

Distributed Intelligence: Overview of the Field and its Application in Multi-Robot Systems

Lynne E. Parker

Abstract—This article overviews the concepts of distributed intelligence, outlining the motivations for studying this field of research. First, common systems of distributed intelligence are classified based upon the types of interactions exhibited, since the type of interaction has relevance to the solution paradigm to be used. We outline three common paradigms for distributed intelligence — the bioinspired paradigm, the organizational and social paradigm, and the knowledge-based, ontological paradigm — and give examples of how these paradigms can be used in multi-robot systems. We then look at a common problem in multi-robot systems — that of *task allocation* — and show how the solution approach to this problem is very different depending upon the paradigm chosen for abstracting the problem. Our conclusion is that the paradigms are not interchangeable, but rather the selection of the appropriate paradigm is dependent upon the specific constraints and requirements of the application of interest. Further work is needed to provide guidance to the system designer on selecting the proper abstraction, or paradigm, for a given problem.

Index Terms—Distributed intelligence, multi-robot systems, multi-agent systems, task allocation.

I. INTRODUCTION TO DISTRIBUTED INTELLIGENCE

DISTRIBUTED Intelligence refers to systems of *entities* working together to reason, plan, solve problems, think abstractly, comprehend ideas and language, and learn. Here, we define an *entity* as any type of intelligent process or system, including agents, humans, robots, smart sensors, and so forth. In these systems, different entities commonly specialize in certain aspects of the task at hand. As humans, we are all familiar with distributed intelligence in teams of human entities. For example, corporate management teams consist of leaders with particular specialties such as Chief Executive Officer (CEO), Chief Operating Officer (COO), Chief Financial Officer (CFO), Chief Information Officer (CIO), and so forth. Oncology patient care teams consist of doctors that specialize in various areas, such as surgical oncology, medical oncology, plastic and reconstructive surgery, pathology, etc. Distributed intelligence is also exhibited in military applications, such as special forces A-Teams, where team members specialize in weapons, engineering, medicine, communications, and so forth. Another military example includes personnel on an aircraft carrier flight deck, who are segmented into the catapult

crew, the landing signal officers, ordnancemen, plane handlers, etc. As humans have clearly learned, these teams can very efficiently solve complex tasks by making use of specialists who work together productively.

The objective of distributed intelligence in computer science (and related fields) is to generate systems of software agents, robots, sensors, computer systems, and even people and animals (such as search and rescue dogs) that can work together with the same level of efficiency and expertise as human teams. Clearly, such systems could address many important challenges, including not only urban search and rescue, but also military network-centric operations, gaming technologies and simulation, computer security, transportation and logistics, and many others.

As a research topic, the study of distributed intelligence has gained much popularity in recent years. Figure 1 shows data from the Web of Science resulting from a keyword search on the terms “distributed intelligence”, “distributed AI”, “distributed artificial intelligence”, “multiagent”, “multi-agent”, “distributed robot”, “multirobot”, and “multi-robot”. Each year’s results show the number of publications that appeared containing these keywords in that year. The search begins in year 1980 — the most recent year with no publications containing any of these keywords — up through 2006. While this is an admittedly incomplete survey of this area of research, the data clearly shows the significantly increasing interest in this research area, as investigators and application developers are recognizing the potential power of distributed intelligence.

What is the potential promise of distributed intelligence? Certainly, some applications can be better solved using a distributed solution approach — especially those tasks that are inherently distributed in space, time, or functionality. Further, if a system is solving various subproblems in parallel, then it offers the potential of reducing the overall task completion time. Any system consisting of multiple, sometimes redundant entities, offers possibilities of increasing the robustness and reliability of the solution, due to the ability for one entity to take over from another failing entity. Finally, for many applications, creating a monolithic entity that can address all aspects of a problem can be very expensive and complex; instead, creating multiple, more specialized entities that can share the workload offers the possibility of reducing the complexity of the individual entities.

Of course, these advantages of distributed intelligence are, to some extent, offset by some disadvantages. For example,

Lynne Parker is with the Distributed Intelligence Laboratory in the Department of Electrical Engineering and Computer Science at University of Tennessee, Knoxville, Tennessee, USA.

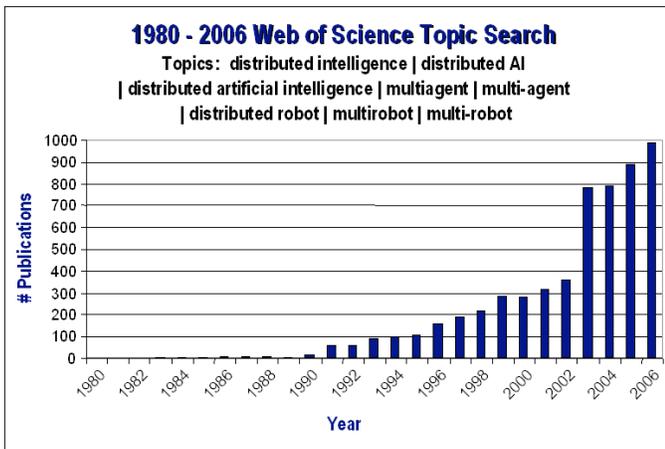


Fig. 1. Web of Science data showing the number of publications appearing per year on topics related to distributed intelligence.

even though the individual entity cost and complexity may be less, determining how to manage the complete system may be more difficult and complex, because of the lack of centralized control or of a centralized repository of global information. Further, distributed intelligent systems may require more communication to coordinate all the entities in the system. Increasing the number of entities can lead to increased interference between entities, as they must act without complete knowledge of the other entities' intents. Finally, systems of multiple entities will typically experience increased uncertainty about the state of the system as a whole.

Overall, however, as new solution approaches are being developed and validated, the research and user community is finding that it is often possible to realize the advantages of distributed intelligence systems while countering many of the possible disadvantages. The challenge is determining how best to properly design the system so as to achieve global coherence through the local interactions of individual entities.

II. THE DOMAIN SPACE OF DISTRIBUTED INTELLIGENCE

As researchers are discovering, there are many possible solution strategies, or paradigms, for achieving distributed intelligence. Not all of these paradigms are appropriate for all types of distributed intelligence. Thus, it is important to understand the various types of distributed intelligence that can occur in different application settings.

There are many ways to classify distributed intelligent systems and multi-robot systems, such as those proposed in [28], [17], [24], [102]. While these existing classifications provide important insight to various characteristics of multi-robot systems, they do not focus specifically on the types of *interactions* exhibited by these systems. For our purposes, we believe it is helpful to view the domain space in distributed intelligence in terms of the types of interactions that can take place between the entities in the system. To help make the distinctions between the various types of interactions clear, we find it useful to define a minimal set of variables (or characteristics) whose values, together, can help us easily categorize the different types of interactions exhibited in multi-robot systems.

Toward this end, as illustrated in Figure 2, we view the types of interactions along three different axes — the *types of goals*, whether entities have *awareness of others* on the team, and whether an entity's actions *advance the goals of others* on the team. In terms of types of goals, we classify systems into two types — those in which each entity has *individual* goals, and those in which the entities have *shared* goals. For the *awareness of others* axis, we divide the systems into those that are *aware* and those that are *not aware*. By *aware* in this context, we refer to whether entities reason about the actions and intentions of their teammates. Robots that are *not aware* may sense the presence of local entities and move so as to maintain a certain distance, for example, but otherwise perform no other reasoning to understand the intent or future plans of the teammates. Often, these “un-aware” systems operate based on the principle of *stigmergy*, in which communication between entities is not direct, but rather through changes made in the environment.

