Cooperative Mobile Robotics: Antecedents and Directions*

Y. Uny Cao, Alex S. Fukunaga, Andrew B. Kahng and Frank Meng
UCLA Computer Science Department, Los Angeles, CA 90024-1596

Abstract

There has been increased research interest in systems composed of multiple autonomous mobile robots exhibiting collective behavior. Groups of mobile robots are constructed, with an aim to studying such issues as group architecture, resource conflict, origin of cooperation, learning, and geometric problems. As yet, few applications of collective robotics have been reported, and supporting theory is still in its formative stages. In this paper, we give a critical survey of existing works and discuss open problems in this field, emphasizing the various theoretical issues that arise in the study of cooperative robotics. We describe the intellectual heritages that have guided early research, as well as possible additions to the set of existing motivations.

1 Preliminaries

There has been much recent activity toward achieving systems of multiple mobile robots engaged in collective behavior. Such systems are of interest for several reasons: (1) tasks may be inherently too complex for a single robot to accomplish, or performance benefits can be gained from using multiple robots; (2) building and using several simple robots can be easier, cheaper, more flexible and more fault-tolerant than having a single powerful robot for each separate task; and (3) insight into social sciences (organization theory, economics), life sciences (theoretical biology, animal ethology) and cognitive science (psychology, learning, artificial intelligence) may be derived from multi-robot experimental systems.

The study of multiple robots naturally extends research on single-robot systems, but is also a discipline unto itself: multiple-robot systems can accomplish tasks that no single robot can accomplish, since ultimately a single robot, no matter how capable, is spatially limited. Multiple-robot systems are also different from other distributed systems because of their implicit "real-world" environment, which is presumably more difficult to model and reason about than traditional components of distributed system environments (i.e., computers, databases, networks).

The term *collective behavior* generically denotes any behavior of agents in a system having more than one agent. *Cooperative behavior*, which is the subject of the present survey, is a subclass of collective behavior that is characterized by cooperation. Webster's dictionary [MW63] defines "cooperate" as "to associate with another or others for mutual, often economic, benefit". Explicit definitions of cooperation in the robotics literature, while surprisingly sparse, include: (a) "joint collaborative behavior that is directed toward some goal in which there is a common interest or reward" [BG91]; (b) "a form of interaction, usually based on communication" [Mat94a]; and (c) "[joining] together for doing something that creates a progressive result such as increasing performance or saving time" [PY90]

These definitions show the wide range of possible motivating perspectives. For example, definitions such as (a) typically lead to the study of task decomposition, task allocation, and other distributed artificial intelligence (DAI) issues (e.g., learning, rationality). Definitions along the lines of (b) reflect a concern with requirements for information or other resources, and may be accompanied by studies of related issues such as correctness and fault-tolerance.[1] Finally, definition (c) reflects a concern with quantified measures of cooperation, such as speedup in time to complete a task. Thus, in these definitions we see three fundamental seeds: the *task*, the *mechanism* of cooperation, and system *performance*.

We define cooperative behavior as follows: *Given some task specified by a designer, a multiple-robot system displays* cooperative behavior *if, due to some underlying mechanism (i.e., the "mechanism of cooperation"), there is an increase in the total utility of the system.* Intuitively, cooperative behavior entails some type of performance gain over naive collective behavior. The mechanism of cooperation may lie in the imposition by the designer of a control or communication structure, in aspects of the task specification, in the interaction dynamics of agent behaviors, etc.[2]

In this paper, we survey the intellectual heritage and major research directions of the field of cooperative robotics. For this survey of cooperative robotics to remain tractable, we restrict our discussion to works involving mobile robots *or simulations of mobile robots*, where a *mobile robot* is taken to be an autonomous, physically independent, mobile robot. Thus, we do not discuss coordination of multiple manipulators, articulated arms, multi-finger hands, etc.[3]

Toward a Picture of Cooperative Robotics

In the mid-1940's Grey Walter, along with Wiener and Shannon, studied turtle-like robots equipped with light and touch sensors; these simple robots exhibited "complex social behavior" in responding to each other's movements [Dor90]. Coordination and interactions of multiple intelligent agents have been actively studied in the field of distributed artificial intelligence (DAI) since the early 1970's [BG88], but the DAI field concerned itself mainly with problems involving software agents. In the late 1980's, the robotics research community became very active in cooperative robotics, beginning with projects such as ACTRESS [AMI89], CEBOT [FN87], GOFER [CCL+90], SWARM [Ben88] and the work at Brussels [Ste90]. This early research was done primarily in simulation;[4] thus, several more recent works (cf. [KZ92, Mat92, Par92]) are significant for their emphasis on the actual physical implementation of cooperative robotic systems.

