

Reactivity and Deliberation: a survey on Multi-Robot Systems

Luca Iocchi, Daniele Nardi, Massimiliano Salerno

Dipartimento di Informatica e Sistemistica Università “La Sapienza”
Via Salaria 113 00198, Roma, Italy.
E-mail: {iocchi,nardi,salerno}@dis.uniroma1.it

Abstract. Multi Robot Systems are, nowadays, an important research area within Robotics and Artificial Intelligence. Although Multi Robot Systems can be regarded as a particular case of Multi Agent Systems, it seems appropriate to study the Multi Robot Systems from a specific viewpoint, because of the issues which arise from the embodiment of agents operating in real environments. In this paper, we present an analysis of Multi Robot Systems by looking at their cooperative aspects. In particular, we propose a taxonomy of Multi Robot Systems and a characterization of reactive and social deliberative behaviors of the Multi Robot System as a whole. Finally, we address some Multi Robot Systems, which we consider representative of the various nodes in our taxonomy.

1 Introduction

Multi Robot Systems (MRS) have been proposed in the last decade in a variety of settings and applied in several tasks. Special attention has been given to MRS developed to operate in a dynamic environment, where uncertainty and unforeseen changes can happen due to robots and other agents that are external to the MRS itself.

In this paper, we regard MRS as a particular form of Multi Agent System (MAS), by specifically addressing reactivity and social deliberation. The work on MRS is pursued based on rather different goals. A significant body of work on MRS has been originated from motivations that are essentially of engineering nature. A 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 modularizing it. Let us first address the improvement of performance. In [14] it is emphasized that there are tasks that require multi robots to be accomplished. This situation can arise not only when the robots can accomplish different functions, but also when they have the same capabilities. 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. With regard to robustness, following [40], we emphasize that, MRS can be designed and implemented in such a way to guarantee two important features: adaptivity and fault tolerance. Adaptivity refers to the ability of the MRS to modify its own behavior over time, depending

on changes due to the dynamic environment, changes in the system mission, changes in the system composition or capabilities, so that the performance of the entire system can either be improved or at least not degraded. Fault tolerance is the ability of the MRS to deal with individual robot failures or communication failures, between two or more robotic agents, that may occur at any time during the mission. MRS which present both the above features can be said robust, where robustness is, therefore, the ability of the MRS to be both adaptive and fault tolerant.

Another significant development of MRS stems from the studies on biological systems or complex models arising in cognitive science and economics (see for example [9]). In this framework the development of simulation models has led also to implementations of real robot systems, built to provide additional evidence to the behaviors studied in the simulations. Many of these systems deal with a large number of robotic agents and the MRS is thus viewed as a swarm [9, 40], or a colony [2] or, more generally, as a robot collective [14]. While in this work we are more interested in the system engineering aspects of the MRS, we consider to some extent also a few biologically inspired approaches.

A significant boost to the work on MRS has recently been given also by the Robotics competitions and RoboCup [27], in particular. In fact, the design of MRS is regarded as one of the major scientific challenges to be developed in the RoboCup environment [3]. In particular, according to the league organization of RoboCup, the real robot leagues as well as the simulated league provide different settings, where different hypotheses underly the design of the MRS. The most distinguishing feature, as compared with previous work on MRS, is that the RoboCup environment is highly dynamic and includes an opponent team; thus the RoboCup setting provides significantly new challenges in the design of MRS.

While the above mentioned motivations may seem very far apart, the technical solutions that are being developed have several common features. However, the discipline is still relatively young and a common framework for the work on MRS is difficult to identify. A MRS cannot be simply regarded as a generalization of the single robot case and the proposed approaches need to be more precisely characterized in terms of assumptions about the environment and in terms of the internal system organization.

Interesting contributions towards a classification of the work on MRS are the surveys by Cao et al. [9] and by Dudek et al. [14]. In [9] several axes for characterizing a MRS are discussed, while in [14] a classification of MRS that is more focussed with the communication and computation aspects of MRS is presented. In addition, MRS can be regarded as MAS and thus many aspects of the system organization of MRS can be analyzed by referring to the literature about MAS.

The aim of the present work is to address the most recent developments of MRS by classifying the proposed approaches in terms of a number of features concerning the system organization and specifically focusing on reactivity and social deliberation.

The paper is organized as follows. In the next section we propose a classification structure that is suitable for characterizing the major features of MRS. In section 3 we specifically focus on the issue of reactivity versus social deliberation. In section 4 we present some of the approaches that we consider as representative of various nodes in the taxonomy, characterizing them according to the task accomplished by the MRS. Section 5 concludes the paper by providing some reflections on the problems to be addressed by future research.

2 Multi Robot System Taxonomy

In this section we present a taxonomy of the Multi Robot Systems (MRS) centered on their cooperative capabilities. The taxonomy allows for a precise classification of different typologies of MRS, with a special interest to those design choices concerning cooperation within the MRS. Following the approach in [9], we propose several classification dimensions. The *classification dimensions* used in the taxonomy are defined via a top down approach by refining the level of the system structure characterization. Four different levels are individuated in the taxonomy and presented in detail: a *Cooperation Level*, a *Knowledge Level*, a *Coordination Level* and an *Organization Level*. Subsequently, we recall two additional dimensions which can characterize the systems, independently of their classification in the taxonomy: communication and system composition. The issue of reactivity versus deliberation is addressed in the following section. In addition, an analysis of some of the most relevant works on the MRS according to our taxonomy will be presented in section 4.

The taxonomy proposed for classification of MRS is shown in **Fig. 1**. The first level of the taxonomy 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.

Cooperation: situation in which several robots operate together to perform some global task that either cannot be achieved by a single robot, or whose execution can be improved by using more than one robot, thus obtaining higher performances (see [39]).

In this work we are interested only in those MRS which present many robotic agents (we also call them a *team*) having, as a whole, a common global goal to achieve. As a consequence, systems consisting of robotic agents that operate in the same environment, but that are not related to each other by a common objective (called *not cooperative* MRS) are not taken into account in this article. Therefore, in the following, the term MRS will refer to a team of cooperative robots.

Among the cooperative systems a first important characterization can be introduced depending on how much knowledge each robot has about the presence of other robots in its own team. This dimension represents the second level of the proposed taxonomy, that is called *Knowledge Level*.

Fig. 1. MRS Taxonomy

Awareness: the property of a robot in the MRS to have knowledge of the existence of the other members of the system [47].