Finally, we segment systems into those in which an entity's actions do advance the goals of others on the team (*yes*) and those that do not (*no*). An example of an entity advancing the goals of others with its actions is a floor cleaning robot, as a member of a floor cleaning robot team. Each robot's actions of cleaning a bit of the floor are helpful to the other teammates, who do not have to repeat the floor cleaning in that particular spot.

Obviously, these segmentations of the domain space are approximate, yet we believe they are helpful in quickly understanding and categorizing the primary types of interactions that can occur in typical applications. Different areas of this subspace represent common types of interactions seen in systems of distributed intelligence. These common forms of interaction are:

- Collective
- Cooperative
- Collaborative
- Coordinative

In the following paragraphs we describe these types of interactions in more detail, discussing their relevance and application to multi-robot systems.

Perhaps the simplest type of interaction is the *collective* interaction, in which entities are not aware of other entities on the team, yet they do share goals, and their actions are beneficial to their teammates. An example of this type of interaction in multi-robot systems is the swarm robotics work of many researchers (e.g., [73], [71], [59]). This work focuses on creating systems of robots that can perform biologically-relevant tasks, such as foraging, swarming, flocking, herding, formation-keeping, and so forth. Robots in these systems typically perform relatively simple local control laws which, when combined with larger numbers of robots, result in the global goal being achieved, often as an emergent property of the local interactions.

The second type of interaction is the *cooperative* interaction, in which entities are aware of other entities, they share goals, and their actions are beneficial to their teammates. In multi-robot systems, an example of this type of interaction is multiple robots working together and reasoning about each

other's capabilities in order to accomplish a joint task, such as pushing a box (e.g., [39]), cleaning up a worksite (e.g., [82]), performing search and rescue (e.g., [77]), or extra-planetary exploration (e.g., [103]). In these systems, robots may at times be working on different parts of the higher level goal, and thus may at times have to ensure that they share the workspace without interfering with each other. However, the majority of the work of the robots is focused on working together to achieve a common goal.

A third type of interaction in systems of distributed intelligence occurs when robots have individual goals, they are aware of their teammates, and their actions do help advance the goals of others. This part of the domain space is typically called *collaborative*, and is characterized by entities helping each other to achieve their individual, yet compatible, goals. While closely associated to the cooperative domain space, we make a distinction here to focus on the ability of entities to work together to help others better achieve their individual goals. In human research teams, we are familiar with the concept of collaboration, in which each person brings unique expertise that helps the team as a whole achieve a broader objective. Each team member has his/her own goal of performing his/her own aspect of the research, but by working together with others with complementary expertise, each can help the other members better achieve their individual goals. Of course, most of these collaborations are also cooperative, and it is possible to turn a collaborative team into a cooperative team by simply viewing the team goals from a higher perspective. A multi-robot example of a collaborative team is a group of robots that each must reach specified goal positions that are unique to each member. If robots are unable to reach their goal positions independently, due to sensor limitations, they could work together with other robots by sharing sensory capabilities to help all team members reach their individual goal locations. This type of collaboration is sometimes called *coalition formation*, and has been illustrated in multi-robot systems in [85], [119].

Finally, the fourth type of interaction relevant to distributed intelligence is what we call *coordinative*. In these systems, entities are aware of each other, but they do not share a common goal, and their actions are not helpful to other team members. In multi-robot systems, these situations often occur when robots share a common workspace. The robots must work to coordinate their actions to minimize the amount of interference between themselves and other robots. Multi-robot path planning techniques (e.g., [56], [43], [1], [25], [86], [88], [19], [107], [53], [93], [63]) or traffic control techniques (e.g., [40], [54], [66], [3], [125], [121]) are commonly used in these domains.

As a side note, we could have extended the third axis of our domain space to categorize systems based on whether they (1) *positively* affect the goals of other entities, (2) have *no effect* on the goals of other entities, or (3) *negatively* effect the goals of other entities. Then, we could create a new type of interaction in which entities have individual goals, they are aware of each other, but their actions have a negative effect on others' goals. This defines the *adversarial* domain, in which entities actively work against each other. In multi-robot

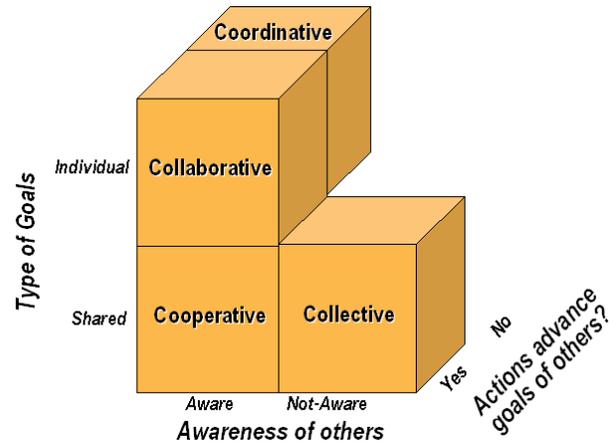


Fig. 2. Categorization of types of interactions in systems of distributed intelligence.

systems, this topic is studied extensively in the multi-robot soccer application domain (e.g., [55], [15], [118], [101]). This form of interaction also has clear relevance for many security and military applications.

Understanding the types of interactions that we want to achieve in a distributed intelligence system can provide insights into the appropriate solution strategy. The following section outlines some common paradigms for distributed intelligence that can achieve these varying types of interactions in multi-robot systems.

III. PARADIGMS FOR DISTRIBUTED INTELLIGENCE

Just as there are many types of interactions in systems of distributed intelligence, there are also many paradigms for achieving distributed intelligence. Each paradigm abstracts the problem space in a different way, enabling the system designer to view the system from a perspective that sheds light on proper solution strategies. Often, these paradigms take inspiration from societies of insects, or societies of humans. Not all paradigms are appropriate for all types of interaction dynamics. In this section, we outline some of the more common paradigms for distributed intelligence, especially focusing on their relevance to multi-robot systems. Note that a fundamental challenge in all of these paradigms is determining how best to achieve global coherence from the interaction of entities at the local level. By abstracting the problem in different ways, alternative solution strategies become apparent that can help address this challenge.

Three commonly used paradigms for building systems of distributed intelligence include:

- Bioinspired, emergent swarms paradigm,
- Organizational and social paradigms, and
- Knowledge-based, ontological, and semantic paradigms.

The following subsections outline these paradigms in more detail.

A. Bioinspired, emergent swarms paradigm

The behavioristic approach to autonomous robot control that gained popularity beginning in the 1980's [13] has its roots

in the observations of animal behavior. Animals, particularly the lower animals, are existence proofs that interesting results can be achieved without the need for a complex, human-level architecture. Many animals appear to be “hard-wired” for certain behaviors, producing very stereotypical reactions to particular stimuli. For instance, a robin begins defending its territory when it sees the red breast of another robin, or even a bunch of red feathers [26]. A pregnant Three-spined Stickleback fish approaches a male Stickleback with a red belly, or even a crude model of a Stickleback, as long as it is painted red underneath [114]. A male grayling butterfly flies up to mate rather large, dark, close, dancing objects, which could include not only female graylings, but also birds, falling leaves, and shadows [115].

Applying animal observations to the realm of autonomous robotics, interesting and seemingly intelligent activities can be obtained by layering behaviors that react to stimuli from the world according to the robot’s current internal state [13]. Rather than decomposing the robot control system based on information processing functions, the behavioristic approach decomposes the control into task achieving behaviors, such as obstacle avoidance, exploration, and map building. The result was a series of autonomous robots that can survive in a dynamic world, avoiding obstacles, exploring the environment, following walls, building maps, climbing over uneven terrain, and so forth [14].

