the field of MAS and MRS is presented along with a conceptual framework to organize the possible systems, while the research topics in the MRS field are discussed in [81].

Although cooperation and coordination are central in many works on MRS and they are addressed in the above cited surveys, a detailed analysis that specifically looks at these aspects can be interesting in at least two respects. First, the complexity of the systems and of the application domains requires more and more sophisticated forms of coordination. Consequently, a systematic analysis of the proposals that

appeared in the recent literature can help the designer of MRS to choose the approach to coordination that is best suited for the application at hand. Second, focusing on coordination it is possible to further analyze the relationship between MAS and MRS: highlighting differences and similarities between the two can lead to new insights that stem from the cross fertilization of the two fields.

The aim of the present work, which is an extended and revised version of [36], is to address the most recent developments of MRS by classifying the proposed approaches, specifically focusing on the coordination aspects of the MRS. To this end, we present a new *taxonomy* for classifying MRS approaches to coordination and a set of system dimensions that address those aspects of the system organization that influence coordination: team size, team composition, communication and architecture. Moreover, we identify the tasks that are faced by MRS in various application domains. Based on the above sketched framework, we provide a precise and fine classification of a large body of recent works in the field and discuss the trends of the research on MRS.

The remainder of the paper is organized as follows: in Section II we present the taxonomy we propose for classifying the approaches to coordination in MRS; in Section III we give an overview of the typical tasks and application domains for MRS; in Section IV we describe several works in MRS and classify them according to our taxonomy; finally, in Section V we discuss the trends emerging in this research field.

II. A TAXONOMY FOR MRS

As already remarked there are several motivations for addressing the design of MRS and, consequently, the work in the area can be classified from several points of view. Our main motivation is the study and evaluation of the ability to take advantage of coordination to improve the system performance. Therefore, the classification we propose is focused on the coordination aspects and thus inspired by the relationship with the field of multi-agent systems (MAS).

In order to provide a classification of recent works on MRS, we first propose a new taxonomy and then put it in perspective with respect to other classifications and recent surveys on MRS. Following the literature, we use the term *dimension* to refer to specific features that are grouped together in the classification.

The taxonomy that we propose for classifying the works on MRS is characterized by two groups of dimensions: *Coordination Dimensions* and *System Dimensions*. Generally speaking, the former aim at characterizing the type of coordination that is achieved in the MRS, while the latter include the system features that influence team development. More specifically, for the *Coordination Dimensions* of our taxonomy, a hierarchical structure is given in Figure II.

For a suitable classification of the works it is important to clearly define the dimensions that are used. In the following of this section we define our classification dimensions and discuss the main differences that arise when we consider MRS instead of MAS.

¹See for example <http://www.andrew.cmu.edu/~tlauwers/robohost/>

²See <http://www.robocup.org/>

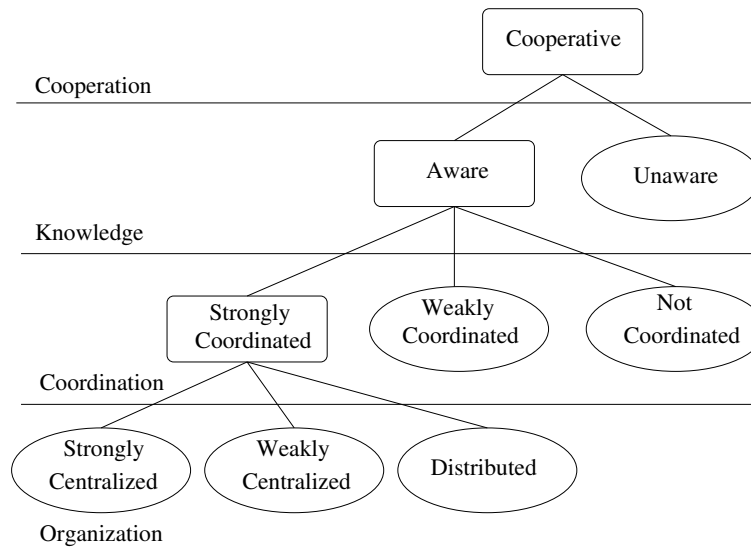


Fig. 1. MRS Taxonomy

<i>Coordination Dimensions</i>	<i>System Dimensions</i>
<i>Cooperation</i>	<i>Communication</i>
<i>Knowledge</i>	<i>Team Composition</i>
<i>Coordination</i>	<i>System Architecture</i>
<i>Organization</i>	<i>Team Size</i>

TABLE I
CLASSIFICATION DIMENSIONS

Cooperation Level: The first level is concerned with the ability of the system to cooperate in order to accomplish a specific task. At the *Cooperation Level* we distinguish cooperative systems from not cooperative ones. A cooperative system is composed of “robots that operate together to perform some global task” [80]. In this work we are interested only in cooperative MRS. Therefore, in the following, the term MRS will refer to a team of cooperative robots. This notion of cooperation is very similar to the ones used for MAS; however, in MAS, cooperation is often compared and merged with competition, which, up to now, has received little attention in the recent works on MRS. A remarkable exception is the work [28] where a free market base approach is used to coordinate multiple robots (see section IV for further details).

Knowledge Level: The second level of the proposed hierarchical structure is concerned with the knowledge that each robot in the team has about its team mates. *Aware* robots have some kind of knowledge of their team mates, while *Unaware* robots act without any knowledge of the other robots in the system. The interest in cooperating unaware MRS (that is not as common in MAS) is motivated from an engineering point of view by the simplicity of such systems, with respect to aware ones. Observe also that the notion of knowledge is not equivalent to communication: in fact, using a communication mechanism does not entail awareness and vice versa, a MRS can be aware even though there is no direct communication among the robots.

Coordination Level: The third level is concerned with the mechanisms used for cooperation. Following the literature on MAS [38], we consider *Coordination* as: “cooperation in which the actions performed by each robotic agent take into account the actions executed by the other robotic agents in such a way that the whole ends up being a coherent and high-performance operation”. However, there are different ways a robot can take into account the actions of the other members of the team. The underlying feature is the *coordination protocol*, that is defined as a set of rules that the robots must follow in order to interact with each other in the environment. Therefore, we can further classify the coordinated MRS based on the type of coordination protocol. We consider *Strong (Weak) coordination* as a form of coordination that relies (does not rely) on a coordination protocol. As a difference with MAS, the approaches based on weak coordination are more interesting in MRS, since in some cases the effective use of a coordination protocol may be difficult.

Organization Level: The fourth level of our hierarchical structure is concerned with the way the decision system is realized within the MRS. The *Organization Level* introduces a distinction in the forms of coordination, distinguishing centralized approaches from distributed ones. A precise characterization of this issue is given for example in [32], where distribution is regarded as the autonomy of each component in the system to take decisions about the actions to perform. In particular, a *Centralized* system has an agent (leader) that is in charge of organizing the work of the other agents; the leader is involved in the decision process for the whole team, while the other members can act only according to the directions of the leader. On the other hand, a *Distributed* system is composed of agents which are completely autonomous in the decision process with respect to each other; in this class of systems a leader does not exist. The classification of centralized systems can be further refined depending on the way the leadership of the group is played. Specifically, *Strong centralization* is used to characterize a system in which decisions are taken by

the same pre-defined leader agent during the entire mission duration, while in a *weakly centralized* system more than one agent is allowed to take the role of the leader during the mission.

Along with the classification introduced to characterize the form of coordination, there are a number of system features that are relevant to the development of the system. We have grouped them in the *System dimensions*, which include: communication, team composition, system architecture and team size.

Communication: Cooperation among robots is often obtained by a communication mechanism that allows the robots to exchange messages. A detailed analysis of the various technical problems related to communication in MAS is given for example in [38]. However, when MRS are considered the communication mechanisms are very different; in addition most of the MRS that we consider in this article operate with a limited number of robots (i.e. less than 10), except for a few recent projects for large-scale MRS that take into account about 100 robots, while in large-scale MAS the number of agents can often be in the order of 10,000-100,000. These observations show that communication issues have, in general, different characteristics for MAS and MRS. Therefore, even though it is possible to have a more precise characterization of communication systems (e.g. regarding topology, range and bandwidth as studied in [31]), in the taxonomy proposed in this article we distinguish two different types of communication depending on the way the robots exchange information: direct or indirect communication. Direct communication makes use of some on board dedicated hardware device, while indirect communication makes use of stigmergy³. The fact that in MRS direct communication is based on a dedicated physical devices, results in a much more expensive and unreliable solution to attain coordination with respect to MAS. Therefore, Indirect communication has received particular attention in MRS literature, to cut implementation and design costs. While in MAS direct communication has been extensively used, remarkable exceptions such as [16] support stigmergic communication, in this work a Pheromone Infrastructure for information sharing is used. Stigmergic communication can both guarantee locality in the interactions among agents, reducing the complexity for the design of large scale systems, and avoid the need of synchronization between the agents, by providing a shared communication structure that each agent can access in a distributed concurrent fashion.

Team Composition: According to team composition MRS can be divided in two main classes, heterogeneous and homogeneous (see for example [99]). Homogeneous teams are composed of team members that have exactly the same hardware and control software, while in heterogeneous teams the robots differ either in the hardware devices or in the software control procedures. This distinction is used also for MAS, but in that case the differences are obviously only in

the software implementation of the agents' behaviors.