Conversely, unaware robots perform their tasks as if they were the only robots present within the system. Cooperation among unaware robotic agents is the weakest form of cooperation. For example in a box-pushing task many robots can contribute to the achievement of the common goal while behaving as single entities, i.e. without taking into account the presence and the actions of the other robots, as in [28].

In the category of aware systems we can identify the modes through which cooperation among robots can be achieved. We therefore introduce a third level in our taxonomy: the *Coordination Level*, that is concerned with the mechanisms used for robot cooperation.

Coordination: 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 (see [15]).

Coordination is not a prerogative of the cooperative MRS, in fact there exist robotic systems in which coordination between the members is required, but the robots have different goals which often are not related to each other. For example, the industrial robots often share tools and coordination is needed to avoid

interferences, which take place when a single indivisible resource is requested by multiple robots. These kind of systems are not considered in the present article, which is concerned with cooperative MRS.

It is obvious that systems of unaware robots cannot be coordinated in the above sense because their members do not have knowledge of the existence of the others.

There are different ways a robot can take into account the actions of the other members of the team; we can further classify the coordinated MRS by introducing a distinction between two different kinds of coordination based on the use of a *coordination protocol*, that is defined as a set of rules that the robots must follow in order to interact each other in the environment.

Strong (Weak) coordination: a form of coordination that relies (does not rely) on a coordination protocol.

Strongly coordinated systems are based on a system of signals by which a robot exerts its influence on the behavior of another. In other words, strong coordination is based on the application of predefined or learned rules concerning the way two or more robots have to interact.

Coordination is not always a necessary property for the cooperative MRS, in fact while it introduces more flexibility within the system organization by allowing for a better use of the available resources, there are tasks which can be accomplished effectively by not coordinated MRS, as in [6] where a formation maintaining task is addressed. The advantage of a not coordinated MRS relies in its less complex design which involves a lower risk of fault, but as a consequence there could be a large waste of resources either because the robots execute contrasting tasks, or because interferences arise. A coordinated MRS, on the contrary, can avoid, or at least reduce, the previous kind of problems, but a more complex design is needed. The weakly coordinated MRS, as in [16], are not dependent on a protocol and consequently more robust to communication failures, but they lack of many of the organizational capabilities which a coordination protocol offers. The more the environment is dynamic and the goal is complex, the more a strongly coordinated MRS is effective in achieving its goal.

The fourth level of our taxonomy is concerned with the way the decision system is realized within the MRS. The *Organization Level* permits a distinction between centralized approaches and distributed ones [9].

Centralization: the organization of a system having a robotic agent (a leader) that is in charge of organizing the work of the other robots; the leader is involved in the decisional process for the whole team, while the other members act according to the directions of the leader.

Distribution: the organization of a system composed by robotic agents which are completely autonomous in the decisional process with respect to each other; in this class of systems a leader does not exist.

A centralized system can have a hierarchical structure, in which the robots operating under control of a leader, can be themselves leaders of other sub team of robots within the multi robot system.

A further classification of centralized systems can be introduced depending on the way the leader role is played.

Strong centralization: centralization in which decisions are taken by a leader that remains the same during the entire mission duration.

It is possible that more than one member in the MRS is allowed to acquire the role of leader, all of them can potentially plan the activity of the other members, but, when the MRS is strongly centralized, at the beginning of the mission the leader role is univocally assigned to only one of the robots, and kept by it till the completion of the mission itself.

Weak centralization: centralization in which more than one robot is allowed to become a leader during the mission.

The role of leader is assigned dynamically during the accomplishment of the mission either depending on environmental changes, or forced by the failure of the current robot leader.

Centralization, both weak and strong, allows for a simpler task assignment between the team member because only one of them, the leader, is in charge of it. Centralized MRS, as the one proposed in [49], have the disadvantage that they strongly rely on communication. Thus, when a communication failure take place, it results in a failure of the entire system. Moreover, a strongly centralized system can fail in accomplishing its task when its leader goes out of order; a weakly centralized system tries to recover from a leader failure by selecting a new leader, as in [39]. An advantage of the strongly centralized MRS is that a well suited robotic agent can be realized to be the leader, for example by having the appropriate computing capabilities to analyze the environmental data used to take decision. On the contrary, the distributed MRS, as in [42], offer a greater robustness to the above mentioned problems by letting each team member decide autonomously, but a greater complexity is required to realize coordination between them.

As already mentioned, along with the classification introduced by the taxonomy, two more dimensions can be identified which are orthogonal to the previous ones: communication and system composition.

Cooperation among robots is often obtained by a communication mechanism that allows the robots to exchange messages among each other. For a detailed analysis of the various technical problems related to communication see [15]. Here we limit ourselves to present few basic aspects about communication, so that its connection with our taxonomy can be illustrated. In fact, strongly coordinated systems must necessarily communicate (in order to execute the coordination protocol), weakly coordinated and not coordinated systems can be developed with or without communication, while unaware systems do not use communication (because any robot does not know of the existence of the others).

We can distinguish two different types of communication depending on the way the robots exchange information. All the communicating systems can use either direct or indirect communication.

Direct communication: communication that makes use of some hardware on board dedicated device to signal something that the other team members can understand.

Indirect communication: communication that makes use of stigmergy¹, both active and passive, for the communication among the team members.

Communication is a very important issue in MRS because many property of the system itself rely on it. Direct communication is the easier and most powerful way to exchange information between the members of a MRS, but due to noises or to hardware failures it often becomes a critical point. To avoid this problem, techniques like stigmergy have been developed and applied, as in [28], so that the members of the system can deal with communication failures. This more robust communication system implies a more complicated MRS design because each member of the team has to interpret the surrounding environment and, moreover, less information can be exchanged between the members themselves.

Finally, a design choice that must be taken into account for developing a MRS is the type of robots composing the system. We can distinguish two different system compositions: heterogeneous, homogeneous (see for example [46]).

Homogeneity: property of a team of robots whose members are exactly the same both in the hardware and in the control software.

Heterogeneity: property of a team of robots whose members have a difference either in the hardware devices or in the software control procedures.

If the robots of a deterministic homogeneous system are subject to the same inputs they have the same behavior, by producing the same outputs. They operate in a different way only when they act under different conditions in the environment in which they operate, or if the system is not deterministic. On the contrary, heterogeneous robots usually have different control mechanisms, different representations of the information, or different behaviors due to the experiences collected during the system existence via learning techniques.