But this same approach — the observation of animal behavior — that has been used for inspiration in the development of individual robots is just as easily used to gain insight into the creation of groups of robots that cooperate toward attaining some goal [80]. By learning how various species of animals function as groups, ideas can be obtained for building a cooperating team of autonomous robots.

1) *Two Types of Animal Societies*: Since there are so many varieties of social behavior in the animal kingdom, a classification of animal societies is useful. One such classification, proposed by Tinbergen [114], is of particular interest for current robotics research in multi-robot systems, as it parallels two possible approaches to achieving multi-robot systems. According to Tinbergen, animal societies can be grouped into two broad categories: those that *differentiate*, and those that *integrate*.

Societies that *differentiate* are realized in a dramatic way in the social insect colonies [124]. These colonies arise due to an innate differentiation of blood relatives that creates a strict division of work and a system of social interactions among the members. Members are formed within the group according to the needs of the society. In this case, the individual exists for the good of the society, and is totally dependent upon the society for its existence. As a group, accomplishments are made that are impossible to achieve except as a whole.

On the other hand, societies that *integrate* depend upon the attraction of individual, independent animals to each other. Such groups do not consist of blood relatives that “stay together”, but instead consist of individuals of the same species that “come together” by integrating ways of behavior [90]. These individuals are driven by a selfish motivation which leads them to seek group life because it is in their own

best interests. Interesting examples of this type of society are wolves, or the breeding colonies of many species of birds, in which hundreds or even thousands of birds congregate to find nesting partners. Such birds do not come together due to any blood relationship; instead, the individuals forming this type of society thrive on the support provided by the group. Rather than the individual existing for the good of the society, we find that the society exists for the good of the individual.

2) *Parallels in Multi-Robot Systems*: In analyzing research in multi-robot systems, a parallel can be drawn with the classifications of animal societies discussed above. A large body of work in robotics, including much of the earliest work (e.g., [21], [112], [98], [23], [70], [10], [59], [99], [33], [122]), involves the study of emergent interactions in colonies of robots — an approach comparable to *differentiating* animal societies. This research emphasizes the use of large numbers of identical robots that individually have very little capability, but when combined with others can generate seemingly intelligent group behavior. This intelligent group behavior is achieved as a side-effect of the individual robot behaviors. This type of interaction is typically called *collective*.

In this paradigm, the need for communication between entities is greatly reduced by assuming the ability of the entities to sense relevant information in their local environments (i.e., *stigmergy*). The application requirements in these problems allow for simple action protocols, or control rules, that are identical on each entity, and that lead to the desired group behavior. An example local control rule under this paradigm that can cause all the agents/robots to aggregate (as in a swarm) is [71]:

```
Aggregate:
  If agent is outside aggregation
    distance
  then turn toward aggregation
    centroid and go.
Else
  stop.
```

This is a powerful paradigm for those applications that require the same task to be performed across a distributed workspace, where the task does not require complex interactions of entities and all entities are interchangeable. Research challenges include developing tools that can predict the global behavior given a set of local control rules, as well as the inverse problem, in which we want to derive the local control rules, given a desired global behavior. This paradigm is relevant for many spatially distributed applications, including flocking, schooling, and formations (e.g., [92], [70], [5], [81], [105], [73], [75], [76], [8]); foraging, coverage, and search (e.g., [87], [31], [120], [34], [4], [104], [94], [106], [16], [69]); target tracking and observation (e.g., [7], [123], [65], [62], [57], [51], [111]); sorting and clumping [8], and so forth. While earlier approaches to these applications were based on human-generated local control rules that were demonstrated to work in practice, more recent work is based on control theoretic principles, with a focus on proving stability and convergence properties in multi-robot team behaviors. Examples

of this work include [47], [9], [116], [29], [68], [35], [36], [108], [2].

Other types of interactions, however, require more complex solution paradigms. Thus, a second approach parallels the *integrative* societies in the animal kingdom. This research aims to achieve higher-level, “intentional” cooperation or collaboration amongst robots. Rather than beginning with robots having very low-level behaviors, individual robots that have a higher degree of “intelligence” and capabilities are combined to achieve purposeful cooperation or collaboration. The goal is to use robots that can accomplish meaningful tasks individually, and yet can be combined with other robots with additional skills to complement one another in solving tasks that no single robot can perform alone. To be purely analogous to the integrative animal societies, robots in this type of interaction would have individual, selfish, motivations which lead them to seek cooperation or collaboration [72]. Such interactions would be sought because it is in the best interests of each robot to do so to achieve its mission. Of course, the possession of a selfish motivation to cooperate or collaborate does not necessarily imply consciousness on the part of the robot. It is doubtful that we would attribute consciousness to all the integrative societies in the animal kingdom; thus, some mechanism must exist for achieving this type of productive interaction without the need for higher-level cognition. A large body of multi-robot systems research falls into this category (e.g., [82], [123], [78], [12], [15], [67], [101], [103], [119], [1]). The next two paradigms discussed in this article — namely, the organizational/social paradigm and the knowledge-based/ontological/semantic paradigm — are additional techniques that have been used to create similar higher-level, intentional cooperation and/or collaboration in multi-robot teams.

The type of approach one should use for the multi-robot solution is dependent upon the applications envisioned for the robot team. The differentiating interaction approach is useful for collective tasks requiring numerous repetitions of the same activity over a relatively large area (relative to the robot size), such as waxing a floor, cleaning barnacles off of ships, collecting rock samples on a distant planet, and so forth. Such applications would require the availability of an appropriate number of robots to effectively cover the work area while continuing to maintain the critical distance separation.

On the other hand, the intentional interaction approach would be required for cooperative or collaborative interactions in applications requiring several distinct tasks to be performed, perhaps in synchrony with one another. Throwing more robots at such problems would be useless, since the individual tasks to be performed cannot be broken into smaller, independent subtasks. Examples of this type of interaction include container management in ports [1], some aspects of extra-planetary exploration [103], some types of search and rescue [49], some aspects of mineral mining [95], transportation [113], industrial and household maintenance [83], construction [96], hazardous waste cleanup [82], security [27], [44], agriculture [89], and warehouse management [45].

Of course, there is overlap in the relevance of these approaches to various applications, and in some instances the dif-

ferences are a matter of degree. For instance, if large numbers of robots are too expensive or are not available to be applied to, say, planetary exploration, then more purposive interaction (i.e., cooperation or collaboration) is required to achieve the goal of the mission. Combinations of the approaches are also possible by using intentionally interacting robots to guide the activities of smaller groups of swarm robots in a coordinated way (e.g., [84]).

B. Organizational and social paradigms

Organizational and social paradigms are typically based on organizational theory derived from human systems. Knowledge from the fields of sociology, economics, and psychology, and related areas, have proven valuable for understanding how to create systems of intelligent artifacts that can work together to solve complex problems. In these approaches, agent/robot interactions are designed by modeling individual and group dynamics as part of an organization. These approaches reduce the communications requirements among entities by making use of models drawn from these fields. This type of approach is commonly used for *cooperative* and *collaborative* types of distributed intelligence. Three examples of organizational theory applied to multi-robot systems are the use of *roles*, the use of *market economies*, and the use of *teamwork models*.

Roles are often used to divide the application into manageable portions of the work that can each be assigned to a different robot in the team. An easy division of work is achieved by assigning roles according to the skills and capabilities of the individual team members. For instance, in multi-robot soccer [100], [67], [118], positions played by the different robots are often defined as roles, such as goal keeper, left defender, right defender, left forward, right forward, and so forth. The robot best suited, and perhaps in closest proximity, to the available roles/positions then selects to perform that role.