System Architecture: System architecture is an important feature for classifying MRS as well as MAS. In this work we always refer to the architecture of the whole MRS and not to the architecture of the single robotic agent. A precise characterization of MRS with respect to reactive or deliberative architectures is presented in [55]. We consider a team architecture as *deliberative* if it allows the team members to cope with the environmental changes by providing a strategy to reorganize the overall team behaviors. On the other hand, in *reactive* team architectures each robot in the team copes with the environmental changes by pursuing an individual approach to reorganize its own task in order to accomplish the goal assigned to it. The main difference between deliberative and reactive team architectures relies on the different approaches adopted by the MRS to recover from an unpredicted situation: in a deliberative MRS a long term plan involving the usage of all the available resources to collectively accomplish a global goal is provided; in a reactive MRS a plan to cope with the problem at hand is provided by the robotic agent directly involved with it.

Team Size: The team size is an important issue for MAS and it is becoming a relevant issue also in MRS development, actually a number of recent works explicitly address large scale MRS [82], [63]. However, the number of robots acting in the same environment is still quite limited with respect to the number of agents in a MAS. Therefore, rather than a quantitative measure of the size of the MRS in our taxonomy we distinguish those approaches that explicitly consider as a design choice the opportunity to deal with a large number of robots, from those that do not.

A. Related work

In the recent literature there are few works that overview and classify the research on MRS [19], [31], [30], [81]. Besides being focused on the most recent developments in the field, the present work aims at providing a novel perspective by focusing on the coordination capabilities of the MRS.

The classification presented by Cao et al. [19] is the closest to ours, since some of the aspects that are relevant to coordination are considered in their taxonomy. In particular, they are referred to as *Group Architecture*: centralization/decentralization of the coordination approach, differentiation among the team composition, communication infrastructure and knowledge of the other agents. Not only we further distinguish between coordination and system dimensions, but we provide for a more refined analysis of the coordination aspects in the *Coordination Dimensions*. On the other hand, we do not take into account other dimensions such as *Origin of Cooperation* and *Geometric problem*. As for *Resource Conflict*, if the resource is the space in which the robots operate, then all the works on MRS implicitly face this problem by providing different solutions. However, when the modality for accessing a shared resource becomes more complex, a coordination protocol and a distributed solution is usually preferred. In fact, there are few works explicitly addressing this issue, some of them are described in Section IV-F.

³Stigmergy is a term coined by the biologist P.Grassé which means to incite work by the effect of previous work [65]. Here, stigmergic communication refers to the sharing of information through modifications in the environment.

The work of Dudek et al. [31], [30] presents a different and rather complementary classification. In particular, the issue of communication is considered in detail with the three dimensions *Communication-Range*, *Communication-Topology* and *Communication-Bandwidth*. In this work we do not explore in detail the communication structure, but we consider the issue of communication by simply distinguishing between implicit and explicit communication, which is most relevant to the kind of MRS dealt with in this article. Specifically, we adopt an abstraction of the communication channel, since it is convenient to decouple the communication layers from the coordination capabilities. We refer the reader to the cited papers [31], [30] for a detailed discussion of the impact of the communication capabilities on the system. Another difference with respect to our work is that, as previously mentioned, our focus is on engineering applications for MRS, and thus we do not consider in detail issues such as *Collective Reconfigurability*, which are more relevant to biologically inspired MRS or *Processing Ability*, that refers to the architecture of the computing system used in each single agent.

A work that takes into account the issues of coordination both in MAS and MRS is [110]. This work also introduces some of our System Dimensions like *Team Composition* and *Communication*, however it is mostly focused on Machine Learning techniques for MAS and MRS. In fact, due to the inherent complexity of MRS, Machine Learning is an issue of great interest [114], [95] and many existing learning techniques can be directly applied in a multi-agent scenario, by delimiting a part of the domain that involves only a single agent, as described by Weiss [43]. Multi-agent learning is however focused on learning techniques which take explicitly advantage from the presence of several agents in the environment, and MAS approaches have been generalized to MRS. For example, Reinforcement Learning is often used in MRS applications in order to improve the performance of cooperating agents: in L-ALLIANCE agents learn how to better estimate the performance of the other agents [87], in [72] reinforcement learning is properly modified to cope with noisy, and dynamic environment, while in [39] Q-Learning is used to approach the Multi-Robot Observation of Multiple Moving Targets. We refer to some of these works for their approaches to coordination, though we do not specifically consider learning techniques in MRS.

In [90] a formal analysis of optimality and complexity for teamwork theories [49], [106], [57] in MAS is presented. The work uses a general framework called the Communicative Multi-agent Team Decision Problem (COMM-MTDP), based on decentralized partially observable Markov decision processes. Using this framework, a classification of teamwork theories is proposed along the dimensions of observability and communication cost, deriving optimality and complexity results. The authors focus on four categories for the observability: *Collective Partial Observability*, where no assumption is made on the agents' observations, *Collective Observability*, where a unique world state can be derived from the collection of the agents' observations, *Individual Observability*, where each agent can derive a unique world state from its own observation and, finally, *Non Observability* where no feed-

back is provided to the agent from the environment. As for communication the authors consider three categories: *General Communication*, where no particular assumptions are made on communication cost, *Free Communication*, where no cost is associated with communication acts, and *No Communication*, where there is no explicit communication. For all the cases that can be characterized by the combination of those categories, time complexity results are given to find an optimal policy for the team. Moreover, the authors report the analysis of a specific domain through an empirical evaluation and comparison of different communication policies taken from the literature, with respect to a local optimal and globally optimal policy inferred by the authors, varying the communication costs and the degree of observability. The comparison pinpoints the various situations (i.e. combination of communication costs and degree of observability) where the communication policies show sub-optimality with respect to the globally optimal policy. The theoretical and experimental results obtained in the paper for teamwork in MAS are very interesting: even if the assumptions made could not always apply to robotic systems (i.e. the absence of delay or noise in the communication channels), they provide useful guidelines for the design of a coordination approach for a MRS based on teamwork theories.

In [69], [71] the problems concerning the analysis and synthesis of intelligent group behaviors in MRS are addressed in detail. Definitions for key concepts related to MRS coordination are given, thus precisely characterizing several important aspects of the problem. In particular, distinctions are made according to the ability of the agents to recognize their kin, and based on whether the agents have an explicit model of their team mates. There is a clear correspondence between our notion of *Awareness* for a robotic agent, and its ability of kin recognition, but we do not classify the *Aware* agents according to their model of team mates, we rather consider the protocol used for coordination. Moreover, the approaches that we consider as *strong coordinated* are necessarily able to recognize their kin and generally a model of the team members is present; on the other hand, systems composed of robotic agents which are able to recognize their kin but do not have a model of the team mates could be classified in our taxonomy as *weakly coordinated* or *aware not coordinated*. Regarding the issue of communication, in [69] a definition similar to the one proposed in this article is given for *Direct Communication* and *Indirect Communication*, however, a further distinction is made between *Direct Communication* and *Directed Communication* (which is a direct communication aimed at a precise receiver). Finally the definition of *Explicit Cooperation*, considered in [69], is similar to the definition of *Coordination* provided in this work, while *Implicit Cooperation* can be related to our definition of *Cooperation*.

The work by Parker [81] is an else survey of the MRS literature, whose goal is to characterize several primary research topics for MRS research. It is focused on distributed robotic systems and in particular on research that has been demonstrated on physical robot implementation. Therefore, in [81] a taxonomy, as above discussed, is not provided: only some of the research issues such as *Communication* and *Architectures* correspond to dimensions in the proposed taxonomy.

Conversely, *Localization, Mapping and Exploration, Object Transportation and Manipulation* and *Motion Coordination*, in our discussion are regarded as general tasks. As already discussed, we do not explicitly address *Biological Inspiration*, but several works that are presented in [81] under the above research topic are also addressed in our classification in light of their approach to coordination. We do not look at *Reconfigurable Robots*, because those works, at present, are mostly focused on hardware development. Nonetheless, the ability to change the number and the capabilities of the robotic agents could be regarded as one of our system dimension.

A recent work focusing on the analysis of task allocation in MRS is presented in [44]. This work presents a taxonomy to analyze the different approaches to the problem of Multi Robot Task Allocation (MRTA) appeared in the literature, proposing a formal framework for the study of this problem. The authors consider three main dimensions: 1) single-task (ST) vs multi-task (MT) robots, based on whether the robots involved in the task assignment process can execute more than one task at a time; 2) single-robot (SR) vs multi-robot (MR) tasks, considering if the tasks to be performed involve one or multiple robots; 3) instantaneous assignment (IA) vs. time-extended assignment (TA), distinguishing whether the information concerning the robots tasks and environment permit only an instantaneous assignment or a more sophisticated planning approach. The authors provide a formal characterization of a wide set of MRTA problems, analyzing and classifying in their taxonomy significant approaches used in MRS literature; for those approaches, bounds to the optimality of the method used with respect to the optimal solution are provided. The focus of this work is posed on a specific problem (MRTA), which represents one, important, aspect of MRS coordination. Being the problem of task assignment generally approached using a well defined coordination protocol, all the works analyzed in [44] can be precisely classified in our taxonomy as *Strongly-Coordinated*. Therefore, the taxonomy presented by the authors can be considered as a more detailed investigation of the approaches that we classify as *Strongly-Coordinated*.