The homogeneity or the heterogeneity of the members can influence the way in which robustness is achieved. In a homogeneous system every agent can execute the same actions as the other team members with the same results, so

¹ Stigmergy is a term coined by the biologist P.Grassé which means to incite work by the effect of previous work [28], in the robotics field there exist two type of stigmergy: active and passive. The first kind occurs when a robot alter the environment so as to affect the sensory input of another robot, the second one occurs when a robot action alter the environment so that the effects of another robot's action change.

that a failure of a member can be easily compensated by another robot in the system, in such a way that fault tolerance is guaranteed. On the contrary adaptivity is weakly ensured because there is no differentiation among the robots. Therefore, homogeneous systems are typically used in the so called swarm-type approach, where a large number of autonomous robots are utilized to realize a distributed system. This method allows the system to achieve the desired goals when the repetitive execution of a set of actions is needed and time is not a critical resource. The accomplishment of the task usually relies on some mathematical convergence results, and the relevant question with these systems is the realization of the correct laws to ensure the achievement of the desired goal.

On the other hand, for assuring adaptivity without giving up fault tolerance, heterogeneous MRS are preferred. Heterogeneous robots allow the MRS to adapt more easily to the different situations which could emerge in a dynamic environment, since they offer a better chance to deal with new and unpredicted needs. The system does not rely on the fact that the robots are all the same, thus making it easier to modify the system composition so as to support a team configuration change by gathering new members within the MRS itself and/or changing some of the existing ones. Moreover, in a heterogeneous system it is still possible for more than one robot to execute the same task, even though with different modalities (because of different hardware and/or software). In this way, when a failure occurs, another member of the MRS can perform the required work to keep achieving the task, thus ensuring fault tolerance. Obviously, working with heterogeneous systems implies a higher effort in realizing the software needed to control the MRS.

3 Social Deliberation vs. Reactivity

Deliberation and reactivity are important features of a robotic system: both in the case of a single robot and in a robotic team. In this section we first recall the architectural issues that have been addressed for a single robot in order to develop a reactive and/or a deliberative behavior. Then, we characterize which behavior is expected from a reactive or a deliberative MRS. Finally, we discuss reactivity and social deliberation with respect to the taxonomy introduced in the previous section.

3.1 Deliberation and Reactivity for a Single Robot

The problem of realizing a reactive and/or a deliberative behavior in a single robot has been addressed by introducing different architectures for the robotic agent, specifically developed to suite the desired behavior. Three main approaches have been largely considered in the literature [33]: the behavior-based, the sense-model-plan-act and the hybrid one.

Behavior-Based Architecture. The subsumption architecture presented in [8] has been the first example of this kind of approach. A behavior-based controller consists of a collection of behaviors that activate and/or maintain goals [33]. Each

behavior is implemented substantially as software or hardware control laws receiving inputs from the robot sensors and/or from other behaviors in the system, and sending outputs to the robot actuators and/or to other behaviors. Thus, the resulting architecture consists in a network of interacting behaviors, whose activation depends on a set of rules (behavior arbitration), such as prefixed priority among behaviors or a voting mechanism. Reactivity is ensured in these systems because of their quick real time responses due to the embedment of a collection of preprogrammed, concurrent condition-action rules. Finally, a behavior-based system is realized in a bottom up fashion and a distributed internal representation of the robot state can be used for building more and more abstract behaviors by laying them one on top of the other.

Sense-Model-Plan-Act Architecture. This approach is based on the explicit representation of the robot internal state and on the realization of four steps which give the architecture the name of sense-model-plan-act architecture [38]. The first step is called sensing and consists in collecting sensor data of the environment where the robot is situated; the data collected are then used to build a model of the environment which will be different depending on the way the system is implemented and on the aspect that are relevant to decide which action to execute. The model is usually constituted by a set of symbols composed by predicates and values which can be manipulated by a logical system, for example propositional or first-order logic. A planner module is in charge of analyzing the information contained in the environment model, by deciding a sequence of actions to be executed to achieve the desired goal, and verifying through logical inference if the sequence itself is compatible with the knowledge base of the system. Once a suitable sequence, called plan, has been found, the high level commands must be converted in low level commands to be sent to the actuators so that the robot can actually execute the plan. The deliberative system is realized in a top down fashion. This architecture is suited to deal with complex situations, but the need of higher computational resources (due to the modelling phase and the planning one), causes severe limitations in the response time (and thus in the reactivity) of the system.

Hybrid Architecture. To use the advantages of both the behavior-based architecture and the deliberative one, several hybrid architectures have been proposed (see for example [17, 18]). This approach is realized by integrating in the same architecture: 1) behaviors that control the low level functions and that are in charge of the reactivity of the system, 2) a decisional module (a planner) in charge of the high level decision making. The hybrid architecture is often composed by two different layers: the planner and the reactive system. Sometimes an intermediate module is used for reconciling the possibly different representations used by the two layers and solving any conflict between the outputs.

The pure sense-model-plan-act architecture, which is used to realize a high level deliberative behavior, is not currently used because of its intrinsic limits,

while the behavior-based and the hybrid architectures are quite common, especially when the robot is situated in a highly dynamic environment, where a quick reaction to a new input is very important, being the environment itself uncertain and unpredictable.

3.2 Social Deliberation and Reactivity in MRS

A precise characterization of MRS with respect to reactivity and social deliberation is difficult to identify in the literature on MRS, although these characteristics can be found in several MRS and are addressed for example in [10]. In addition, in the case of MAS reactivity and social deliberation are addressed either through the study of a specific agent architecture (see for example [34, 30]), or by considering two different kinds of cooperation called reactive and intentional, as described in [15].

We propose to address these notions based on the ability of the team of robots to act in response to the dynamics of the environment. We thus characterize social deliberation and reactivity for MRS as follows.

MRS social deliberation: a system behavior that allows the team to cope with the environmental changes by providing a strategy that can be adopted to reorganize the team members' tasks, so as to use all the resources available to the system itself to effectively achieve the global goal.

MRS reactivity: a system behavior in which every single robot in the team copes with the environmental changes by providing a specific solution to reorganize its own task in order to fulfill the accomplishment of its originally assigned goal.

The main difference between social deliberation and reactivity in the proposed definition relies on the different modalities applied by the MRS to recover from an unpredicted situation: in a social deliberative MRS a long term plan involving the usage of all the available resources to 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. In a social deliberative system a global representation of the environment that is shared by all the team members may exist, but it is not required. In fact, a system can be deliberative without a global reconstruction of the environment, by imposing some constraints on the overall system behavior. While a reactive system architecture allows for the implementation of MRS that are able to quickly react to environmental changes, without affecting the entire MRS organization.