Market economies are used in multi-robot systems as a paradigm for *task allocation*, which we discuss further in the next section. In brief, task allocation is the problem of mapping tasks to robots, such that the most suitable robot is selected to perform the most appropriate task, leading to all tasks being satisfactorily completed. Market-based approaches to task allocation (e.g., [22], [126], [39], [12]) make use of the theory of market economies to determine how best to allow robots to negotiate on responsibilities in the mission. Market economy approaches define methods for how to manage bids, how to handle multiple bids in parallel, how to consider multiple tasks at once, and so forth. More discussion on this topic is given in the next section.

Teamwork models allow agents/robots to explicitly reason about coordination and communication. In dynamic environments, the ability to reason about the interactions of agents/robots can enable the team members to reorganize themselves as needed to address new situations that arise. Many teamwork models, such as Tambe’s STEAM model [109], are based on the concept of *joint intentions* [20], which models the joint mental state of the team. The idea is that a team is jointly committed to a task if its individual team

members are committed to the task and believe that they are executing the task. Protocols for establishing team member commitments are defined as part of this general model. A related concept is that of *shared plans* [41], which represent agents' intentions to perform particular tasks.

C. Knowledge-based, ontological, and/or semantic paradigm

A third paradigm commonly used for developing systems of distributed intelligence is the *knowledge-based, ontological, and/or semantic* paradigm. The focus in these approaches is on knowledge sharing between heterogeneous robots/agents, with the objective of easily allowing these entities to share and understand knowledge from disparate sources. Often, knowledge is defined as an *ontology*, which specifies a common vocabulary and semantics for the knowledge in the system. Such approaches require a language for representing knowledge, such as the Knowledge Interchange Format (KIF) [37], as well as a language for communicating knowledge, such as the Knowledge Query and Manipulation Language (KQML) [30]. This paradigm achieves communication reduction by making use of shared assumptions of vocabulary and semantics.

This type of paradigm can be used for many types of interactions, including cooperative, collaborative, and coordinative. However, while this paradigm has become perhaps the dominant paradigm in multi-agent systems, it is not commonly used in multi-robot systems, at least in the form of full-fledged ontologies. More than likely, this is because physical robot systems are more challenged by noise and uncertainty in sensing and actuation, as well as low-bandwidth communications, limited power, and limited computation. As such, the limiting bottleneck in multi-robot systems is not typically the semantics of the shared knowledge, but rather dealing with these uncertainties. However, this does not mean that multi-robot systems do not use knowledge-based approaches. On the contrary, many approaches do model information about the system and about the teammates in order to more effectively cooperate, collaborate, and coordinate.

IV. CONTRASTING PARADIGMS FOR A TYPICAL MULTI-ROBOT CHALLENGE: TASK ALLOCATION

Having explored three common paradigms in systems of distributed intelligence, we now compare and contrast these paradigms in their approach to a common challenge in multi-robot systems — that of *task allocation*. As previously introduced, task allocation arises in many multi-robot applications in which the mission of the team is defined as a set of tasks that must be completed. Each task can usually be addressed by a variety of different robots; conversely, each robot can usually work on a variety of different tasks. Independent tasks can be achieved concurrently, while dependent tasks must be achieved according to their interdependence constraints. Once the set of tasks has been defined, the challenge is to determine the preferred mapping of robots to tasks that optimizes some objective function. This is the task allocation problem. (See Gerkey [38] for a taxonomy and formal analysis of the computational complexity of variants of the Multi-Robot Task Allocation (MRTA) problem.) The general task

allocation problem is known to be NP-hard [38], meaning that optimal solutions cannot be found quickly for large problems. Therefore, solutions to this problem are typically approximations that are acceptable in practice. The following subsections examine how each of the paradigms we have discussed would handle the multi-robot task allocation problem.

A. Bioinspired approach to task allocation

The bioinspired approach to task allocation typically assumes large numbers of homogeneous robots that are all interchangeable. In this situation, any robot that is available and senses the need for a task to be performed can select to perform that task (i.e., the task is *allocated* to that robot). Because of *stigmergy*, robots do not have to explicitly communicate to decide which task to undertake. Or, alternatively, robots may broadcast minimal information about their state or their environment without undertaking any conversations with other robots about which robot should perform which tasks. Robots that fail can be replaced by any other available robot. If all robots operate under this principle, then the entire mission is typically accomplished.

An example of this type of task allocation is the work of Balch and Arkin [6]. While this work focused on the effect of communication on a society of robots, their work also achieved implicit task allocation through stigmergy and minimal inter-robot communication of robot state and/or goal information. The tasks studied in this work were *Forage*, *Consume*, and *Graze*. The *Forage* task requires robots to wander in the environment, seeking out an *attractor* object. Once found, the robot attaches to the attractor and returns it to home base. The *Consume* task is similar to the *Forage* task, except that the robot performs work in place on the object, rather than return it home. Finally, the *Graze* task requires the robots to collectively visit all areas of the environment. To accomplish these tasks, the robots were programmed with a schema-based reactive control system that enabled robots to avoid obstacles, detect attractors, move to a desired goal destination (e.g., Home or an attractor), etc. Through a variety of simulations that varied the settings of parameters in the system, Balch and Arkin showed that the robots could successfully complete their set of tasks through stigmergy, when available, or through the use of minimal state communication when sensing through the environment (i.e., stigmergy) was not possible.

Most other bioinspired approaches to achieving robot collectives also implicitly achieve task allocation through stigmergy. Examples of these approaches include [59], [60], [70], [74], [87], [99], [105], [106].

B. Organizational approach to task allocation

An organizational approach to task allocation could make use of *roles*, as we described previously for multi-robot soccer. Each role encompasses several specific tasks, and robots select roles that are best suited for their capabilities. In this case, robots need not be homogeneous, but instead can have a variety of different sensing, computation, and effector capabilities.

An example of role-based task allocation is the work of Simons, et al. [96], who have developed the Distributed Robot

Architecture (DIRA), which allows autonomy in individual robots and facilitates explicit coordination among robots. Their approach is based on a layered architecture, in which each robot's control architecture consists of a *planning* layer that decides how to achieve high-level goals; an *executive* layer that synchronizes agents, sequences tasks and monitors task execution; and a *behavioral* layer that interfaces to the robot's sensors and effectors. Each of these layers interacts with those above and below it. Additionally, robots can interact with each other via direct connections at each of the layers. Task allocation is achieved through a role-based approach, in which each robot takes on one of the roles necessary for the task. This architecture has been demonstrated in a team of three robots — a crane, a roving eye, and a mobile manipulator — performing a construction assembly task. This task requires the robots to work together to connect a beam at a given location. Each robot fulfills a critical role of the overall task. In these demonstrations, a software agent acting in the “foreman” role decides which robot should move the beam at which times. Initially, the crane moves the beam to the vicinity of the emplacement based on encoder feedback. The foreman then sets up a behavioral loop between the roving eye and the crane robot to servo the beam closer to the point of emplacement. Once the beam is close enough, the foreman tasks the roving eye and the mobile manipulator to servo the arm to grasp the beam. After contact is made, the foreman tasks the roving eye and the mobile manipulator to coordinate to servo the beam to the emplacement point, thus completing the task.

Teamwork models for cooperation and collaboration in multi-robot/multi-agent teams often also use roles to achieve task allocation. For example, in the work of [109], roles are used to simplify the assignment and execution monitoring of the activities of agents on the team during the mission. Other role-based approaches include [101], [117], [48], [18], [79].