Finally, a chapter of [4] is devoted to social behavior for MRS. Several interesting characteristics and fundamental problems for MRS behaving in a society are pointed out. The chapter is mostly focused on several aspects of MRS social structure, such as social organization, communication, distributed perception and social learning, providing a broad perspective on MRS.

III. TASKS AND DOMAINS FOR MRS

Although we do not consider the application domain as a dimension in our taxonomy, we believe it is useful to give an overview of the specific test-beds which are commonly used in the MRS literature for validating and evaluating the proposed coordination techniques. In particular, we divide each of the described test-bed in a generic description of the abstract task the robot should be able to execute, and the real application domains that are strictly tied to it. By abstract task we refer to a general description of the goal the MRS should pursue, without considering the details of the application, while the

application domains tied to the test-bed are the real-world applications that require similar coordination capabilities.

Foraging and Coverage: The foraging task requires the components of the MRS to pick up objects scattered in the environment; foraging is a test-bed often used for MRS, because of its analogies with tasks like toxic waste cleaning, mine cleaning, service robots [78], [59]. A major issue in this test-bed is to avoid interferences among the robots during the task execution. The coverage task is very similar to foraging, since it requires the robots to process all the points of the free space in the environment [51]. The main issue for coverage is therefore to find effective techniques for cooperatively scanning all the environment. Applications for coverage are: demining, snow removal, lawn mowing, car-body painting, etc.

Multi-Target Observation: The task of multi-target observation requires a team of robots to detect and track a set of moving objects. The robots have to maximize the time during which each of the moving target is being observed by at least one of the robotic agents within the MRS. The multi-target observation (also known as CMOMMT: Cooperative Multi-Robot Observation of Multiple Moving Targets) is a very recent MRS test-bed, first introduced in [86]. Besides systems composed of mobile robots, sensor nets are often used to accomplish this kind of task [34], [15], [75]. Such systems are typically composed of devices with very limited computational capabilities, and are able to communicate among them and monitor the environment. Although sensor nets represent an effective solution to this task, as previously remarked, they are not in the scope of the MRS we are considering in this paper. Multi-Robot Observation has many connections with security, surveillance and recognition problems [86], [116], where targets moving around in a bounded area must be observed.

Box pushing and Object Transportation: The task of box pushing requires the robotic agents to cooperatively push boxes in order to reach a desired configuration. The box pushing task has analogies with problems like stockage, truck loading and unloading. While in the box pushing task the objects are assumed to be on a plane, object transportation focuses on lifting and carrying objects [97], thus substantially increasing the complexity of the task. Applications involved are transportation of heavy objects in industrial environments or assembly of large-scale structures, such as terrestrial buildings or planetary habitat [97]. In most of the applications it is assumed that each robot can not carry the object alone, thus object transportation is frequently used as test-bed for issues like motion coordination and formation control.

Exploration and Flocking: Under the label of exploration and flocking different tasks can be grouped: these tasks differ in the way they are realized, but have the common feature to require MRS members to coordinate their movements in the environment. Behaviors like flocking, formation maintenance or map building can be considered in the same class. In the exploration task the robots must be spread in the environment in order to collect as much information as possible about the surrounding area. In the flocking task the goal for the robotic agents is to move together, such as in a

flock. The formation task is focused on having the robots move in the environment forming particular shapes. Cooperation among the robotic agents is also used to localize each other, and to fuse information acquired from the environment. Map building of unknown environments is a common issue related to exploration, and in particular, a very interesting topic is the cooperative simultaneous localization and mapping, in which the robots need to localize while moving and building the map of the environment [37]. The problem of exploration and flocking is related with several applications such as transshipment operations in harbors, airports and marshalling yards, motion coordination in industrial application and exploration of dangerous environments. Another example of exploration task is given by the RoboCup *Robot-Rescue* league [62], that is a setting for experimenting MRS involved in searching victims in an unknown environment representing a disaster scenario.

Soccer: Robotic soccer has been considered in the last years as an interesting test-bed for research in multi-agent and multi-robot cooperation [61]. The uncertain dynamics and hostile environment in which the robots operate makes coordination of the multi-robot system a real challenging problem. While in the early years of the robotic league competitions the focus has been on improving the single robot capabilities, only recently coordination in the MRS has become a central issue. The different settings of each of the robotic leagues present several issues for coordination in MRS. In particular, in the *Middle-Size league* and the *Four Legged league*, all the robot sensors must be on board; therefore robots are more autonomous and have to deal with high uncertainty in reconstructing global information about the environment. On the other hand in the *small size league* the robotic agents can take advantage of a top view of the environment provided by a camera on the top of the field, therefore the coordination approaches in this league are mostly centralized. The use of coordination in the soccer domains has demonstrated significant improvement in the performance of the teams.

IV. SYSTEM CLASSIFICATION

In this section we describe several works related to MRS, by collecting them according to their position inside our taxonomy. By discussing concrete examples of implemented MRS we aim at characterizing in deeper detail the proposed taxonomy. Moreover, we describe the constraints imposed by the different approaches to coordination presented, and consider the tasks and domains that are addressed within each class.

A. Unaware systems

Unaware systems are characterized by the fact that each member of the MRS executes its own task without any knowledge about the other members of the team. Obviously, in this setting, coordination is not possible, while cooperation among the robots can still be obtained in a goal driven manner. Due to the fact that each member of the team does not have knowledge of the other robots, the communication among the robotic agents can not be direct. It is worth mentioning that

all the works we classified as unaware can not be considered as *deliberative*, because a deliberative approach requires a reorganization of the whole team and thus the robots must be aware of each other.

Unaware approaches are frequently adopted in biologically inspired MRS [65], [25]. The robots achieve cooperation by using only very simple basic behaviors and exploiting only local interactions. These kind of systems are well suited for large scale development, but they are normally used only for very simple tasks such as foraging [25] or box pushing [65]. In particular, in [25] stigmergic communication is used for a team of robots collecting objects by using the simple rule of transporting objects near other objects. In [65] the authors use a cooperative object transportation model inspired by ant colonies; stigmergic communication is achieved through the perceived stimuli on the item being transported. Another biologically inspired model, based on the behavior of *E. coli* bacteria, is presented in [88], where a cooperative foraging task is addressed. *E. coli* bacteria are characterized by a very simple and effective foraging strategy, which relies on the sensing of nutritional and toxic substances. The author uses this model to design a non-gradient optimization algorithm that drives the bacteria in their search for the optimal solution of the objective function. This model has been used in a preliminary study for controlling groups of Unmanned Autonomous Aerial Vehicle [35], that should be able to cooperatively search for interesting target in a dangerous environment. Also Pheromone based models have been considered to design cooperative unaware systems. In particular in [17] multiple pheromone types are used by static entities (named pumps) to guide moving entities (walkers). The goal of the moving entities is to reach the positions occupied by the pumps, resulting in a special kind of foraging task with application to military Air campaign. The Pheromone types have different propagation radii, in order to provide a better guidance function over all the working space. Pure pheromone approaches (such as the one discussed above) can be considered as unaware; however, a protocol based information exchange on an underlying pheromone infrastructure can provide the system with more sophisticated coordination capabilities (see Section F).

Notice that, as previously described approaches show, in unaware systems interactions among entities are possible through the use of stigmergy. However, each active entity interprets the modifications of the environment, without having any model of the cause of those modifications. In particular, if an entity A observes a particular modification of the environment, it has no ability to distinguish if the source that caused that modification, is an entity cooperating with A, an exogenous event (i.e. a human) or the entity A itself. Therefore, the behavior of A will always be independent of the source of the observed modification. This property of the system can be considered as the main criterion to classify such approaches as *Unaware*.

In many of the unaware systems, the MRS is composed of homogeneous robots; however, [5] makes a comparison between homogeneous and heterogeneous unaware MRS in a multi foraging task (i.e. foraging with different kinds of objects) by considering the relation between performance and

a metric, called *social entropy*, which denotes the degree of diversity within the system. The results presented show that even an unaware approach can benefit from the heterogeneity of the team members, and that this benefit is tightly related to the inherent complexity of the task.

In [117] a robot team is used to cooperatively transport an object. The team is composed of a leader robot and some followers (three in the example reported in the paper): the leader has different hardware, and is in charge of observing the environment and plan the object motion. The followers can detect the object motion and react to the object movement based only on their local perception. The idea of the proposed system is that all robots support the object together to share the load and keep the balance of the object, while each follower tries to minimize the inter-force with the object to reduce the driving effort of the leader. Although the overall system requires the presence of a leader and several followers, each robot acts as it is the only agent in the system, thus the system can be classified as *unaware*.

B. Aware, not coordinated systems

Aware systems are characterized by the fact that the robots of the team have knowledge of the presence of other robots in the environment, and act together in order to accomplish the same global goal. However, a robot may not take into account the actions performed by other robots in order to accomplish its task, and in this case we consider it as aware not coordinated. It is not always easy to give a general criterion to precisely distinguish whether each robot is taking into account other's robot actions during its task execution. In general, aware not coordinated approaches are characterized by simple methods to reduce interference among robots executing a cooperative task, while avoiding the use of a specified protocol. It is worth noticing that avoiding interferences among robots is definitely different from simple obstacle avoidance, because obstacle avoidance does not require the robot to recognize its team mates and because interference can occur in more general settings, as compared with conflicts on physical space, such as sensorial interference.