*Partial support for this work was provided by NSF Young Investigator Award MIP-9257982; the UCLA Commotion Laboratory is supported by NSF CDA-9303148.

[1] This is a characterization of the definition itself; the leading motivation in [Mat94a] was actually the social nature of intelligence and its manifestation in group behaviors.

[2] In this work, we do not discuss the *competitive* subclass of collective behavior, which includes pursuit-evasion [Rey94, MC94] and one-on-one competitive games [AUN+94]. Note that a cooperative team strategy for, e.g., the robot soccer league recently proposed by [Kit94] would lie within our present scope.

[3] Even with this restriction, we find that in 7 years (1987 - 1994) well over 200 papers have been published in this field of cooperative (mobile) robotics, encompassing theories from such diverse disciplines as artificial intelligence, game theory/economics, theoretical biology, distributed computing/control, animal ethology and artificial life.

[4] While CEBOT, ACTRESS and GOFER have all had physical implementations (with $\leq$ 3 robots), in some sense these implementations were presented by way of proving the simulation results.

The rapid progress of cooperative robotics since the late 1980's has been an interplay of *systems*, *theories* and *problems*: to solve a given problem, systems are envisioned, simulated and built; theories of cooperation are brought from other fields; and new problems are identified (prompting further systems and theories). Since so much of this progress is recent, it is not easy to discern deep intellectual heritages from within the field. More apparent are the intellectual heritages from other fields, as well as the canonical task domains which have driven research. Three examples of the latter are:

• **Traffic Control.** When multiple agents move within a common environment, they typically attempt to avoid collisions. Fundamentally, this may be viewed as a problem of *resource conflict*, which may be resolved by introducing, e.g., traffic rules, priorities, or communication architectures. From another perspective, path planning must be performed taking into consideration other robots and the global environment; this multiple-robot path planning is an intrinsically *geometric* problem in configuration space-time. Note that prioritization and communication protocols – as well as the internal modeling of other robots – all reflect possible variants of the *group architecture* of the robots.

• **Box-Pushing.** Many works have addressed the box-pushing (or couch-pushing) problem, for widely varying reasons. The focus in [Par94] is on task allocation, fault-tolerance and (reinforcement) *learning*. By contrast, [DJR94] studies two box-pushing protocols in terms of their intrinsic communication and hardware requirements, via the concept of information invariants. Other works in the box-pushing/object genre manipulation include [WNM94, MHB94, SB93, JB94].

• **Foraging.** In foraging, a group of robots must pick up objects scattered in the environment; this is evocative of toxic waste cleanup, harvesting, search and rescue, etc. The foraging task is one of the canonical testbeds for cooperative robotics [Ste90, Ark92, GD92, LB92, DF93, AH93, Mat94a, BHD94].[5] A wide variety of techniques have been proposed, ranging from simple stigmergy [6] [BHD94] to more complex algorithms in which robots form chains along which objects are passed to the goal [DF93]. Again, group architecture and learning are major research themes in addressing this problem.

Organization of Paper

With respect to our above definition of cooperative behavior, we find that the great majority of the cooperative robotics literature centers on the *mechanism* of cooperation. (I.e., few works study a task without also claiming some novel approach to achieving cooperation.) Thus, our study has led to the synthesis of five "Research Axes" which we believe comprise the major themes of investigation to date into the underlying mechanism of cooperation.[7]

Section 2 of this paper describes these axes, which are: 2.1 Group Architecture, 2.2 Resource Conflict, 2.3 Origin of Cooperation, 2.4 Learning, and 2.5 Geometric Problems. In Section 3, we present more synthetic reviews of cooperative robotics: Section 3.1 discusses constraints arising from technological limitations; and Section 3.2 discusses possible lacunae in existing work (e.g., formalisms for measuring performance of a cooperative robot system), then reviews three fields which we believe must strongly influence future work. We conclude in Section 4 with a list of key research challenges facing the field.

2 Research Axes

Seeking a *mechanism* of cooperation may be rephrased as the "cooperative behavior design problem": *Given a group of robots, an environment, and a task, how should cooperative behavior arise?*[8] In some sense, every work in cooperative robotics has addressed facets of this problem, and the major research axes of the field follow from elements of this problem.