An alternative organizational approach to task allocation is the market-based approach. In these approaches, robots explicitly communicate to bid for tasks according to their expected contribution to those tasks. Assignments are typically made by greedily assigning each task to the robot that can perform it with the highest utility. The fundamental paradigm for interaction in this case is based upon the Contract Net Protocol [97], which was the first to address the problem of how agents can negotiate to collectively solve a set of tasks. The use of a market-based approach specifically for multi-robot task allocation was first developed in the M+ architecture [12]. In the M+ approach, robots make incremental choices of tasks to perform from the set of *executable* tasks. Tasks are considered executable if all of their predecessor tasks have been accomplished or are underway. For these tasks, robots create their own individual plans, estimating their costs for executing these tasks. The robots then compare their costs to offers announced by other robots. The robot selects to perform the task of lowest cost that it can perform that is better than the cost announced by any other robot. Other aspects of the M+ architecture allow robots to negotiate with other teammates to incrementally adapt their actions to suit the team as a whole, through the use of social rules that facilitate the merging of plans.

A significant amount of recent work addresses more powerful market-based approaches for task allocation [39], [58], [96], [117], [61], [11], [91], [127], [42], [52], [64], [32], [50], [85]. These approaches can deal with more complex task representations, more coupling between tasks, dynamic events, combinatorial auctions, and so forth. See [22] for a complete survey of these approaches.

C. Knowledge-based approach to task allocation

The knowledge-based approach is also used for task allocation in multi-robot teams, through the modeling of teammate capabilities. Many variations are possible, such as in the ALLIANCE approach by Parker [82], in which robots model the ability of team members to perform the tasks of the system by observing team member performance and collecting relevant task quality statistics, such as time to task completion. Robots then use these models to select tasks to perform that benefit the group as a whole. In this approach, explicit communication is not required for the selection of task assignments.

Other techniques are also possible that make use of learned models of teammate capabilities. For example, the COBOS work of Fua and Ge [32] maintains a *task suitability matrix*, which is a team model representing the suitability of each robot on the team to perform each task. Tasks can be broken down into subtasks, and the suitability of a robot to perform a macro task can be calculated from its suitability for performing lower-level *basis* tasks. In this work, suitability is calculated based on the compatibility of a robot's *intrinsic* abilities fundamental to the task; *extrinsic* factors, such as time and distance, can also be incorporated into the task description.

Work by Parker and Tang [85], [110] in the development of ASyMTRe is an example of a *semantic* approach to task allocation. In ASyMTRe, each robot team member possesses a set of building-block capabilities, called *schemas*. Examples of schemas include *perceptual schemas*, which extract information from sensors, and *motor schemas*, which convert information from perceptual schemas to motor command outputs. In this approach, each schema's inputs and outputs represent semantic *information types* that must be obtained (in the case of inputs), or are generated (in the case of outputs). This semantic information can be obtained from any source, such as from a schema on the same robot, or a schema on a different robot. By determining a valid information flow through a combination of schemas within, and across, robot team members, the team as a whole is able to compose *coalitions* for solving given tasks. The use of semantic information in this approach enables sensor sharing across heterogeneous team members in a manner that is flexible and extendible to a variety of team compositions and team tasks.

D. Summary of contrasting paradigms for task allocation

As we see through these task allocation examples, a specific problem in multi-robot systems can be addressed in many different ways, based upon the paradigm selected for abstracting the problem at hand. Each paradigm has its own advantages and disadvantages, which may be specific to the application. The paradigms are therefore not interchangeable for many

applications, with the most suitable approach depending upon the relevant constraints and requirements of the application.

V. CONCLUSIONS

In this paper, we have outlined aspects of the field of distributed intelligence, focusing on the types of interactions that can occur in such systems, and some common paradigms used to achieve distributed intelligence. To explore the challenges, we have used examples from the field of multi-robot systems to illustrate, compare, and contrast the alternative interactions and paradigms. The main message of these discussions is that the choice of paradigm is not always obvious, and is dependent upon the requirements of the application to be addressed. We also note that complex systems of multiple robots can make use of several different paradigms simultaneously. For example, a large-scale exploration, mapping, deployment, and detection problem, such as that described in [46], can make use of an organizational paradigm to define roles for the high-level abstraction, an application-specific knowledge-based approach for multi-robot mapping, a knowledge-based modeling approach for mobile network deployment, and a bioinspired approach for creating a mobile sensor network. The challenge as system designers is to create and make use of the appropriate paradigms that best address the specific constraints and challenges of the application at hand.