A clear example of the differences among unaware, aware not coordinated and weakly coordinated approaches is presented in [11], where three different approaches are used for the coverage of an unknown environment by a team of cooperative robots. One of these approaches can be considered unaware, the second aware not coordinated, while the third is weakly coordinated. The overall design choice is to have behavior-based robots using only local sensing: in the unaware approach, each robot chooses the most promising direction of motion simply evaluating through vision the direction that maximizes the frontal visibility; in the aware not coordinated approach, the robots recognize each other and choose the direction of motion which is opposite to the average angle subtended by all its neighbors in its visual field; in the weakly coordinated approach, when two robots recognize each other they form a coalition, and calculate the most promising direction of motion as that direction that maximizes the coalition sensorial coverage. The experiments presented in the

article show that the aware not coordinated approach slightly outperforms the weakly coordinated and that both approaches outperform the unaware one. This work shows clearly the difference between an aware, not coordinated approach and a weakly coordinated one. Moreover, it shows by means of experiments that, in some cases, an aware not coordinated approach can outperform a weakly coordinated one.

Particular cases of aware not coordinated approaches might end up in competitive settings; while those kind of approaches do not fall in our taxonomy that is explicitly focused toward cooperative approaches, it seems worthwhile to report an example of such approaches that clearly pinpoint the difference between competitive and cooperative coordinated approaches. In [107] the authors present a very interesting solution to the problem of interferences among aircrafts using two different levels: at the first level a competitive approach is used while at the second level coordination among aircrafts is required. The authors address the problem of air traffic management in free flight mode⁴. At the first level interference among aircrafts are avoided following a competitive approach: each aircraft is modeled as a player in a n-player zero-sum game, each aircraft models the action of the others as a known set of disturbance values, ignoring which particular value will be actually used in the particular situation. Each aircraft solves the game considering the worst possible disturbance values, if a saddle solution exist and this solution is within the safe requirement for each aircraft then the aircraft need not to coordinate and not even to cooperate. However, if this condition is not met, then the conflict resolution is addressed at the second level: using a coordinated approach the aircrafts perform predefined *a-priori* safe maneuver.

In aware not coordinated systems, cooperation among robotic agents is often considered as an *emergent* property of the system that results from the interaction between the system and its dynamic environment [98]. The collective task is therefore designed using an interaction loop between the system and the environment, ultimately converging towards the desired performance. In [66] such an approach is presented in several different tasks. In particular, a box pushing task is implemented with robots whose collective behavior is achieved by providing each of them with basic behaviors. The basic behaviors are activated according to the progress of the overall task, which is monitored by perception of each single robot. The approach achieves coordination by exploiting only local information that the robots can acquire from the environment and therefore it is suitable for coordination of large size teams.

C. Weakly coordinated systems

MRS may present a form of weak coordination that does not rely on the application of an explicit predefined coordination protocol. By coordination protocol we mean a set of explicit predefined rules, which are followed by all the robots of the system, that clearly define the behavior of the robots depending on the messages exchanged among the team mates.

⁴In free flight mode each aircraft can autonomously execute small detour from the assigned route to achieve better flight condition or avoid conflict due to schedule delay

By behavior we mean a high level action the robot can accomplish: typical examples of behaviors are going in a predefined position, pick up an object, track an object and so on. Weak coordination does not pose any constraints on the *System* dimensions defined in our taxonomy. Thus all the combinations of *Communication*, *Architecture Composition* and *Size* are possible.

In particular, in [105], a weakly coordinated approach for object transportation, avoiding the use of communication is presented. Two robots cooperatively transport an object using free rotational joints, that keep the object tied to the robots; the joints are equipped with a force sensor, able to measure the relative angle between the axes of the object and the robot. One of the robots is the leader and is in charge of executing a desired given trajectory for the object, while the follower estimates the motion of the object and avoids the obstacles present in the environment. No communication is held between the leader and the follower, however the follower estimates the motion of the leader and use this information for better accomplishing its task realizing a coordinated approach. The two robot coordinate themselves without the use of communication and therefore, without an explicit protocol, thus the approach is to be considered weakly coordinated.

In [111] another weakly coordinated approach that avoids communication is proposed. The system is used to cooperatively clean a room, the robots are homogeneous and each one performs the next move based only on its local perception, thus the system can be considered reactive. Cooperation is achieved by applying an algorithm that tells each robot to clean a location only if it is not critical: a non critical location is a location that does not disconnect the current graph of dirty grid points. The robots stop when no dirty points exist in the environment, the communication is stigmergic in the sense that the dirty points can be considered as markers, that the robots use to cooperate in the cleaning process. In this case coordination among robots is achieved, by following a general rule that constrains the robot behaviors, integrating the actions of the robots in a convenient indirect way.

In [6] the task is formation control and communication is used to exchange only the global position among the robots; considering those information and the formation to be achieved each robot tries to adjust its relative distance with respect to its team mates. The way robots achieve the desired distance is not specified by a coordination protocol and thus the approach can be considered weakly coordinated. This approach can be further classified as reactive because each robot reacts to the unexpected environmental changes, in order to maintain the predefined formation, without reorganizing the overall team. A deliberative approach for a weakly coordinated system is also possible, as presented in [41]. To avoid interference among the robots, the environment is divided into regions of work, the number and the size of the regions depend on the number of working robots. At the beginning, each robot is assigned to one region; during the task execution each robot broadcasts a diagnostic message, that means the robot is working correctly; on the basis of the number of the diagnostic messages received, the working area division is reassigned. This work is an example of a deliberative approach, because the regions of

competence are assigned on the basis of information regarding the overall team (number of working robots), although the communication protocol is rather simple. The information derived from the messages are used to change some parameters for the behaviors (i.e. the boundaries of the working area for each robots) and not the behaviors themselves (i.e. searching for an object, picking up the object and bring it to the deposit position), therefore the approach is to be considered aware not coordinated.

D. Strongly coordinated, strongly centralized systems

Strongly coordinated MRS are based on a system of signals by which the robots in a team exchange information, according to a predefined coordination protocol, concerning the way the robots have to interact. Among these systems a further classification can be done depending on the way the protocol is implemented. In strongly centralized systems a particular robot (called *leader*) is in charge of organizing the work of the entire team, while the other members act according to its directions. Although in most works direct communication is used, it is possible to conceive a system where a leader uses stigmergic communication in order to give “commands” to the other robots. Moreover, most works in this class present a deliberative approach, that is not a direct consequence of the choice of realizing a strongly coordinated, strongly centralized MRS, but just a very convenient way of implementing it.

The MRS presented in [96] deals with exploration and formation maintaining. A predefined leader is in charge of selecting one formation among a set of predefined ones, according to the current situation, and then communicating to the other robotic agents their dislocation within the environment. The strongly coordinated strongly centralized approach is frequently used in systems designed for space missions; in [10] a team of cooperating rovers is based on the use of the central station that coordinates the rovers during the task execution. To cope with dynamic changes of the environment a continual planning approach is used, and three different approaches to coordination are presented. In the first approach the continual planning is performed on the central station, and the commands are sent to the rover; in the second approach tasks are assigned to the rover from the central station and each task is planned by the rovers; finally, in the third approach the central station rules an auction and the rovers bid to get the assignment of the tasks. This work shows how in a strongly centralized approach a leader can take into account different levels of autonomy for the team members.

Another example of centralized system is SAMON [67], where the missions of multiple Autonomous Underwater Robots Vehicles (AUV) are coordinated through a behavior-based architecture. In this case, several layers of coordination are used: Tactical Coordinator, which issues missions orders, Supervisory AUVs, which distributes the subordinates AUV the search regions, AUVs, which decide their own itinerary to collect data from the fixed sensory packages.

Also most teams in the RoboCup *Small-Size league* (e.g. [109], [13]) use a strongly coordinated strongly centralized approach, since robots are usually controlled by a remote

host through a global vision system. In fact, the availability of global information about the environment and the use of a remote host for robot control naturally lead to the implementation of centralized strategies.

In most cases centralized systems are not well suited for coordination of large scale MRS, both for the communication overhead among the team members, and for the high computation demand required by the leader.

E. Strongly coordinated, weakly centralized systems

Strongly centralized MRS are not robust to failures in communication and to incorrect operation of the leader. Therefore, in many applications a different kind of centralized system has been preferred. Weakly centralized systems are characterized by the fact that the leader is not chosen a priori, but it is selected dynamically during the mission depending on the current situation of the team and the environment.

Several policies have been used to decide which robot should become the leader; in particular, in [73], a fixed priority is set among the robotic agents, and if one of them is not available as leader, the next agent in the priority order takes the control of the team. This is possible because the robots are homogeneous. In [80] the leadership is given to a robot depending on its specific characteristics (the robots are heterogeneous); for example, a box pushing task is accomplished cooperatively by a supervisor and a pusher; the supervisor is the leader, because it can monitor the overall process with a camera. In [97] the robots start an auction and “bid” in order to become the team leader. In [28] the authors propose a free market based approach to coordinate a group of robots. Costs and revenues are associated to the tasks, and robots can trade the tasks allocation trying to maximize their revenue. The authors show that the use of a leader robot proposing the tasks allocation to a group of robots, can enhance the team performance. The leader of a group is chosen according to the quality of the allocation that the candidate leader proposes to the group: if the proposed allocation reduces the associated cost of the team the resulting excess profit can be redistributed within the group (including the leader), in such a way that the robots in the group accept the leader’s allocation.