To clarify this distinction we can consider a case where the MRS composed, for instance by three robots, has the goal to open a room door and to move the robots into the room. Let us suppose that all the robots can accomplish the same tasks, even if in different ways. One of the robots has the specific task to open the door and then enter the room, while the others have the task to go directly to the door and, when open, enter the room. If the first robot encounters

an obstacle two different things can happen depending on the type of system organization.

In a reactive system the robot that encountered the obstruction tries to avoid it, so that it can continue the execution of its original task without affecting the behavior of the other team members. If it fails to bypass the obstacle, the other robots would remain in front of the closed door, until one of them tries to exit from the impasse by opening the door by itself, but without a global task redistribution.

In a social deliberative system the same situation would be faced by deciding to reorganize the entire mission for example by exchanging the task allocation among the robots, in any case by elaborating a new strategy involving several members of the team.

3.3 Reactivity and Social Deliberation in the Proposed Taxonomy

As previously argued, when looking at reactivity versus deliberation, there is the need to distinguish between the case in which these concepts are referred to a MRS and that in which they are referred to a single robot. In fact, the former can be either social deliberative or reactive independently of the kind of architecture used to realize the single robots. In the previous example the way the robot reorganizes its own task to open the door is irrelevant with respect to its behavior as a member of the system. The single robot architecture could be either behavior-based, or hybrid, without affecting the reactive or social deliberative behavior of the entire system. On the contrary, in order to properly characterize the system, it is necessary to consider other features different from the single robot architecture.

As shown in **Fig. 2**, in order to have social deliberation in a MRS it is necessary to consider whether a system is coordinated or not, in particular, whether it is strongly coordinated or not. As evidenced by the taxonomy all the systems that are not strongly coordinated are reactive, while the strongly coordinated ones can be either reactive or social deliberative. This happens because of the need to have a coordination protocol for reorganizing the MRS task allocation. Among the strongly coordinated MRS only the distributed ones can be reactive, because centralization implies that the task allocation among the different team members is executed by the robotic agent which takes the role of the leader, thus enforcing the behavior of the overall system, by providing a strategy.

As a consequence, among the orthogonal dimensions of communication and system composition, only communication is a required feature for social deliberative systems which rely on strong coordination, while composition of the team is not a constraining characteristic.

4 MRS application domain

In this section we are going to discuss several works related to MRS. We collect them according to the task used to test the techniques proposed, so that a

Fig. 2. MRS Taxonomy

comparison among them is more natural. It is worth emphasizing that our goal is not to provide a complete survey of existing approaches. Rather, we aim at characterizing more precisely some aspects of the taxonomy presented in Section 2, by providing concrete examples of implemented MRS.

We consider five different kinds of tasks for MRS: foraging, multi target observation, box pushing, exploration and soccer. These domains are representative of a large number of issues which arise from the study of cooperative MRS.

4.1 Foraging

Foraging is a testbed often used for MRS because of its analogies with tasks like rescue and search operations, toxic waste cleaning, mine cleaning [4, 24, 32, 40, 12]. The task requires the MRS's robotic agents to pick up objects scattered in the environment; along with the foraging task also the multi foraging has been defined [4], which is a variation of the first one where different kinds of objects must be collected. Different techniques to cope with problems like interferences (especially when a homogeneous MRS is used), communication and knowledge sharing are involved in this task.

In [4] an unaware, reactive MRS, composed by behavior based robotic agents, is proposed to deal with a multi foraging task, in which two typology of objects must be collected and also delivered to different locations depending on their own type. Three strategies are presented and the results obtained with the different settings are compared by considering the relation between performance

and a metric, called social entropy, introduced by the author. The social entropy measures the amount of the robotic agent diversity within the MRS. In the first setting, a homogeneous MRS is realized; each robot is capable of collecting and delivering both the types of objects. In the second setting, the MRS is heterogeneous and the robotic agents can only pick up either one type of object or the other. In the third case, the MRS is again heterogeneous, but this time the specialization is concerned with the task the robots can execute: by using a territorial division, some robots are in charge of collecting the objects around the environment and to drop them on the edge of an area which include the retrieving location. Inside the area, the remaining robotic agents collect the dropped objects and deliver them to the predefined areas. Experimental results show that a homogeneous MRS performs better than a heterogeneous one as long as the complexity of the task does not require a diversification of the capabilities of the MRS members.

In [24] a reactive, strongly coordinated, distributed MRS is presented. The heterogeneous robotic agents, realized in a behavior based architecture, have the common goal to clean an office like environment. One of the robotic agent is the vacuum and the other is the sweeper. The work studies the possibility to realize a grounded symbolic knowledge and communication. Three levels of representation are proposed, one on top of the other: the iconic representation, the indexical representation and the symbolic representation. Each of the levels is learned by the robots. The cooperation is realized with the superposition of four different modules, each of which is related to a representation level. The first module makes use of no representation and no communication, so that each robot acts for itself, by just executing correctly its own task. The input sensor data are the icons for the environmental objects (the sensors are the same so that the robots have the same icon). The second module lets the robots have knowledge of each other. An indexical representation is realized by the robotic agents by associating an icon to a robot. The sweeper becomes the index for the dust, so that the vacuum relates the sweeper position to the dust pile position, by watching where the sweeper stops. The third module introduces a direct communication among the agent so that the indexical representation can be shared. In the final module the symbolic representation is introduced, so that a more complete information can be exchanged by the robots. It lets the robotic agents be able to accomplish the task even if they are not in a direct view.

In [32] a reactive, distributed, strongly coordinated MRS is proposed to accomplish a foraging task. The system is composed by behavior-based homogeneous robots. The objective of the work is to obtain social behaviors via learning techniques, which allows for reducing the interferences between the robotic agents. The main problem is to find the appropriate reward, so that greedy robots act in a social way even if the effect of the acquired behavior are not of immediate benefit for the single robotic agent. Three different types of reinforcement are proposed to let the MRS present the social rule of yielding. The first is based on the perception that each robot has about the way its own task is being accomplished. The second kind of reinforcement, called observational rein-

forcement, encourages the robotic agent to repeat the behaviors observed for the other robots. The third type of reinforcement is obtained directly from the other robotic agents, based on their experiences. This kind of reinforcement does not require the robotic agent to model the internal state of the other robots, because, if the social rules are consistent, the punishment or the reward received would have been the same if the robotic agent were directly involved in the situation which generated them.