REFERENCES

- [1] R. Alami, S. Fleury, M. Herrb, F. Ingrand, and F. Robert. Multi-robot cooperation in the MARTHA project. *Robotics and Automation Magazine*, 5(1):36–47, 1998.
- [2] G. Antonelli and S. Chiaverini. Kinematic control of platoons of autonomous vehicles. *IEEE Transactions on Robotics*, 22(6):1285–1292, 2006.
- [3] H. Asama, K. Ozaki, H. Itakura, A. Matsumoto, Y. Ishida, and I. Endo. Collision avoidance among multiple mobile robots based on rules and communication. In *Proceedings of IEEE/RJS International Conference on Intelligent Robots and Systems*. IEEE, 1991.
- [4] T. Balch. The impact of diversity on performance in robot foraging. In *Proceedings of the Third Annual Conference on Autonomous Agents*, pages 92–99. ACM Press, 1999.
- [5] T. Balch and R. Arkin. Behavior-based formation control for multi-robot teams. *IEEE Transactions on Robotics and Automation*, 14(6):926–939, December 1998.
- [6] T. Balch and R. C. Arkin. Communication in reactive multiagent robotic systems. *Autonomous Robots*, 1(1):27–52, 1995.
- [7] R. W. Beard, T. W. McLain, and M. Goodrich. Coordinated target assignment and intercept for unmanned air vehicles. In *Proceedings of IEEE International Conference on Robotics and Automation*. IEEE, 2002.
- [8] R. Beckers, O. Holland, and J. Deneubourg. From local actions to global tasks: Stigmergy and collective robotics. In R. Brooks and P. Maes, editors, *Proceedings of the 14th International Workshop on Synthesis and Simulation of Living Systems*, pages 181–189. MIT Press, 1994.
- [9] C. Belta and V. Kumar. Abstraction and control for groups of robots. *IEEE Transactions on Robotics*, 20(5):865–875, 2004.
- [10] G. Beni and J. Wang. Swarm intelligence in cellular robotics systems. In *Proceedings of NATO Advanced Workshop on Robots and Biological System*, 1989.
- [11] M. Berhault, H. Huang, P. Keskinocak, S. Koenig, W. Elmaghraby, P. Griffin, and A. Kleywegt. Robot exploration with combinatorial auctions. In *Proceedings of IEEE/RJS International Conference on Intelligent Robots and Systems*, pages 1957–1962. IEEE, 2003.
- [12] S. Botelho and R. Alami. M+: A scheme for multi-robot cooperation through negotiated task allocation and achievement. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 1234–1239. IEEE, 1999.
- [13] R. A. Brooks. A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, RA-2(1):14–23, March 1986.
- [14] Rodney A. Brooks. New approaches to robotics. *Science*, 253:1227–1232, September 1991.
- [15] B. Browning, J. Bruce, M. Bowling, and M. Veloso. STP: Skills, tactics and plays for multi-robot control in adversarial environments. *IEEE Journal of Control and Systems Engineering*, 219:33–52, 2005.
- [16] Z. J. Butler, A. A. Rizzi, and R. L. Hollis. Cooperative coverage of rectilinear environments. In *Proceedings of IEEE International Conference on Robotics and Automation*. IEEE, 2000.
- [17] Y. Cao, A. Fukunaga, and A. Kahng. Cooperative mobile robotics: Antecedents and directions. *Autonomous Robots*, 4:1–23, 1997.
- [18] L. Chaimowicz, B. Grocholsky, J. F. Keller, V. Kumar, and C. J. Taylor. Experiments in multirobot air-ground coordination. In *Proceedings of IEEE International Conference on Robotics and Automation*. IEEE, 2004.
- [19] C. M. Clark, S. M. Rock, and J.-C. Latombe. Motion planning for multiple mobile robot systems using dynamic networks. In *Proceedings of IEEE International Conference on Robotics and Automation*, pages 4222–4227, 2003.
- [20] P. R. Cohen and H. Levesque. Teamwork. *Nous*, 25(4):487–512, 1991.
- [21] J. Deneubourg, S. Goss, G. Sandini, F. Ferrari, and P. Dario. Self-organizing collection and transport of objects in unpredictable environments. In *Japan-U.S.A. Symposium on Flexible Automation*, pages 1093–1098, Kyoto, Japan, 1990.
- [22] B. Dias, R. Zlot, N. Kalra, and A. Stentz. Market-based multirobot coordination: A survey and analysis. *Proceedings of the IEEE*, 94(7):1257–1270, July 2006.
- [23] A. Drogoul and J. Ferber. From Tom Thumb to the Dockers: Some experiments with foraging robots. In *Proceedings of the Second International Conference on Simulation of Adaptive Behavior*, pages 451–459, 1992.
- [24] G. Dudek, M. Jenkin, E. Milios, and D. Wilkes. A taxonomy for multi-agent robotics. *Autonomous Robots*, 3:375–397, 1996.
- [25] M. Erdmann and T. Lozano-Perez. On multiple moving objects. *Algorithmica*, 2:477–521, 1987.
- [26] W. Etkin. *Social Behavior from Fish to Man*. The University of Chicago Press, Chicago, 1964.
- [27] H. R. Everett, R. T. Laird, D. M. Carroll, G. A. Gilbreath, T. A. Heath-Pastore, R. S. Inderieden, T. Tran, K. J. Grant, and D. M. Jaffee. Multiple Resource Host Architecture (MRHA) for the Mobile Detection Assessment Response System (MDARS). In *SPAWAR Systems Technical Document 3026, Revision A*, 2000.
- [28] A. Farinelli, L. Iocchi, and D. Nardi. Multi-robot systems: A classification focused on coordination. *IEEE Transactions on Systems, Man, and Cybernetics B*, 34(5):2015–2028, 2004.
- [29] J. A. Fax and R. M. Murray. Information flow and cooperative control of vehicle formations. *IEEE Transactions on Automatic Control*, 49(9), 2004.
- [30] T. Finin, Y. Labrou, and J. Mayfield. KQML as an agent communication language. In J. Bradshaw, editor, *in Software Agents*. MIT Press, 1995.
- [31] M. Fontan and M. Mataric. Territorial multi-robot task division. *IEEE Transactions on Robotics and Automation*, 15(5):815–822, 1998.
- [32] C.-H. Fua and S. S. Ge. COBOS: Cooperative Backoff Adaptive Scheme for Multirobot Task Allocation. *IEEE Transactions on Robotics*, 21(6):1168–1178, 2005.
- [33] T. Fukuda, S. Nakagawa, Y. Kawauchi, and M. Buss. Self organizing robots based on cell structures — CEBOT. In *Proceedings of IEEE International Workshop on Intelligent Robots and Systems*, pages 145–150. IEEE, 1988.
- [34] Douglas Gage. Randomized search strategies with imperfect sensors. In *Proceedings of SPIE Mobile Robots VIII*, pages 270–279. SPIE, September 1993.
- [35] V. Gazi. Swarm aggregations using artificial potentials and sliding-mode control. *IEEE Transactions on Robotics*, 21(6):1208–1214, 2005.
- [36] S. S. Ge and C.-H. Fua. Queues and artificial potential trenches for multirobot formations. *IEEE Transactions on Robotics*, 21(4):646–656, 2005.
- [37] M. R. Genesereth and R. E. Fikes. *Knowledge Interchange Format, Reference Manual*. Computer Science Department, Univ. of Stanford, 1992.
- [38] B. Gerkey and M. J. Mataric. A formal analysis and taxonomy of task allocation in multi-robot systems. *International Journal of Robotics Research*, 23(9):939–954, 2004.