In [80] and [97] weakly centralized approaches adopting a hybrid architecture are presented. By hybrid architecture here we mean an architecture that can be both *social deliberative* and *reactive*. Specifically, both those approaches are characterized by organizing the architecture of each single robotic agent in several layers (three in both the works) in a hierarchy. Each layer is able to communicate with its peers among the robotic agents, exchanging different type of information according to the level in the hierarchy. Depending on the level of information exchanged among the different layers, and thus on the protocol used, it is possible to achieve a *social deliberative* or *social reactive* approach.

F. Strongly coordinated, distributed systems

In distributed systems each team member is executing a coordination protocol, while taking decisions in a completely autonomous fashion. These systems are generally more robust

to communication failures and robot malfunctioning, even though these problems may affect the overall performance of the team in the accomplishment of the task. The large number of works that adopt this approach shows that there is a clear interest among the researchers to approach the problems of cooperation in a distributed fashion. The *strongly coordinated distributed* approach entails that some kind of communication has to be used, and leaves unconstrained the other System Dimensions.

A very interesting approach is the one used by Parker in the ALLIANCE architecture [85]. ALLIANCE is a framework for coordinating MRS composed of heterogeneous behavior-based robotic agents. All the robotic agents have sets of behaviors, which are controlled by modules called *motivational behaviors* that can cross inhibit each other. Motivational behaviors are based on two parameters: *impatience* and *acquiescence*; they are updated based on the data that each robot acquires from its sensors or from other robots. In [86] the ALLIANCE architecture is used in a Multi-Target Observation. The approach of the ALLIANCE architecture is reactive in the sense that each member of the team decides whether to employ itself in accomplishing a task, without any need to reorganize the other members activity.

In [116] the problem of multi-target observation is also addressed. The robotic agents are behavior-based and homogeneous, but the technique proposed is perfectly applicable to a heterogeneous system as well. The proposed architecture, called Broadcast of Local Eligibility (BLE), is an extension of the subsumption architecture, to enable coordination between the robotic agents. Each behavior of each robot has a function that locally evaluates the robot’s eligibility to accomplish a given task; the values are then exchanged among the “peer behaviors” of the robotic agents. The robot, whose behavior has the highest value, inhibits the corresponding behaviors on the other system members, thus advocating the task. This process, called cross-inhibition, can be executed only among peer behaviors. The approach followed in this work is social deliberative in the sense that when a robot starts accomplishing a task, it inhibits the peer behaviors of the other members. Thus, when something happens, which imposes a selection of a different action to be executed by a member of the MRS, all the other members will be involved in this reorganization and the new action will be executed by the robots which best fit the requirements, thus obtaining a new strategy.

Communication in the *strongly coordinated distributed* approach does not need to be direct, a coordination protocol can be realized using stigmergic communication. An interesting example of this approach is presented in [16], where a MAS is used for the control of a manufacturing system. The work uses a Pheromone Infrastructure (PI) as stigmergic information sharing mechanism. This approach is designed for a MAS, therefore the stigmergic communication is realized as a shared software framework (i.e. the PI), and not using physical signals that can be perceived through the environment. In [54] an example of a strongly coordinated approach for a MRS making use of physical signals for stigmergic communication is presented. The work is explicitly designed for large scale systems, the test-bed is composed of a group of very small

autonomous mobile robots, equipped with a gripping device; the robots have to cooperate in pulling some sticks out of the ground. The sticks are too long for being pulled by one robot alone, therefore the robots must pull the same stick together. The robots walk randomly in the environment and when they find a stick, they try to pull it up. If the stick is not being pulled by another robot they wait a certain amount of time (gripping time) for help, otherwise they help another robot in pulling the stick. This work shows how a task that needs a tightly coupled coordination can be performed using only stigmergic communication.

Strong coordination often requires some kind of synchronization on the use of the available resources among the robotic agents. Among others, the works in [2] and [59] explicitly address the issue of conflict resolution on shared resources in a MRS providing a distributed approach. In [2] the approach is based on a central station that plans a mission for all the agents. The mission assigned to the agents consists of very high level instructions, leaving each robotic agent the execution and coordination of the high level actions. To avoid conflicts on shared resources the plans of the agents are merged; in particular, any time an agent needs to execute a new plan that uses shared resources, it executes a Plan Merging Operation, trying to set timing constraints on the actions of all the plans that access the shared resources. In [59] the application domain is a cooperative cleaning task, two robots, a vacuum and a sweeper, have the common goal to clean an office like environment. The architecture is called ABBA (Architecture for Behavior-Based Agents) and is based on a network consisting of two types of nodes: Competence Modules (CM) and Feature Detection (FD). The edges between the nodes represent relationships of different kind (successor, predecessor, conflictor, precondition, positive or negative correlation), which can be learned and modified during the task execution. Each node has a level of activation that is controlled through the relationships, conflict resolution is achieved by establishing the relationships among the nodes in a proper way.

Finally, an example of distributed MRS in the soccer domain is given by the RoboCup Azzurra Robot Team (ART) that has implemented a distributed heterogeneous robotic soccer team [23], based on a simple and flexible coordination protocol. The approach makes use of a *formation/role* mechanism and of dynamic assignment of roles. Role assignment is obtained by explicit communication of information about the status of the environment. A simple form of negotiation is used in order to realize a deliberative, distributed MRS that does not require a global representation of the environment. Each robot has the knowledge necessary to play any role, and robots switch roles on the fly, when a distributed agreement on the actions to be performed is achieved. This work presents an effective example of dynamic task assignment for a MRS performing a very complex task in a highly dynamic and hostile environment.

V. TRENDS AND CONCLUSIONS

In this paper we have addressed the recent developments in the field of MRS, focusing on those approaches that are

targeted to specific applications and motivated by engineering considerations. Specifically, we have presented a taxonomy with the aim of highlighting the coordination aspects of the recent proposals in the literature: we have defined a set of *Coordination Dimensions* for the classification of the approaches to team coordination, together with a set of *System Dimensions* that account for the design choices that are more relevant to the team organization. Moreover, we have identified some prototypical tasks that characterize various application domains. Finally, we have classified some of the recent works on MRS in terms of the proposed taxonomy; a summary of the classification is presented in Table II, where the approaches described are positioned according to the task accomplished and their placement inside our taxonomy.

Although the classification is focused on coordination and, consequently, biased towards a class of MRS, a few reflections on the outcomes of the analysis performed seem worthwhile. A first observation is that both *Unaware* and *Aware not coordinated* approaches can achieve very interesting results in the execution of tasks, such as *Foraging* or *Box Pushing*. Therefore, cooperation without coordination can be successful in MRS, when the form of cooperation can be obtained with a loose coupling of the agents. Arguably, approaches that fall in this segment of the taxonomy are more frequent than in the case of MAS.

However, the analysis of the recent works in the literature shows that for more complex tasks (e.g. Soccer, Rescue missions, etc.), where the unpredictable, uncertain and sometimes competing environment requires both a very effective performance and high robustness, more complex coordination capabilities are required. In particular, among the *Strongly Coordinated* approaches, all the possible organizations have been extensively used, however a trend towards the development of distributed approaches is not surprising. Distributed approaches are generally more flexible, robust and less computational demanding. In the case of aware coordinated approaches, the techniques for coordination have been largely inspired by the literature on MAS. However, the major problems that have been tackled are concerned with the application of the coordination techniques in critical conditions due to the uncertainty of the environment and to the limitations and inaccuracies in the sensing capabilities of the robots. We argue that further research developments along these lines may eventually lead to coordination models, that are explicitly designed for the robotic scenario. Moreover, as the tasks to be completed by the MRS become more and more challenging, thus requiring tightly coupled coordination, the issue of *Conflict Resolution* is likely to receive increasing attention.

Summarizing, the research on MRS addressed in the present work covers a broad range of approaches showing that the form of coordination can vary significantly depending on the task to be performed. The complexity of the tasks in which robots are involved (e.g. building patrolling, large-scale assembly, rescue operations) entails increasingly complex capabilities both in software and in hardware. Furthermore, the research on MRS is nowadays facing *Large Scale* systems [82], [63]: MRS with a large number of robots capable of

cooperative localization, long term autonomy, task assignment and conflict resolutions. In fact, teams formed by a large number of robots will impact also on the other *System Dimensions* of our taxonomy.

In this context, social deliberative architectures become more complex to realize, possibly requiring the introduction of different levels of organization; the adoption of a team strategy seems nonetheless needed, especially in uncertain and dynamic environments. Direct communication is an obvious choice to achieve coordination, when the application domain does not pose additional constraints, but it is easy to foresee that the increase in the size of the robotic team will introduce constraints in the structure of the communication network. Interestingly, among the MRS several alternative approaches have been proposed such as stigmergic communication or coordination without communication at all.

In the recent efforts on large scale systems, heterogeneity is often chosen in order to exploit different robot capabilities and reduce the cost of the overall system. Moreover, in this setting a *Weakly Centralized* approach has been adopted, where the more capable robot assumes the role of leader for the most complex tasks (cooperative localization, autonomous re-charging, autonomous teleoperation). This notwithstanding, within large scale systems distributed approaches can be advantageous also in the case of heterogeneous systems.