First, the realization of cooperative behavior must rely on some infrastructure, the **group architecture**. This encompasses such concepts as robot heterogeneity/homogeneity, the ability of a given robot to recognize and model other robots, and communication structure. Second, for multiple robots to inhabit a shared environment, manipulate objects in the environment, and possibly communicate with each other, a mechanism is needed to resolve **resource conflicts**. The third research axis, **origins of cooperation**, refers to how cooperative behavior is actually achieved.[9] Because adaptability and flexibility are essential traits in a task-solving group of robots, we view **learning** as a fourth key to achieving cooperative behavior.[10] Whereas the first four axes are related to the *generation* of cooperative behavior, our fifth and final axis – **geometric problems** – covers research issues that are tied to the embedding of robot tasks in a two- or three-dimensional world. These issues include multi-agent path planning, moving to formation, and pattern generation.

2.1 Group Architecture

The *architecture* of a computing system has been defined as "the part of the system that remains unchanged unless an external agent changes it" [Van91]. The *group architecture* of a cooperative robotic system provides the infrastructure upon which collective behaviors are implemented and determines the capabilities and limitations of the system. We now briefly discuss some of the key architectural features of a group architecture for mobile robots: centralization/decentralization, differentiation, communications, and

[5] The task is interesting because it can be performed by each robot independently (i.e., the issue is whether multiple robots achieve a performance gain). There are some conceptual overlaps with the related task of materials handling in a manufacturing workcell [DA93].

[6] [BHD94] defines stigmergy as "the production of a certain behaviour in agents as a consequence of the effects produced in the local environment by previous behaviour". This is actually a form of "cooperation without communication", which has been the stated object of several foraging solutions since the corresponding formulations become nearly trivial if communication is used. However, note also that stigmergy may not satisfy our above definition of cooperation, since there is no performance improvement over the "naive algorithm" (in this particular case, the proposed stigmergic algorithm *is* the naive algorithm).

[7] Note that our survey concentrates on fundamental *theoretical* issues that impinge on cooperative robotics. Thus, we do not discuss several important practical concerns, such as the user interface issues that arise with multiple-robot systems [YSA+94, AA94, NR91]. In addition, emerging technologies such as nanotechnology [Dre92] and Micro Electro-Mechanical Systems [MGT88] that are likely to be very important to cooperative robotics are beyond the scope of this paper.

[8] Note that certain basic robot interactions are not task-performing interactions per se, but are rather basic primitives upon which task-performing interactions can be built (e.g., following ([Con87, DJR94] and many others) or flocking [Rey87, Mat94a]). It might be argued that these interactions entail "control and coordination" tasks rather than "cooperation" tasks, but our treatment does not make such a distinction.

[9] We do not discuss instances where cooperation has been "explicitly engineered" into the robots' behavior since this is the default approach. Instead, we are more interested in biological parallels (e.g., to social insect behavior), game-theoretic justifications for cooperation, and concepts of emergence.

[10] One important mechanism in generating cooperation, namely, task decomposition and allocation, is *not* considered a research axis since (i) very few works in cooperative robotics have centered on task decomposition and allocation (with the notable exception of [Par94]), (ii) cooperative robot tasks (foraging, box-pushing) in the literature are simple enough that decomposition and allocation are not required in the solution, and (iii) the use of decomposition and allocation depends almost entirely on the group architectures (e.g. whether it is centralized or decentralized).

the ability to model other agents. We then describe several systems that have addressed these specific problems.

2.1.1 Centralization/Decentralization

The most fundamental decision that is made when defining a group architecture is whether the system is centralized or decentralized, and if it is decentralized, whether the system is hierarchical or distributed. *Centralized* architectures are characterized by a single control agent. *Decentralized* architectures lack such an agent. There are two types of decentralized architectures: *distributed* architectures in which all agents are equal with respect to control, and *hierarchical* architectures which are locally centralized.

Currently, the dominant paradigm is the decentralized approach.[11] The behavior of decentralized systems is often described using such terms as "emergence" and "self-organization."[12] It is widely claimed that decentralized architectures (e.g., [BHD94, Ark92, Ste94, Mat94a]) have several inherent advantages over centralized architectures, including fault tolerance, natural exploitation of parallelism, reliability, and scalability.