- [39] B. P. Gerkey and M. J. Mataric. Sold! auction methods for multi-robot coordination. *IEEE Transactions on Robotics and Automation*, 18(5):758–768, 2002.
- [40] D. Grossman. Traffic control of multiple robot vehicles. *IEEE Journal of Robotics and Automation*, 4:491–497, 1988.
- [41] B. J. Grosz and S. Kraus. Collaborative plans for complex group action. *Artificial Intelligence*, 86:269–357, 1996.
- [42] J. Guerrero and G. Oliver. Multi-robot task allocation strategies using auction-like mechanisms. In *Proc. of Sixth Congr. Catalan Association for Artificial Intelligence*, pages 111–122, 2003.
- [43] Y. Guo and L. E. Parker. A distributed and optimal motion planning approach for multiple mobile robots. In *Proceedings of IEEE International Conference on Robotics and Automation*, 2002.
- [44] Y. Guo, L. E. Parker, and R. Madhavan. Towards collaborative robots for infrastructure security applications. In *Proceedings of International Symposium on Collaborative Technologies and Systems*, pages 235–240, 2004.
- [45] C. Hazard, P. R. Wurman, and R. D’Andrea. Alphabet soup: A testbed for studying resource allocation in multi-vehicle systems. In *Proceedings of AAAI Workshop on Auction Mechanisms for Robot Coordination*, pages 23–30. AAAI, 2006.
- [46] A. Howard, L. E. Parker, and G. S. Sukhatme. Experiments with a large heterogeneous mobile robot team: Exploration, mapping, deployment, and detection. *International Journal of Robotics Research*, 25:431–447, 2006.
- [47] A. Jadbabaie, J. Lin, and A. S. Morse. Coordination of groups of mobile autonomous agents using nearest neighbor rules. *IEEE Transactions on Automatic Control*, 48(6), 2003.
- [48] J. Jennings and C. Kirkwood-Watts. Distributed mobile robotics by the method of dynamic teams. In T. Leuth, R. Dillman, P. Dario, and H. Worn, editors, *Proc. of Fourth International Symposium on Distributed Autonomous Robotic Systems*. Springer, 1998.
- [49] J. S. Jennings, G. Whelan, and W. F. Evans. Cooperative search and rescue with a team of mobile robots. In *Proceedings of the 8th International Conference on Advanced Robotics*, pages 193–200, 1997.
- [50] E. G. Jones, B. Browning, M. B. Dias, B. Argall, M. Veloso, and A. Stentz. Dynamically formed heterogeneous robot teams performing tightly-coupled tasks. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 570–575. IEEE, 2006.
- [51] B. Jung and G. Sukhatme. Tracking targets using multiple mobile robots: The effect of environment occlusion. *Autonomous Robots*, 13(3):191–205, 2002.
- [52] N. Kalra, D. Ferguson, and A. Stentz. Hoplitest: A market-based framework for planned tight coordination in multirobot teams. In *Proceedings of the IEEE International Conference on Robotics and Automation*. IEEE, 2005.
- [53] K. Kant and S. W. Zucker. Toward efficient trajectory planning: the path-velocity decomposition. *The International Journal of Robotics Research*, 5(3):72–89, 1986.
- [54] S. Kato, S. Nishiyama, and J. Takeno. Coordinating mobile robots by applying traffic rules. In *Proceedings of the 1992 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1535–1541, July 1992.
- [55] H. Kitano, M. Asada, Y. Kuniyoshi, I. Noda, E. Osawa, and H. Matsuura. Robocup: A challenge problem of ai. *AI Magazine*, 18(1):73–86, 1997.
- [56] S. Kloder and S. Hutchinson. Path planning for permutation-invariant multirobot formations. *IEEE Transactions on Robotics*, 22(4):650–665, 2006.
- [57] A. Kolling and S. Carpin. Multirobot cooperation for surveillance of multiple moving targets – a new behavioral approach. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 1311–1316. IEEE, 2006.
- [58] H. Kose, U. Tatlidede, C. Mericli, K. Kaplan, and H. L. Akin. Q-learning based market-driven multi-agent collaboration in robot soccer. In *Proceedings of the Turkish Symposium on Artificial Intelligence and Neural Networks*, pages 219–228, 2004.
- [59] C. R. Kube and H. Zhang. Collective robotics: From social insects to robots. *Adaptive Behavior*, 2(2):189–219, 1993.
- [60] Masao Kubo and Yukinori Kakazu. Learning coordinated motions in a competition for food between ant colonies. In D. Cliff, P. Husbands, J.-A. Meyer, and S. Wilson, editors, *Proceedings of the Third International Conference on Simulation of Adaptive Behavior*, pages 487–492. MIT Press, 1994.
- [61] M. Lagoudakis, E. Markakis, D. Kempe, P. Keshinocak, A. Kleywegt, S. Koenig, C. Tovey, A. Meyerson, and S. Jain. Auction-based multi-robot routing. In *Robotics: Science and Systems I*. MIT Press, 2005.
- [62] S. M. LaValle, H. H. Gonzalez-Banos, C. Becker, and J.-C. Latombe. Motion strategies for maintaining visibility of a moving target. In *Proceedings of the 1997 IEEE International Conference on Robotics and Automation*, pages 731–736. IEEE, 1997.
- [63] S. M. LaValle and S. A. Hutchinson. Optimal motion planning for multiple robots having independent goals. *IEEE Transactions on Robotics and Automation*, 14:912–925, 1998.
- [64] L. Lin and Z. Zheng. Combinatorial bids based multi-robot task allocation method. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 1145–1150. IEEE, 2005.
- [65] S. Luke, K. Sullivan, L. Panait, and G. Balan. Tunably decentralized algorithms for cooperative target observation. In *Proceedings of the fourth international joint conference on Autonomous Agents and Multiagent Systems*, pages 911–917. ACM Press, 2005.
- [66] V. J. Lumelsky and K. R. Harinarayan. Decentralized motion planning for multiple mobile robots: The cocktail party model. *Autonomous Robots*, 4(1):121–135, 1997.
- [67] S. Marsella, J. Adibi, Y. Al-Onaizan, G. Kaminka, I. Muslea, and M. Tambe. On being a teammate: Experiences acquired in the design of RoboCup teams. In O. Etzioni, J. Muller, and J. Bradshaw, editors, *Proceedings of the Third Annual Conference on Autonomous Agents*, pages 221–227, 1999.
- [68] J. A. Marshall, M. E. Broucke, and B. R. Francis. Formations of vehicles in cyclic pursuit. *IEEE Transactions on Automatic Control*, 49(11), 2004.
- [69] M. Mataric. Behavior-based control: Examples from navigation, learning, and group behavior. *Journal of Experimental and Theoretical Artificial Intelligence*, 19(2-3):323–336, 1997.
- [70] M. J. Mataric. Designing emergent behaviors: From local interactions to collective intelligence. In J. Meyer, H. Roitblat, and S. Wilson, editors, *Proceedings of the Second International Conference on Simulation of Adaptive Behavior*, pages 432–441. MIT Press, 1992.
- [71] M. J. Mataric. Issues and Approaches in the Design of Collective Autonomous Agents. *Robotics and Autonomous Systems*, 16:321–331, 1995.
- [72] D. McFarland. Towards robot cooperation. In D. Cliff, P. Husbands, J.-A. Meyer, and S. Wilson, editors, *Proceedings of the Third International Conference on Simulation of Adaptive Behavior*, pages 440–444. MIT Press, 1994.
- [73] J. McLurkin. Stupid robot tricks: Behavior-based distributed algorithm library for programming swarms of robots. In *M.S. Thesis, Massachusetts Institute of Technology*, 2004.
- [74] J. McLurkin and J. Smith. Distributed algorithms for dispersion in indoor environments using a swarm of autonomous mobile robots. In *Symposium on Distributed Autonomous Robotic Systems*. Springer, 2004.
- [75] A. I. Mourikis and S. I. Roumeliotis. Optimal sensor scheduling for resource-constrained localization of mobile robot formations. *IEEE Transactions on Robotics*, 22(5):917–931, 2006.
- [76] A. I. Mourikis and S. I. Roumeliotis. Performance analysis of multirobot cooperative localization. *IEEE Transactions on Robotics*, 22(4):666–681, 2006.
- [77] R. R. Murphy. Marsupial robots for urban search and rescue. *IEEE Intelligent Systems*, 15(2):14–19, 2000.
- [78] R. R. Murphy, C. Lisetti, R. Tardif, L. Irish, and A. Gage. Emotion-based control of cooperating heterogeneous mobile robots. *IEEE Transactions on Robotics and Automation*, 18(5):744–757, 2002.
- [79] E. Pagello, A. D’Angelo, and E. Menegatti. Cooperation issues and distributed sensing for multirobot systems. *Proceedings of the IEEE*, 94:1370–1383, 2006.
- [80] L. E. Parker. Adaptive action selection for cooperative agent teams. In Jean-Arcady Meyer, Herbert Roitblat, and Stewart Wilson, editors, *Proceedings of the Second International Conference on Simulation of Adaptive Behavior*, pages 442–450. MIT Press, 1992.
- [81] L. E. Parker. Designing control laws for cooperative agent teams. In *Proceedings of the IEEE Robotics and Automation Conference*, pages 582–587. IEEE, 1993.
- [82] L. E. Parker. ALLIANCE: An architecture for fault-tolerant multi-robot cooperation. *IEEE Transactions on Robotics and Automation*, 14(2):220–240, 1998.
- [83] L. E. Parker and J. Draper. Robotics applications in maintenance and repair. In S. Nof, editor, *Handbook of Industrial Robotics*, pages 1023–1036. Wiley Publishers, 2nd edition, 1999.
- [84] L. E. Parker, B. Kannan, F. Tang, and M. Bailey. Tightly-coupled navigation assistance in heterogeneous multi-robot teams. In *Proceedings of IEEE International Conference on Intelligent Robots and Systems*. IEEE, 2004.