In conclusion, our analysis of the literature indicates that the problem of coordination will be central to the design of MRS, especially when dealing with complex tasks and large scale systems. In this respect, teamwork theories and team oriented programming [26], [106] could play increasing role in order to obtain more effective and general coordination frameworks.

VI. ACKNOWLEDGMENTS

This research is partially supported by MIUR (Italian Ministry of Education, University and Research) under the project RoboCare (A Multi-Agent System with Intelligent Fixed and Mobile Robotic Components).

REFERENCES

- [1] M. N. Ahmadabadi, S. M. Rushan, Z. Wang, and E. Nakano, "A constrain-move based distributed cooperation strategy for four object lifting robots," in *Proc. of the Int. Conf. on Intelligent Robots and Systems (IROS'00)*, 2000.
- [2] R. Alami, S. Fleury, M. Herrb, F. Ingrand, and F. Robert, "Multi robot cooperation in the martha project," *IEEE Robotics and Automation Magazine*, 1997.
- [3] Y. Arai, H. Asama, and H. Kaetsu, "Map acquisition in multi-robot systems based on time sharing scheduling," in *Distributed Autonomous Robotic Systems*, H. Asama, T. Arai, T. Fukuda, and T. Hasegawa, Eds. Springer, 2002, vol. 5, pp. 383–392.
- [4] R. C. Arkin, *Behavior Based Robotics*. MIT press, 1998.
- [5] T. Balch, "The impact of diversity on performance in multi-robot foraging," in *Proc. of Agents '99*, 1999.
- [6] T. Balch and R. C. Arkin, "Behavior-based formation control for multi-robot teams," *IEEE transactions on robotics and automation*, vol. 14, no. 6, 1998.
- [7] T. Balch and M. Hybinette, "Social potentials for scalable multirobot formations," in *Proc. of IEEE Intl. Conf. on Robotics and Automation (ICRA-2000)*, 2000.
- [8] T. Balch and L. E. Parker, Eds., *Robot Teams: From Diversity to Polymorphism*. A K Peters Ltd, 2002.
- [9] T. R. Balch, G. Boone, T. R. Collins, H. Forbes, D. MacKenzie, and J. C. Santamaria, "Io, Ganymede, and Callisto a multiagent robot trash-collecting team," *AI Magazine*, vol. 16, no. 2, pp. 39–51, 1995.
- [10] A. Barrett, G. Rabideau, T. Estlin, and S. Chien, "Coordinated continual planning methods for cooperating rover," in *Proc. of IEEE Aerospace Conf.*, vol. 1, 2001, pp. 1/133 – 1/140.
- [11] M. A. Batalin and G. S. Sukhatme, "Spreading out: A local approach to multi-robot coverage," in *Distributed Autonomous Robotic Systems*, H. Asama, T. Arai, T. Fukuda, and T. Hasegawa, Eds. Springer, 2002, vol. 5, pp. 373–382.
- [12] R. W. Beard, T. W. McLain, and M. Goodrich, "Coordinated target assignment and intercept for unmanned air vehicles," in *Proceedings of the 2002 IEEE Int. Conf. on Robotics and Automation*, Washington, DC, May 2002.
- [13] S. Behnke, R. Rojas, and G. Wagner, "A hierarchy of reactive behaviors handles complexity," in *Proc. of ECAI2000 Workshop on Balancing Reactivity and Social Deliberation in Multi-Agents Systems*, 2000.
- [14] A. Brendenfeld and H. U. Kobialka, "Team cooperation using dual dynamics," in *Proc. of ECAI2000 Workshop on Balancing Reactivity and Social Deliberation in Multi-Agent Systems*, 2000.
- [15] R. Brooks, P. Ramanathan, and A. Sayeed, "Distributed target classification and tracking in sensor networks," in *To appear in Proceedings of the IEEE*, 2003.
- [16] S. Brueckner, "Return from the ant - synthetic ecosystems for manufacturing control," Ph.D. dissertation, Humboldt University, 2000.
- [17] S. Brueckner and V. D. Parunak, "Multiple pheromones for improved guidance," in *Symposium on Advanced Enterprise Control*, 2000.
- [18] Z. J. Butler, A. A. Rizzi, and R. L. Hollis, "Cooperative coverage of rectilinear environments," in *Int. Conf. on robotics and Automation*, San Francisco, CA, April 2000.
- [19] Y. U. Cao, A. Fukunaga, and A. Kahng, "Cooperative mobile robotics: Antecedents and directions," *Autonomous Robots*, vol. 4, pp. 1–23, 1997.
- [20] S. Carpin, C. Ferrari, F. Montesello, E. Pagello, and P. Patuelli, "Scalable deliberative procedures for efficient multi-robot coordination," in *Proc. of ECAI2000 Workshop on Balancing Reactivity and Social Deliberation in Multi-Agent Systems*, 2000.
- [21] S. Carpin and L. E. Parker, "Cooperative leader following in a distributed multi-robot system," *Proceedings of IEEE International Conference on Robotics and Automation*, 2002.
- [22] —, "Cooperative motion coordination amidst dynamic obstacles," in *Distributed Autonomous Robotic Systems*, H. Asama, T. Arai, T. Fukuda, and T. Hasegawa, Eds. Springer, 2002, vol. 5, pp. 145–154.
- [23] C. Castelpietra, L. Iocchi, D. Nardi, M. Piaggio, A. Scalzo, and A. Sgorbissa, "Coordination among heterogenous robotic soccer players," in *Proc. of the Int. Conf. on Intelligent Robots and Systems (IROS'02)*, 2000.
- [24] L. Chaimowicz, T. Sugar, V. Kumar, and M. F. M. Campos, "An architecture for tightly coupled multi-robot cooperation," in *Proc. of IEEE Int. Conf. on Robotics and Automatio*, Seoul, Korea, may 2001, pp. 2292–2297.
- [25] F. Chantemargue and B. Hirsbrunner, "A collective robotics application based on emergence and self-organization," in *Proc. of Fifth Int. Conf. for Young Computer Scientists (ICYCS'99)*, 1999.
- [26] P. R. Cohen and H. J. Levesque, "Teamwork," *Special Issue on Cognitive Science and Artificial Intelligence.*, vol. 25, pp. 486–512, 1991.
- [27] A. Das, R. Fierro, V. Kumar, J. Ostrowski, J. Spletzer, and C. Taylor, "A framework for vision based formation control," *IEEE Transactions on Robotics and Automation*, vol. XX, 2001.
- [28] M. B. Dias and A. T. Stentz, "Opportunistic optimization for market-based multirobot control," in *IROS 2002*, September 2002.
- [29] B. Donald, L. Garipey, and D. Rus, "Distributed manipulation of multiple objects using ropes," in *Proceedings of the 2002 NRL Workshop on Multi-Robot Systems*, A. C. Shultz and L. E. Parker, Eds. Washington, DC: Kluwer Academic Publishers, 2002.
- [30] G. Dudek, M. Jenkin, and E. Milios, *A Taxonomy of Multirobot Systems*. AK Peters, 2002, ch. in Robot Teams: From Diversity to Polymorphism, Tucker Balch and Lynne E. Parker, eds.
- [31] G. Dudek, M. Jenkin, E. Milios, and D. Wilkes, "A taxonomy for multi-agent robotics," *Autonomous Robots*, vol. 3, no. 4, pp. 375–397, 1996.
- [32] E. Durfee, V. Lesser, and D. Corkill, "Trends in cooperative distributed problem solving," *IEEE Transactions on Knowledge and Data Engineering*, vol. KDE-1, no. 1, pp. 63–83, March 1989.
- [33] R. Emery, K. Sikorski, and T. Balch, "Protocols for collaboration, coordination and dynamic role assignment in a robot team," in *Proceedings of the 2002 IEEE Int. Conf. on Robotics and Automation*, Washington, DC, May 2002.