In [40] a reactive, distributed, strongly coordinated MRS, composed by heterogeneous behavior-based robotic agents, is tested with a foraging task. The focus of the work presented is the ALLIANCE architecture. All the robotic agents have sets of behaviors, each of which lets the robot itself to execute a corresponding task for which it is allowed. These behavior sets inhibit more basic and reactive behaviors, and are controlled by modules called motivational behaviors, which can cross inhibit each other. ALLIANCE's motivational behaviors are activated on the basis of a threshold mechanism; two parameters regulate their activation modalities, by characterizing the speed growth to the thresholds of the control functions. The first parameter is representative of the impatience of a robotic agent to take care of a task that is already assigned to an other agent within the system, so as to cope with the fault tolerance aspect. The second parameter is representative of the acquiescence of a robotic agent to give up the task which it is executing, so that it can accomplish another task for which it is more suitable and leaving the current one to a more adequate robot. Inputs to the control functions are the information that the robotic agents exchange, and the data coming from the each robot's sensors.

Other MRS which address the foraging task that are based on the swarm approach are described in [12, 31, 1]. All these works presented reactive MRS.

According to the characterization of social deliberation and reactivity proposed in the previous section, all the MRS presented in the above works are classified as reactive. In fact, in presence of an unforeseen event, each team member elaborates an individual solution, based on its own capabilities, to cope with the new situation. Therefore a single robot never modifies the global strategy to achieve the team's goal. In [4] the MRS is reactive because of its unawareness, the team cannot reorganize itself since each member ignores the existence of the others, hence only individual solutions can be addressed. The other works, on the contrary, present strongly coordinated MRS, but in each of them the social deliberative aspect is not realized, letting the robots face the unpredicted situations by providing individual specific tactics.

As highlighted by above discussion, foraging tasks have been usually addressed by reactive MRS, extending previous work on behavior-based architectures for a single robot. To the best of our knowledge deliberative approaches have not been proposed for this task. This may be explained by the fact that the major issue in foraging is to avoid spatial interferences among team members, which may be effectively addressed without the need to realized a social deliberative MRS.

4.2 Multi Target Observation

The multi target observation (also known as CMOMMT: Cooperative Multi Robot Observation of Multiple Moving Targets) is a very recent MRS testbed, first introduced by Parker in [41]. The task consists in maximizing the time during which each of the moving target is being observed by, at least, one of the robotic agents within the MRS. There are many connections with security, surveillance and recognition problems [41, 50, 23], where targets moving around in a bounded area must be observed. Problems like those connected with communication, sensor fusion, cooperation and coordination are involved in the multi target observation.

In [41], a reactive, distributed, strongly coordinated MRS, composed by homogeneous, behavior-based robotic agents (even if the system is allowed to cope with heterogeneous robots), is proposed to deal with the CMOMMT. The system is realized through the behavior based ALLIANCE architecture, combined with techniques based on potential fields and a target seeking system. In [50] the multi target observation task is achieved by using a social deliberative, distributed, strongly coordinated MRS. 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 for coordination between the robotic agents. Each behavior of each robot has a function which locally evaluates the robot's eligibility to accomplish a given task; the values are then exchanged between the robotic agents "peer behaviors". The robot whose behavior performed the highest value, inhibits the corresponding behaviors on the other system's members, thus claiming the task. This process, called cross-inhibition, can be executed only among peer behaviors. Internally to each agent the subsumption architecture selects the current behavior to be activated, by making use of a behavior arbitration possibly tuned so that priority is given to the task each robot is specialized for. Both the internal subsumption and the cross subsumption are sufficient to permit a heterogeneous MRS assigning the optimal task to each robotic agent within it.

The two above described works pinpoint a very interesting aspect: the difference between a reactive MRS and a reactive single robot architecture. They both present strongly coordinated MRS whose members are behavior based robotic agents, but while the ALLIANCE architecture implements a reactive system, with the BLE architecture a social deliberative MRS is realized. ALLIANCE ensures the robustness of the system by letting each member decide when to employ itself in accomplishing a task, without any need to reorganize the other members activity. A robot can take over the task assigned to another team member even if that member is still attempting to accomplish it. In the BLE architecture this is not possible, because 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 better fit the requirements, thus obtaining a new strategy.

Another work related to Multi-Robot target observation is presented in [23], where the authors describe a preliminary system that is focussed on tracking only one moving target (Paparazzi problem).

The multi target observation task is similar to the foraging one, with the addition of dynamic targets that must be continuously tracked. The more complex environment becomes an interesting testbed where deliberative systems could be effectively used.

4.3 Box pushing

The box pushing task has analogies with problems like, for example, stockage or truck loading and unloading [28, 42, 39]. Moreover, for a MRS, different issues, like task allocation, robustness or communication, are related to it; thus the box pushing task is a very common MRS testbed to verify all those techniques which cope with such issues.

In [28] a reactive, unaware MRS is realized. The system proposed is a classical example of swarm approach; the robotic agents are homogeneous and rely on a behavior based architecture. The authors implemented a set of behaviors inspired by the way some ant species organize themselves to accomplish a pushing task. The system's members do not have direct communication capabilities, and they select their actions driven by their own reactivity. Three different primitive behaviors are implemented: a behavior which provides a pose (position and orientation) modification attracted by the input stimuli, a behavior which acts exactly the opposite of the previous one and, finally, a behavior which activates a set of preprogrammed actions used to escape from a deadlock. The robots, by using behaviors relying on the previous ones, manifest a real cooperative effort in the box pushing task.

In [42] the box pushing task is used as a testbed for a reactive, distributed, strongly coordinated MRS. The system is composed by heterogeneous robotic agents realized with the behavior based architecture L-ALLIANCE, developed by the author. L-ALLIANCE extends ALLIANCE capabilities by introducing learning techniques to adjust the parameters which regulate the ALLIANCE working modes. These parameters, called impatience and acquiescence motivations, control the capability of a robot to handle situations when other robots (outside itself) fail in performing a given task, and the capability of a robot to handle situations in which it fails in accomplishing a task respectively. The way the entire MRS operates depends on the values assigned to the above mentioned parameters. Because in a long term mission the robotic agents can degrade their operational capabilities, an external robotic agent should be in charge of setting the motivation values depending on the actual state of the MRS. This is needed to avoid, by relying on a learning technique, to adapt the parameters to the current system configuration.

In [39] the task of box pushing is proposed to test a strongly coordinated, weakly centralized, social deliberative MRS, whose heterogeneous robotic agents

are realized in a hybrid architecture. The cooperative aspect of the MRS is here faced by dividing it in four different sequential phases: decomposition, allocation, local planning and execution. In the first phase a leader robot decomposes the global task into subtasks executable by a single robot or a group of them. Once the subtasks are individuated, they are scheduled and the robot's capabilities that are necessary to accomplish them are pinpointed. In the second phase the leader selects, by making use of a negotiation protocol, the available robots which fit the required capabilities. In the third phase the leader assigns them their own tasks, and each subordinated robot decomposes it into an action sequence which it can execute. If conflicts due to some kind of interference among the MRS's members arise, they must be solved in this phase. In the fourth phase the actions planned are executed by taking care of the possible unpredictable faults that could arise. If a leader presents some malfunction it can be replaced by another member of the team.