- [85] L. E. Parker and F. Tang. Building multi-robot coalitions through automated task solution synthesis. *Proceedings of the IEEE, special issue on Multi-Robot Systems*, 94(7):1289–1305, 2006.
- [86] D. Parsons and J. Canny. A motion planner for multiple mobile robots. In *Proceedings of IEEE International Conference on Robotics and Automation*, pages 8–13, 1990.
- [87] K. Passino. Biomimicry of bacterial foraging for distributed optimization and control. *IEEE Control Systems Magazine*, 22(3):52–67, 2002.
- [88] M. Peasgood, C. Clark, and J. McPhee. A complete and scalable strategy for coordinating multiple robots within roadmaps. *IEEE Transactions on Robotics*, to appear, 2008.
- [89] T. Pilarski, M. Happold, H. Pangels, M. Ollis, K. Fitzpatrick, and A. Stentz. The demeter system for automated harvesting. In *Proceedings of the 8th International Topical Meeting on Robotics and Remote Systems*, 1999.
- [90] A. Portmann. *Animals as Social Beings*. The Viking Press, New York, 1961.
- [91] G. Rabideau, T. Estlin, S. Schien, and A. Barrett. A comparison of coordinated planning methods for cooperating rovers. In *Proceedings of AIAA Space Technology Conference*, 1999.
- [92] C. W. Reynolds. Flocks, herds and schools: A distributed behavioral model. *ACM SIGGRAPH Computer Graphics*, 21:25–34, 1987.
- [93] M. Rude. Collision avoidance by using space-time representations of motion processes. *Autonomous Robots*, 4:101–119, 1997.
- [94] P. Rybski, S. Stoeter, C. Wyman, and M. Gini. A cooperative multi-robot approach to the mapping and exploration of mars. In *Proceedings of AAAI/IAAI-97*. AAAI, 1997.
- [95] G. Shaffer and A. Stentz. A robotic system for underground coal mining. In *Proceedings of IEEE International Conference on Robotics and Automation*, pages 633–638, 1992.
- [96] R. Simmons, S. Singh, D. Hershberger, J. Ramos, and T. Smith. First results in the coordination of heterogeneous robots for large-scale assembly. In *Proc. of the ISER Seventh International Symposium on Experimental Robotics*. Springer-Verlag, 2000.
- [97] R. G. Smith. The Contract Net Protocol: high-level communication and control in a distributed problem solver. *IEEE Transactions on Computers*, C-29(12), December 1980.
- [98] L. Steels. Cooperation between distributed agents through self-organization. In Yves Demazeau and Jean-Pierre Muller, editors, *Decentralized A.I.* Elsevier Science, 1990.
- [99] D. Stilwell and J. Bay. Toward the development of a material transport system using swarms of ant-like robots. In *Proceedings of IEEE International Conference on Robotics and Automation*, pages 766–771, 1993.
- [100] P. Stone and M. Veloso. A layered approach to learning client behaviors in the robocup soccer server. *Applied Artificial Intelligence*, 12:165–188, 1998.
- [101] P. Stone and M. Veloso. Task decomposition, dynamic role assignment, and low-bandwidth communication for real-time strategic teamwork. *Artificial Intelligence*, 110(2):241–273, 1999.
- [102] P. Stone and M. Veloso. Multiagent systems: A survey from a machine learning perspective. *Autonomous Robots*, 8(3):345–383, 2000.
- [103] A. Stroupe, A. Okon, M. Robinson, T. Huntsberger, H. Aghazarian, and E. Baumgartner. Sustainable cooperative robotic technologies for human and robotic outpost infrastructure construction and maintenance. *Autonomous Robots*, 20(2):113–123, 2006.
- [104] K. Sugawara and M. Sano. Cooperative behavior of interacting simple robots in a clockface arranged foraging field. In H. Asama, T. Arai, T. Fukuda, and T. Hasegawa, editors, *Distributed Autonomous Robotic Systems*, pages 331–339. Springer, 2002.
- [105] K. Sugihara and I. Suzuki. Distributed algorithms for formation of geometric patterns with many mobile robots. *Journal of Robotic Systems*, 13(3):127–139, 1996.
- [106] S. Sun, D. Lee, and K. Sim. Artificial immune-based swarm behaviors of distributed autonomous robotic systems. In *Proceedings of IEEE International Conference on Robotics and Automation*, pages 3993–3998. IEEE, 2001.
- [107] P. Svestka and M. Overmars. Coordinated path planning for multiple robots. *Robotics and Autonomous Systems*, 23:125–152, 1998.
- [108] P. Tabuada, G. Pappas, and P. Lima. Motion feasibility of multi-agent formations. *IEEE Transactions on Robotics*, 21(3):387–392, 2005.
- [109] M. Tambe. Towards flexible teamwork. *Journal of Artificial Intelligence Research*, 7:83–124, 1997.
- [110] F. Tang and L. E. Parker. A complete methodology for generating multi-robot task solutions using asymptotic and market-based task allocation. In *Proceedings of IEEE International Conference on Robotics and Automation (ICRA)*, 2007.
- [111] Z. Tang and U. Ozguner. Motion planning for multitarget surveillance with mobile sensor agents. *IEEE Transactions on Robotics*, 21(5):898–908, 2005.
- [112] Guy Theraulaz, Simon Goss, Jacques Gervet, and Jean-Louis Deneubourg. Task differentiation in *Polistes* wasp colonies: a model for self-organizing groups of robots. In *Proceedings of the First International Conference on Simulation of Adaptive Behavior*, pages 346–355, 1990.
- [113] C. Thorpe, T. Jochem, and D. Pomerleau. The 1997 automated highway free agent demonstration. In *Proceedings of IEEE Conference on Intelligent Transportation System*, pages 496–501, 1997.
- [114] N. Tinbergen. *Social Behavior in Animals*. Chapman and Hall LTD, Great Britain, 1953.
- [115] N. Tinbergen. *Animal Behavior*. Time-Life Books, New York, 1965.
- [116] C. M. Topaz and A. L. Bertozzi. Swarming patterns in two-dimensional kinematic model for biological groups. *SIAM Journal of Applied Math*, 65(1):152–174, 2004.
- [117] D. Vail and M. Veloso. Multi-robot dynamic role assignment and coordination through shared potential fields. In A. Schultz, L. E. Parker, and F. Schneider, editors, *Multi-Robot Systems: From Swarms to Intelligent Automata: Proceedings of 2003 International Workshop on Multi-Robot Systems*, pages 87–98. Springer, 2003.
- [118] M. Veloso, P. Stone, and K. Han. The CMUnited-97 robotic soccer team: Perception and multiagent control. *Robotics and Autonomous Systems*, 29(2-3):133–143, 1999.
- [119] L. Vig and J. A. Adams. Multi-robot coalition formation. *IEEE Transactions on Robotics*, 22(4):637–649, 2006.
- [120] I. Wagner, M. Lindenbaum, and A. M. Bruckstein. Mac vs. PC – determinism and randomness as complementary approaches to robotic exploration of continuous unknown domains. *International Journal of Robotics Research*, 19(1):12–31, 2000.
- [121] Jing Wang. Fully distributed traffic control strategies for many-AGV systems. In *Proceedings of the IEEE International Workshop on Intelligent Robots and Systems*, pages 1199–1204. IEEE, 1991.
- [122] Jing Wang and Gerardo Beni. Distributed computing problems in cellular robotic systems. In *Proceedings of the IEEE International Workshop on Intelligent Robots and Systems*, pages 819–826, Tsuchiura, Japan, 1990.
- [123] B. B. Werger and M. J. Matarić. Broadcast of local eligibility for multi-target observation. In L. E. Parker, G. Bekey, and J. Barhen, editors, *Distributed Autonomous Robotic Systems 4*, pages 347–356. Springer, 2000.
- [124] E. Wilson. *The Insect Societies*. The Belknap Press, Cambridge, 1971.
- [125] S. Yuta and S. Premvuti. Coordinating autonomous and centralized decision making to achieve cooperative behaviors between multiple mobile robots. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1566–1574. IEEE, 1992.
- [126] R. Zlot and A. Stentz. Market-based multirobot coordination for complex tasks. *International Journal of Robotics Research*, 25(1):73–101, January 2006.
- [127] R. Zlot, A. Stentz, M. B. Dias, and S. Thayer. Multi-robot exploration controlled by a market economy. In *Proceedings of IEEE International Conference on Robotics and Automation*, pages 3016–3023. IEEE, 2002.