Coord. Dim.	Foraging	Box pushing	MTO	Exploration	Soccer
Unaware	[5][25][9][48][17][88]	[65][64][117][48]			[115]
Aware not Coord.		[66][105]		[7][11]	
Weakly Coord.	[41][111][112][102]	[113][101][52]		[6][91][76][27][89]	[40]
Strongly Coord. Strongly Centr.	[10][93]			[96][50][94][53][67]	[109][13]
Strongly Coord. Weakly Centr.	[73][74][28]	[80][97][24]		[42]	[79][118]
Strongly Coord. Distributed	[103][78][60][48][18] [59][72][85][20][58]	[54][1][87] [87][84][45][46]	[12][116] [86][56]	[83][100][29][3][104] [92][2][21][22]	[108][68][33] [23][14]

TABLE II
OVERVIEW OF THE CLASSIFIED MRS

- [34] D. Estrin, D. Culler, K. Pister, and G. Sukhatme, "Connecting the physical world with pervasive networks," *IEEE Pervasive Computing*, vol. 1, no. 1, pp. 59–69, 2002.
- [35] P. et al., *Cooperative Control for Autonomous Air Vehicles*. Dordrecht, Netherlands: Kluwer Academic Publishers, 2002, ch. 12.
- [36] A. Farinelli, L. Iocchi, and D. Nardi, "An analysis of coordination in multi-robot systems," in *IEEE Int. Conf. on Systems, Man & Cybernetics (to appear)*, Washington D. C., October 2003, pp. 1487–1492.
- [37] J. W. Fenwick, P. M. Newman, and J. J. Leonard, "Cooperative concurrent mapping and localization," in *Proceedings of the 2002 IEEE Int. Conf. on Robotics and Automation*, Washington, DC, May 2002.
- [38] J. Ferber, *Multi-Agent Systems*. Addison-Wesley, 1999.
- [39] F. Fernandez and L. E. Parker, "Learning in large cooperative multi-robot domains," *International Journal of Robotics and Automation*, vol. 16, no. 4, pp. 217–226, 2001.
- [40] M. Ferraresso, C. Ferrari, E. Pagello, R. Polesel, R. Rosati, A. Speranzon, and W. Zhanette, "Collaborative emergent actions between real soccer robots," in *RoboCup-2000: Robot Soccer World Cup IV*, 2000.
- [41] M. Fontan and M. Mataric, "Territorial multi-robot task division," *IEEE Transactions on Robotics and Automation*, vol. 15, no. 5, 1998.
- [42] T. Fukuda, I. Y., F. Arai, and T. Higashi, "Graduated spatial pattern formation of micro robot group," in *Distributed Autonomous Robot Systems*, H. Asama, T. Arai, T. Fukuda, and T. Hasegawa, Eds. Springer, 2002, vol. 5, pp. 340–348.
- [43] W. G., "Distributed reinforcement learning," *Robotics and Automation Systems*, vol. 15, pp. 135–142, 1995.
- [44] B. Gerkey and J. M. Mataric, "Multi-robot task allocation: Analyzing the complexity and optimality of key architectures," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA 2003)*, Taipei, Taiwan, Sep 14 - 19 2003.
- [45] B. P. Gerkey and M. J. Mataric, "Principled communication for dynamic multi-robot task allocation," in *Proceedings of the Int. Symposium on Experimental Robotics*, Waikiki, Hawaii, Dec. 2000.
- [46] —, "Pusher-watcher: An approach to fault tolerant tightly-coupled robot coordination," Technical report, Institute for Robotics and Intelligent Systems, University of southern California, March 2001.
- [47] D. Goldberg, V. Ciciello, M. Dias, R. Simmons, S. Smith, T. Smith, and A. Stentz, "A distributed layered architecture for mobile robot coordination: Application to space exploration," in *Proceedings of the 3rd Int. NASA Workshop on Planning and Scheduling for Space*, 2002.
- [48] D. Goldberg and M. J. Mataric, *Design and Evaluation of Robust Behavior-Based Controllers for Distributed Multi-Robot Collection Tasks*. AK Peters, 2002, ch. in Robot Teams: From Diversity to Polymorphism, Tucker Balch and Lynne E. Parker, eds.
- [49] B. Grosz and S. Kraus, "Collaborative plans for complex group actions," *Artificial Intelligence*, vol. 86, pp. 269–358, 1996.
- [50] D. Guzzoni, A. Cheyer, L. Julia, and K. Konolige, "Many robots make short work," *AI Magazine*, 1997.
- [51] C. H., "Coverage for robotics - a survey on recent results," *Annals of Mathematics and Artificial Intelligence*, vol. 31, pp. 113–126, 2001, kluwer academic Publishers.
- [52] Y. Hirata, K. Kosuge, H. Asama, H. Kaetsu, and K. Kawabata, "Coordinated transportation of a single object by multiple mobile robots without position information of each robot," in *Proc. of the Int. Conf. on Intelligent Robots and Systems (IROS'00)*, 2000.
- [53] H. Howard, M. Mataric, and G. S. Sukhatme, "An incremental deployment algorithm for mobile robot teams," in *Proc. of the Int. Conf. on Intelligent Robots and Systems (IROS'00)*, 2002.
- [54] A. J. Ijspeert, A. Martinoli, A. Billard, and L. M. Gambardella, "Collaboration through the exploitation of local interactions in autonomous collective robotics: The stick pulling experiment," *Autonomous Robots*, vol. 11, no. 2, pp. 149–171, 2001, kluwer Academic Publishers.
- [55] L. Iocchi, D. Nardi, and M. Salerno, "Reactivity and deliberation: a survey on multi-robot systems," in *Balancing Reactivity and Deliberation in Multi-Agent Systems (LNAI 2103)*, E. P. M. Hannebauer, J. Wendler, Ed. Springer, 2001, pp. 9–32.
- [56] M. Jenkin and G. Dudek, "The paparazzi problem," in *Proc. of the Int. Conf. on Intelligent Robots and Systems (IROS'00)*, 2000.
- [57] N. Jennings, "Controlling cooperative problem solving in industrial multi-agent systems using join intentions," in *Artificial Intelligence*, 1995.
- [58] D. Jung and A. Zelinsky, "An architecture for distributed cooperative planning in a behaviour-based multi-robot system," *Journal of Robotics and Autonomous Systems* 26, pp149-174, 1999.
- [59] —, "Grounded symbolic communication between heterogeneous cooperating robots," *Autonomous Robots*, vol. 8, no. 3, pp. 269–292, 2000.
- [60] A. Khoo and D. Horswill, "An efficient coordination architecture for autonomous robot teams," in *Proc. of the Int. Conf. on Intelligent Robots and Systems (IROS'02)*, Washington, DC, May 2002.
- [61] H. Kitano, M. Asada, Y. Kuniyoshi, I. Noda, E. Osawa, and H. Matsubara, "Robocup: A challenge problem for ai and robotics," in *Lecture Note in Artificial Intelligence*, vol. 1395, 1998, pp. 1–19.
- [62] H. Kitano and S. Tadokoro, "RoboCup-Rescue: A grand challenge for multi-agent and intelligent systems," *AI Magazine*, vol. 22, no. 1, pp. 39–52, 2001.
- [63] K. Konolige et al., "Centibots: Large scale robot teams," in *Proceedings of the 2003 NRL Workshop on Multi-Robot Systems*, A. C. Shultz and L. E. Parker, Eds., vol. II. Washington, DC: Kluwer Academic Publishers, 2003.
- [64] C. Kube, H. Zhang, and X. Wang, "Controlling collective tasks with an ALN," in *IEEE/RSJ IROS*, pages 289–293, 1993.
- [65] C. R. Kube and E. Bonabeau, "Cooperative transport by ants and robots," *Robotics and Autonomous Systems*, vol. 30, no. 1, pp. 85–101, 2000.
- [66] C. R. Kube and H. Zhang, "Collective robotic intelligence," in *Int. Conf. on Simulation of Adaptive Behaviour*, 1992, pp. 460–468.
- [67] R. Kumar and J. A. Stover, "A behavior-based intelligent control architecture with application to coordination of multiple underwater vehicles," *IEEE Transactions on Systems, Man, and Cybernetics—Part A: Cybernetics*, vol. 30, no. 6, pp. pages 767–784, Nov. 2000.
- [68] T. C. Liang and J. S. Liu, "Motion control realizing cyclic ball passing strategy among multiple mobile robots in robot soccer games," in *Proceedings of the 2002 IEEE Int. Conf. on Robotics and Automation*, Washington, DC, May 2002.
- [69] M. J. Mataric, "Interaction and intelligent behavior," Ph.D. dissertation, MIT, 1994.
- [70] —, "Designing and understanding adaptive group behavior," *Adaptive Behavior*, vol. 4, no. 1, 1995.
- [71] —, "Issues and approaches in the design of collective autonomous agents," *Robotics and Autonomous Systems*, vol. 16, no. 2-4, pp. 321–331, 1995.
- [72] —, "Learning social behavior," *Robotics and Autonomous Systems*, vol. 20, pp. 191–204, 1997.
- [73] D. Miller, "Multiple behavior-controlled microrobots for planetary surface missions," in *IEEE Int. Conf. on Systems, Man, and Cybernetics*, Los Angeles, USA, 1990, pp. 289–292.