The work on box pushing is based on both reactive and social deliberative MRS. The reactive systems are implemented relying on various kind of techniques, in fact in [28] the system is unaware and the cooperation is obtained by letting each team member achieve its individual goal, without realizing any explicit coordination, thus the system reactivity is inherited by the single agent reactivity. In [42] the MRS is reactive, but it is also strongly coordinated, so as to reduce the situations in which the robotic agents could interfere with each other. Anyway each robot takes decision about how to cope with unpredicted events autonomously, by considering its knowledge about the environment and the team global state. In [39] each team member can notify an unforeseen situation to its leader, so as to reorganize the roles of the robots and the task assigned to each of them. This MRS is then a social deliberative system, able to eventually create a new strategy to deal with the environment modification or to provide robustness.

Applications in the box pushing domain have been realized by using both reactive and deliberative approaches. Indeed, as compared with the foraging and multi target observation problems, in box pushing one may take advantage of explicit cooperation in the accomplishment of the task to the extent in which box maneuvering requires coordination.

4.4 Exploration and Flocking

Under the label of exploration and flocking different tasks can be grouped: these tasks differ on the way they are realized, but have the common feature to require MRS's members to move around in the environment. Behaviors like flocking, formation maintenance or map building [6, 39, 22, 20, 45] are, thus, part of the same task class. In the exploration task there is the need to distribute the robots as sparsely as possible so that the information they can collect does not overlap, while in the flocking task the goal is that the robotic agents move around together, such as in a flock. The formation task can be seen as a way to realize either the exploration one or the flocking one, in fact, depending on which kind

of formation is chosen, the robotic agents can be as far as possible from each other, or as close as possible to each other.

In [6] a reactive, aware, not coordinated, homogeneous MRS, composed by not communicating behavior based robotic agents, is proposed to cope with the formation maintaining task, by using the potential function technique. The authors, inspired by the way the molecules snap into place as they form crystals, set each robotic agent with "attachment sites" to which the other system's members are drawn. The attachment site is not a physical device, but a position in the space surrounding a robotic agent. In this work the robotic agents are not forced to occupy a prefixed position within the formation, but they have different attachment sites disposition and, depending on which one is selected, they search for the first available position in the formation. The choice of one position respect to another depends on the robot's overall behavior generated by the superposition of the potential functions, called by the authors social potentials.

In [22] an exploration task in a office like environment is used to test a distributed, reactive MRS, designed to study the problem of grounding representations and communications of the robotic agents in real perceptions. The approach is strongly coordinated since, even though the robots are not coordinated for navigation, a protocol for exchanging representations of perceptual data is used. The author presents a technique to build a primitive language composed by a set of symbols obtained by classifying the robotic agent perception with a self-organizing classifier. The MRS's members cooperate in building a common knowledge about the environment by broadcasting their high-level representation of the perception data over the entire MRS. The process of building the set of symbols is divided into four subprocesses: perception, broadcast, singularity filtering and classification. In the first process each robot explores its portion of the environment producing perceptions, then, in the second phase, it broadcasts them to the other team members. The third phase consists in filtering the collected perceptions by retaining only those which are sufficiently singular with respect to the whole set. In the classification phase the robotic agents classify the set of perceptions so as to obtain the same resulting class.

The architecture presented in [39] is also used for a formation maintaining task. The resulting MRS is an example of strongly coordinated, weakly centralized, social deliberative MRS, composed by heterogeneous robotic agents realized in a hybrid architecture. In this work two robots travel in a office like environment in a convoy formation. Cooperation is obtained by dividing the task in four different sequential phases: decomposition, allocation, local planning and execution. A leader robot decomposes the global task into subtasks executable by a single robot or a group of them and directs them to the execution of these subtasks.

Other effective social deliberative MRS applications in the exploration domain have been realized by using strongly coordinated and centralized approaches [20, 45], or distributed ones [44].

The works concerned with exploration and flocking task realize both reactive and social deliberative systems. In [6] the MRS is not coordinated and thus it is

classified as reactive since its robotic agents execute their task by filling the gaps available in the formation, without communicating with each other. The work in [22] tests the possibility to realize a grounded language within the team, so it is not important to implement a MRS that presents a social deliberative behavior. Each team member can react to environmental changes by taking individual decisions without any necessity to reorganize the entire system strategy, while strong coordination is an essential property to realize the common grounded knowledge base. If the main task had been the exploration finalized to map building instead of a grounded language realization, a social deliberative system would have been more effective by letting the MRS better employ its resources. Finally in [39], as previously said, the system is social deliberative.

MRS for exploration or map building can be based on both reactive and social deliberative approaches. For this kind of tasks, with the increasing of the complexity of the environment social deliberative approaches could become more effective in ensuring a stronger level of cooperation.

4.5 Soccer

Robotic soccer is one of the tasks that has been considered in the last years as an interesting testbed for research in multi-agent and multi-robot cooperation [3, 13, 27]. Indeed the soccer games played by robotic teams provide an interesting setting in which cooperation among robots is fundamental for an effective accomplishment of the task in a highly dynamic and uncertain environment.

Depending on the different RoboCup leagues, different issues in multi-agent cooperation have been studied. The *Simulation league* is concerned with simulated agents, that are implemented by separate programs and can communicate among each other. Because of the lack of global information about the environment for the single agent, this setting provides an interesting environment for experiments in coordinating a multi-agent system [21, 35]. The other RoboCup leagues are instead composed by actual robotic teams and the dynamic and uncertain environment in which they operate makes coordination of the multi-robot system a more challenging problem.

In the *Small-Size league* robots are small and very fast, they are usually controlled by a remote host that can take advantage of a global vision system (a camera positioned above the field). The availability of global information about the environment and the use of a remote host for robots' control allows for the implementation of centralized strategies [49]. Most of these systems are strongly centralized and deliberative, since the remote host can be considered as an agent acting as the leader of the team. However, reactive approaches can also be found in the behavior based hierarchical control architecture proposed in [43].