- [74] D. P. Miller, C. Stein, A. Wright, and R. Sargen, "Are you being served?" in *Proceedings of the AAAI-97/IAAI-97 Conf.*, 1997.
- [75] J. P. Modi, H. Jung, W. M. Tambe, M. Shen, and S. Kulkarni, "A dynamic distributed constraint satisfaction approach to resource allocation," in *Proceedings of Seventh Int. Conf. on Principles and Practice of Constraint Programming*, 2001.
- [76] S. Monteiro and E. Bicho, "A dynamical systems approach to behavior-based formation control," in *Proceedings of the 2002 IEEE Int. Conf. on Robotics and Automation*, Washington, DC, May 2002.
- [77] R. Murphy, J. G. Blitch, and J. L. Casper, "Robocup/aaai urban search and rescue events: Reality and competition," *AI Magazine*, In press.
- [78] R. R. Murphy, C. Lisetti, R. Tardif, L. Irish, and A. Gage, "Emotion-based control of cooperating heterogeneous mobile robots," in *IEEE Transactions on Robotics and Automation*, 2002, special issue on Multi-robot Systems.
- [79] B. Nebel, J. S. Gutmann, and W. Hatzack, "CS Freiburg '99," in *RoboCup-99: Robot Soccer World Cup III*, 1999.
- [80] F. R. Noreils, "Toward a robot architecture integrating cooperation between mobile robots: Application to indoor environment," *Int. Journal of Robotics Research*, vol. 12, no. 1, pp. 79–98, 1993.
- [81] L. E. Parker, "Current state of the art in multi-robot teams," in *Distributed Autonomous Robotic Systems 4*. Springer, 2000, pp. 3–12.
- [82] —, "The effect of heterogeneity in teams of 100+ mobile robots," in *Proceedings of the 2003 NRL Workshop on Multi-Robot Systems*, A. C. Shultz and L. E. Parker, Eds., vol. II. Washington, DC: Kluwer Academic Publishers, 2003.
- [83] L. E. Parker, K. Fregene, Y. Guo, and R. Madhavan, "Distributed heterogeneous sensing for outdoor multi-robot localization, mapping and path planing," in *Proceedings from the 2002 NRL Workshop on Multi-Robot Systems*, A. C. Shultz and L. E. Parker, Eds. Washington, DC: Kluwer Academic Publishers, 2002.
- [84] L. E. Parker, Y. Guo, and D. Jung, "Cooperative robot teams applied to the site preparation task," in *Proceedings of 10th Int. Conf. on Advanced Robotics*, 2001, pp. 71–77.
- [85] L. E. Parker, "ALLIANCE: An architecture for fault tolerant multirobot cooperation," *IEEE Transactions on Robotics and Automation*, vol. 14, no. 2, pp. 220–240, 1998.
- [86] —, "Cooperative robotics for multi-target observation," *Intelligent Automation and Soft Computing*, vol. 5, no. 19, 1999.
- [87] —, "Lifelong adaption in heterogeneous multi-robot teams: Response to continual variation in individual robot performance," *Autonomous Robots*, vol. 8, no. 3, pp. 239–267, 2000.
- [88] K. Passino, "Biomimicry of bacterial foraging for distributed optimization and control," *IEEE Control Systems Magazine*, vol. 22, no. 3, pp. 52–67, June 2002.
- [89] M. Polycarpou, Y. Yang, and K. Passino, "A cooperative search framework for distributed agents," in *Proceedings of the 2001 IEEE International Symposium on Intelligent Control*, 2001, pp. 1–6.
- [90] D. V. Pynadath and M. Tambe, "The communicative multiagent team decision problem: Analyzing teamwork theories and models," *Journal of Artificial Intelligence Research*, vol. 16, pp. 389–423, 2002.
- [91] M. Quinn, "A comparison of approaches to the evolution of homogeneous multirobot teams," in *Proc. Int. Congress on Evolutionary Computing*, 2001.
- [92] I. Rekleitis, G. Dudek, and E. Milios, "Multi-robot collaboration for robust exploration," in *Proceedings of Int. Conf. in Robotics and Automation*, 2000.
- [93] P. Rybski, S. Stoeter, C. Wyman, and L. M. Gini, "A cooperative multi-robot approach to the mapping and exploration of mars," in *Proceedings of AAAI/IAAI-97*, Providence, 1997.
- [94] A. Saffiotti, N. Zumel, and E. Ruspini, "Multi-robot team coordination using desirabilities," in *Proc of the 6th Intl Conf on Intelligent Autonomous Systems (IAS)*, Venice, Italy, 2000, pp. 107–114.
- [95] S. Sen, "Adaptation, coevolution and learning in multiagent systems: Papers from the 1996 aaai spring symposium," AAAI Press, Menlo Park, CA, AAAI Technical Report SS-96-01, March 1996.
- [96] R. Simmons, D. Apfelbaum, D. Fox, R. P. Goldman, K. Z. Haigh, D. J. Musliner, M. Pelican, and S. Thrun, "Coordinated deployment of multiple, heterogeneous robots," in *Proc. of the Int. Conf. on Intelligent Robots and Systems (IROS'00)*, 2000.
- [97] R. Simmons, S. Singh, D. Hershberger, J. Ramos, and T. Smith, "First results in the coordination of heterogeneous robots for large-scale assembly," in *Proceedings of the Int. Symposium on Experimental Robotics (ISER)*, Honolulu Hawaii, December 2000.
- [98] L. Steel, "Toward a theory of emergent functionality," in *First Int. Conf. on Simulation of Adaptive Behavior*. MIT Press, 1990, pp. 451–461.
- [99] P. Stone, *Layered Learning in Multiagent Systems*. MIT Press, 2000.
- [100] A. W. Stroupe and T. Balch, "Mission relevant collaborative observation and localization," in *Proceedings from the 2002 NRL Workshop on Multi-Robot Systems*, A. C. Shultz and L. E. Parker, Eds. Washington, DC: Kluwer Academic Publishers, 2002.
- [101] T. Sugar, J. P. Desai, V. Kumar, and J. P. Ostrowski, "Coordination of multiple mobile manipulators," in *Proc. of IEEE Int. Conf. on Robotics and Automation*, Seoul, Korea, May 2001.
- [102] K. Sugawara and M. Sano, "Cooperative behavior of interacting simple robots in a clockface arranged foraging field," in *Distributed Autonomous Robotic Systems*, H. Asama, T. Arai, T. Fukuda, and T. Hasegawa, Eds. Springer, 2002, vol. 5, pp. 331–339.
- [103] S. Sun, D. Lee, and K. Sim, "Artificial immune-based swarm behaviors of distributed autonomous robotic systems," in *Proceedings of IEEE Int. Conf. of Robotics and Automation (ICRA 2001)*, 2001, pp. 3993 – 3998.
- [104] K. Takahashi and M. Kakikura, "Research on cooperative capture by multiple mobile robots - a proposition of cooperative capture strategies in the pursuit problem-," in *Distributed Autonomous Robotic Systems*, H. Asama, T. Arai, T. Fukuda, and T. Hasegawa, Eds. Springer, 2002, vol. 5, pp. 393–402.
- [105] H. Takeda, Y. Hirata, Z. D. Wang, and K. Kosuge, "Collision avoidance algorithm for two tracked mobile robots transporting a single object in coordination based on function allocation concept," in *Distributed Autonomous Robotic Systems*, H. Asama, T. Arai, T. Fukuda, and T. Hasegawa, Eds. Springer, 2002, vol. 5, pp. 155–164.
- [106] M. Tambe, "Towards flexible teamwork," *Journal of Artificial Intelligence Research (JAIR)*, vol. 7, pp. 83–124, 1997.
- [107] C. Tomlin, G. J. Pappas, and S. Sastry, "Conflict resolution for air traffic management: A study in multi-agent hybrid systems," *IEEE Transactions on Automatic Control*, vol. 43, no. 4, April 1998.
- [108] E. Uchibe, T. Kato, M. Asada, and K. Hosoda, "Dynamic task assignment in a multiagent/multitask environment based on module conflict resolution," in *Proc. of the Int. Conf. on Intelligent Robots and Systems (IROS'01)*, May 2001, Seoul, Korea.
- [109] M. Veloso and P. Stone, "Individual and collaborative behaviors in a team of homogeneous robotic soccer agents," in *Proceedings of the Third Int. Conf. on Multi-Agent Systems*, 1998, pp. 309–316.
- [110] —, *A Survey of Multiagent and Multirobot Systems*. AK Peters, 2002, ch. in *Robot Teams: From Diversity to Polymorphism*, Tucker Balch and Lynne E. Parker, eds.
- [111] I. Wagner and A. Bruckstein, "Cooperative cleaners: a study in ant robotics," Technion, IIT, Haifa, Tech. Rep. 9512, 1995.
- [112] I. A. Wagner, M. Lindenbaum, and A. M. Bruckstein, "Mac vs. pc - determinism and randomness as complementary approaches to robotic exploration of continuous unknown domains," *Int. Journal of Robotics Research*, vol. 19, no. 1, pp. 12–31, 2000.
- [113] Z. D. Wang and V. Kumar, "A decentralized test algorithm for object closure by multiple cooperating robots," in *Distributed Autonomous Robotic Systems*, H. Asama, T. Arai, T. Fukuda, and T. Hasegawa, Eds. Springer, 2002, vol. 5, pp. 165–174.
- [114] G. Weiß and S. Sen, Eds., *Adaptation and Learning in Multi-Agent Systems*, ser. Lecture Notes in Artificial Intelligence. Berlin, Germany: Springer-Verlag, 1996, vol. 1042.
- [115] B. B. Werger, "Cooperation without deliberation: A minimal behavior-based approach to multi-robot teams," *Artificial Intelligence*, vol. 110, pp. 293–320, 1999.
- [116] B. B. Werger and M. J. Mataric, "Broadcast of local eligibility for multi-target observation," in *Proc. of 5th Int. Symposium on Distributed Autonomous Robotic Systems (DARS)*, 2000.
- [117] X. Yang, K. Watanabe, K. Kiguchi, and K. Izumi, "Coordinated transportation of a single object by a group of nonholonomic mobile robots," in *Distributed Autonomous Robotic Systems*, H. Asama, T. Arai, T. Fukuda, and T. Hasegawa, Eds. Springer, 2002, vol. 5, pp. 175–184.
- [118] K. Yokota, K. Ozaki, N. Watanabe, A. Matsumoto, D. Koyama, T. Ishikawa, K. Kawabata, H. Kaetsu, and H. Asama, "Utori united: Cooperative team play based on communication," in *RoboCup-98: Robot Soccer World Cup II*, 1998.