On the contrary, in the *Middle-Size league*, all the robot's sensors must be on board; therefore robots are more autonomous and have to deal with high uncertainty in reconstructing global information about the environment. The approach in [16] represents an example of reactive, weakly coordinated system, in which two robots are able to exchange the ball by using direct communication of their roles and a field-vector based collision avoidance that takes into

account in a proper way the role of the robots. Other behavior based reactive MRS are presented in [7, 10]. On the other hand, successful deliberative centralized approaches in the Middle-Size league are used when it is possible to reliably reconstruct global information about the environment. For example the CS Freiburg team [19] makes use of very precise measurements given by laser scanners to reconstruct a global representation of the field and a social deliberative weakly centralized approach is obtained by dynamic selection of the robot that acts as the leader. However, possible communication failures, that are very common in this environment, as well as the general difficulty of reconstructing a global reliable view of the environment, require full autonomy on each robot and distributed approaches are to be preferred. Distributed heterogeneous robots cooperating by using explicit communication is described in [51]. The Azzurra Robot Team also presents a distributed heterogeneous robotic soccer team [11] based on a simple and flexible coordination protocol. From a technical perspective the proposed approach makes use of a *formation/role* mechanism and of dynamic assignment of roles 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. In fact 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. A distributed architecture has the advantage that, in case of communication failures, the system degrades its performance, but it still keeps on the execution of the task. Finally, in the *Sony Legged league* three Sony AIBOs robots are used to form a team. Therefore the robots are the same in hardware and heterogeneity can be only given by different software control. At this time, mostly because of the lack of network devices on the robots, there are very simple forms of cooperation: in SPQR team [36] the robots are unaware of each other, but cooperation is still obtained by a static assignment of the roles (i.e. GoalKeeper, Defender and Forward). Other teams (for example [29, 48]) have developed a different software only for the goalkeeper, while the other two robots are the same. For an effective game the robots should not obstruct each other when approaching the ball. Using different roles can help to accomplish this task; when the robots have exactly the same behavior it is not possible to guarantee that robots do not interfere.

The RoboCup teams have chosen reactive or deliberative approaches for their MRS, mainly depending on the capabilities of the robots involved in the competitions. In fact, the global vision system allowed in the *Small-Size league* has led to centralized and deliberative approaches; in the *Middle-Size league*, in which robots are fully autonomous, the deliberative approach has been mostly used, but also reactive MRS are present; finally, in the *Sony Legged league* we have only reactive MRS, due to the difficulties of communication among robots.

In general, the need of dynamically adapting different strategies during the game, has been considered an important aspect of social deliberation. For example, in order to realize in a robotic team the capabilities of passing the ball, a strong cooperation is required and social deliberative approaches based on

centralized techniques [37] as well as on learning [46] have been proved to be successful.

5 Conclusions

In this paper we have presented a classification of MRS with the goal of highlighting reactivity and social deliberation. It is worth emphasizing that, although a MRS can be regarded as a kind of multi-agent system, the need to cope with the acquisition of knowledge from the environment, makes it more challenging to build actual experimentation settings for MRS. 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.

We have first introduced a set of dimensions for the classification of MRS and provided a taxonomy of the type of cooperations that can be found in MRS. In addition, we have discussed two dimensions that can be considered orthogonal to the taxonomy in the sense that they do not interact with the features used to build the taxonomy. Then, we have focussed our attention on the issue of reactivity and social deliberation, by distinguishing the single robot from the MRS setting, and by proposing a characterization of MRS in terms of reactivity and social deliberation. In addition, we have seen that reactive approaches are viable in the settings where no coordination takes place, for example when the robots are unaware of the presence of other robots, as well as in distributed, strongly coordinated system organizations. On the other hand, social deliberation is based on coordination and it is implicit in the centralized approaches, while being suited also to distributed ones.

Finally, we have presented some of the MRS described in the literature by arranging them according to the task accomplished. As it turns out, reactive and social deliberative approaches are pursued in order to solve any of the problems examined, possibly with the exception of the foraging task. The outcomes of this analysis show that reactive approaches tend to be used when there are several robots and the task can be independently achieved by a single robot. However, when the task requires some degree of coordination social deliberation can be effectively applied. The models that are used for coordination in MRS are often not complex: the critical aspect is rather on their actual application, given all the difficulties arising in coping with dynamic environments.

We believe that, given the growing interest in practical application of MRS, we will have a significant increase of the work in this field. A major issue to be addressed by future work is the definition of suitable evaluation methodologies, in order to assess the adequacy and effectiveness of various forms of cooperation in MRS (see for example [5]). A major role in this respect can be played by the competitions [26, 25], which provide a common set up that is a basic requirement for comparing different approaches.

6 Acknowledgments

This research has been carried out within the project “Robot Calciatori” funded by Italian Consiglio Nazionale delle Ricerche and the projects RAMSETE and MISTRAL funded by Italian Ministero dell’Università e della Ricerca Scientifica e Tecnologica. The authors also wish to thank the reviewers for their detailed comments which helped to improve the quality of this article.

References

1. R. C. Arkin. Cooperation without communication: Multiagent schema-based robot navigation. *Journal of Robotic Systems*, 9(3), 1992.
2. R. C. Arkin and G. A. Bekey. *Robot Colonies*. Kluwer Academic Publishers, 1997.
3. M. Asada. The RoboCup physical agent challenge: Goals and protocols for Phase-I. In H. Kitano, editor, *RoboCup-97: Robot Soccer World Cup I*, 1998.
4. T. Balch. The impact of diversity on performance in multi-robot foraging. In *Proc. of Agents '99*, 1999.
5. T. Balch. Hierarchic social entropy: an information theoretic measure of robot team diversity. *Autonomous Robots*, 8(3), 2000.
6. T. Balch and M. Hybinette. Social potentials for scalable multirobot formations. In *Proc. of IEEE International Conference on Robotics and Automation (ICRA-2000)*, 2000.
7. 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.
8. Rodney A. Brooks. A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, RA-2(1), 1986.
9. Y. Uny Cao, A. Fukunaga, and A. Kahng. Cooperative mobile robotics: Antecedents and directions. *Autonomous Robots*, 4:1–23, 1997.
10. 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.
11. C. Castelpietra, L. Iocchi, D. Nardi, M. Piaggio, A. Scalzo, and A. Sgorbissa. Coordination among heterogenous robotic soccer players. In *Proc. of International Conference on Intelligent Robots and Systems (IROS'2000)*, 2000.
12. F. Chantemargue and B. Hirsbrunner. A collective robotics application based on emergence and self-organization. In *Proc. of Fifth International Conference for Young Computer Scientists (ICYCS'99)*, 1999.
13. S. Coradeschi, L. Karlsson, P. Stone, T. Balch, G. Kraetzschmar, and M. Asada. Overview of RoboCup-99. *A.I. Magazine*, 1999.
14. D. Dudek, M. Jenkin, E. Miliotis, and D. Wilkes. A taxonomy for multi-agent robotics. *Autonomous Robots*, 3(4):375–397, 1996.
15. J. Ferber. *Multi-Agent Systems*. Addison-Wesley, 1999.
16. M. Ferrarezzo, C. Ferrari, E. Pagello, R. Polesel, R. Rosati, A. Speranzon, and W. Zanette. Collaborative emergent actions between real soccer robots. In *RoboCup-2000: Robot Soccer World Cup IV*, 2000.

17. Erann Gat. Integrating planning and reacting in a heterogeneous asynchronous architecture for controlling real-world mobile robots. In *Proceedings of the Tenth National Conference on Artificial Intelligence (AAAI'92)*, 1992.
18. M. Georgeff and A. Lansky. Reactive reasoning and planning. In *Proceedings of the Sixth National Conference on Artificial Intelligence (AAAI'87)*, 1987.
19. J.-S. Gutmann, T. Weigel, and B. Nebel. Fast, accurate, and robust self-localization in the robocup environment. In *RoboCup-99: Robot Soccer World Cup III*, 1999.
20. D. Guzzoni, A. Cheyer, L. Julia, and K. Konolige. Many robots make short work. *AI Magazine*, 1997.
21. M. Hannebauer, J. Wendler, P. Gugenberger, and H. Burkhard. Emergent cooperation in a virtual soccer environment. In *Proc. of DARS-98*, 1998.
22. L. Hugues. Grounded representations for a robots team. In *Proc. of the 2000 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2000.
23. M. Jenkin and G. Dudek. The paparazzi problem. In *Proc. of International Conference on Intelligent Robots and Systems (IROS'2000)*, 2000.
24. David Jung and Alexander Zelinsky. Grounded symbolic communication between heterogeneous cooperating robots. *Autonomous Robots*, 8(3):269–292, 2000.
25. H. Kitano. Robocup rescue: A grand challenge for multi-agent systems. In *Proc. of ICMAS 2000*, 2000.
26. 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*, volume 1395, pages 1–19. Springer Verlag, 1998.
27. H. Kitano, E. Pagello, and M. Veloso, editors. *RoboCup-99: Robot Soccer World Cup III*. Springer-Verlag, 1999.
28. C. Ronald Kube and E. Bonabeau. Cooperative transport by ants and robots. *Robotics and Autonomous Systems*, 30(1):85–101, 2000.
29. M. Lawther and J. Dalglish. Unsw united. In *RoboCup-99: Robot Soccer World Cup III*, 1999.
30. J. Malek. On augmenting reactivity with deliberation in a controlled manner. In *Proc. of ECAI2000 Workshop on Balancing Reactivity and Social Deliberation in Multi-Agent Systems*, 2000.
31. M. Mataric. *Interaction and Intelligent Behavior*. PhD thesis, MIT, 1994.
32. Maja J. Mataric. Learning social behavior. *Robotics and Autonomous Systems*, 20:191–204, 1997.
33. Maja J. Mataric. Behavior-based robotics. In Robert A. Wilson and Frank C. Keil, editors, *MIT Encyclopedia of Cognitive Sciences*, pages 74–77. MIT Press, 1999.
34. V. K. Mavromichalis and G. Vouros. ICAGENT: Balancing between reactivity and deliberation. In *Proc. of ECAI2000 Workshop on Balancing Reactivity and Social Deliberation in Multi-Agent Systems*, 2000.
35. F. Montesello, A. D'Angelo, C. Ferrari, and E. Pagello. Implicit coordination in a multi-agent system using a behavior-based approach. In *Proc. of DARS-98*, 1998.
36. D. Nardi, C. Castelpietra, A. Guidotti, M. Salerno, and C. Sanitati. S.P.Q.R. In *RoboCup-2000: Robot Soccer World Cup IV*. Springer-Verlag, 2000.
37. B. Nebel, J. S. Gutmann, and W. Hatzack. CS Freiburg '99. In *RoboCup-99: Robot Soccer World Cup III*, 1999.
38. N. J. Nilsson. Shakey the robot. Technical Report 323, SRI Artificial Intelligence Center, 1984.
39. Fabrice R. Noreils. Toward a robot architecture integrating cooperation between mobile robots: Application to indoor environment. *International Journal of Robotics Research*, 12(1):79–98, 1993.

40. Lynne E. Parker. ALLIANCE: An architecture for fault tolerant multirobot cooperation. *IEEE Transactions on Robotics and Automation*, 14(2):220–240, 1998.
41. Lynne E. Parker. Cooperative robotics for multi-target observation. *Intelligent Automation and Soft Computing*, 5(19), 1999.
42. Lynne E. Parker. Lifelong adaption in heterogeneous multi-robot teams: Response to continual variation in individual robot performance. *Autonomous Robots*, 8(3):239–267, 2000.
43. G. Wagner S. Behnke, R. Rojas. A hierarchy of reactive behaviors handles complexity. In *Proc. of ECAI2000 Workshop on Balancing Reactivity and Social Deliberation in Multi-Agent Systems*, 2000.
44. A. Saffiotti, N.B. Zumel, and E.H. Ruspini. Multi-robot team coordination using desirabilities. In *Proc of the 6th Intl Conf on Intelligent Autonomous Systems (IAS)*, pages 107–114, Venice, Italy, 2000.
45. 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 International Conference on Intelligent Robots and Systems (IROS'2000)*, 2000.
46. P. Stone. *Layered Learning in Multiagent Systems*. MIT Press, 2000.
47. C. F. Touzet. Robot awareness in cooperative mobile robot learning. *Autonomous Robots*, 2, 2000.
48. M. Veloso, S. Lenser, E. Winner, and J. Bruce. Cm-trio-99. In *RoboCup-99: Robot Soccer World Cup III*, 1999.
49. M. Veloso and P. Stone. Individual and collaborative behaviors in a team of homogeneous robotic soccer agents. In *Proceedings of the Third International Conference on Multi-Agent Systems*, 1998.
50. B. B. Werger and M. J. Mataric. Broadcast of local eligibility for multi-target observation. In *Proc. of DARS*, 2000.
51. K. Yokota, K. Ozaki, N. Watanabe, A. Matsumoto, D. Koyama, T. Ishikawa, K. Kawabata, H. Kaetsu, and H. Asama. Uttori united: Cooperative team play based on communication. In *RoboCup-98: Robot Soccer World Cup II*, 1998.