2.1.2 Differentiation

We define a group of robots to be *homogeneous* if the capabilities of the individual robots are identical, and *heterogeneous* otherwise. In general, heterogeneity introduces complexity since task allocation becomes more difficult, and agents have a greater need to model other individuals in the group. The literature is currently dominated by works that assume homogeneous groups of robots. However, some notable architectures can handle heterogeneity, e.g., ACTRESS and ALLIANCE (see Section 2.1.5 below). In heterogeneous groups, task allocation may be determined by individual capabilities, but in homogeneous systems, agents may need to *differentiate* into distinct roles that may be known at design-time, or rise dynamically at run-time.

2.1.3 Communication Structures

The communication structure of a group determines the possible modes of interagent interaction. We characterize three types of interactions that can be supported.

Interaction via environment

The simplest, most limited type of interaction occurs when the environment itself is the communication medium (in effect, a shared memory), and there is no explicit communication or interaction between agents. Systems that depend on this form of interaction include [GD92, BHD94, Ark92, Ste90, TK93, Tun94, SSH94].

Interaction via sensing

Corresponding to arms-length relationships in organization theory [Hew93], this refers to local interactions that occur between agents as a result of agents sensing one another, but without explicit communication. Collective behaviors that can use this kind of interaction include flocking and pattern formation (keeping in formation with nearest neighbors). This type of interaction requires that agents can distinguish between other agents in the group and other objects in the environment.

Interaction via communications

The third form of interaction involves explicit communication with other agents, by either directed or broadcast intentional messages.[13] Because architectures that enable this form of communication are similar to networks, many standard issues from the field of networks arise, including the design of network topologies and communications protocols.[14]

2.1.4 Modeling of Other Agents

Modeling the intentions, beliefs, actions, capabilities, and states of other agents can lead to more effective cooperation between robots. Communications requirements can also be lowered if each agent has the capability to model other agents. Note that the modeling of other agents entails more than implicit communication via the environment or perception: modeling requires that the modeler has some representation of another agent, and that this representation can be used to make inferences about the actions of the other agent.

In cooperative robotics, agent modeling has been explored most extensively in the context of manipulating a large object. Many solutions have exploited the fact that the object can serve as a common medium by which the agents can model each other.[15]

2.1.5 Representative Architectures

All systems implement some group architecture. We now describe several particularly well-defined representative architectures, along with works done within each of their frameworks. It is interesting to note that these architectures encompass the entire spectrum from traditional AI to highly decentralized approaches.

SWARM

A *SWARM* is a distributed system with a large number of autonomous robots [JLB94]. *SWARM intelligence* is "a property of systems of non-intelligent robots exhibiting collectively intelligent behavior" [HB91]. [16] *Self-organization* in a SWARM is the ability to distribute itself "optimally" for a given task, e.g., via geometric pattern formation or structural organization. SWARM exhibits a distributed architecture, usually with no differentiation among members.[17] Interaction takes place by each cell reacting to the state of its nearest neighbors.[18] Examples for possible applications include large-scale displays and distributed sensing [HW88]. Some-

[11] We have not found any instances of systems that are centralized.

[12] The definitions of these notions have been the subject of much intellectual debate. We do not define them here, and appeal to their intuitive definition.

[13] In other words, the recipient(s) of the message may be either known or unknown.

[14] For example, [Wan94] proposes inter-robot communication using a media access protocol (similar to Ethernet). In [IHH93], robots with limited communication range communicate to each other using the "hello-call" protocol, by which robots with limited communication ranges establish "chains" in order to extend their effective ranges. [Gag93] describes methods for communicating to many ("zillions") robots. [AYO93] proposes a communication protocol modeled after chemical diffusion. Similar communications mechanisms are studied in [LB92, GD92, DF93].

[15] The second of two box-pushing protocols in [DJR94] can achieve "cooperation without communication" since the object being manipulated also functions as a "communication channel" that is shared by the robot agents; other works capitalize on the same concept to derive distributed control laws which rely only on local measures of force, torque, orientation, or distance, i.e., no explicit communication is necessary (cf. [SB93] [HO93]). In a two-robot bar carrying task, Fukuda and Sekiyama's agents [FS94] each uses a probabilistic model of the other agent. When a risk threshold is exceeded, the agent communicates with its partner to maintain coordination. In [Don93a, Don93b], the theory of information invariants is used to show that extra hardware capabilities can be added in order to infer the actions of the other agent, thus reducing communication requirements. This is in contrast to [SSH94], where the robots achieve box pushing but are not aware of each other at all.

[16] The work on SWARM systems began as work on Cellular Robotic Systems (CRS), where many simple agents occupied one -or two- dimensional environments and were able to perform tasks such as pattern generation and self-organization.

[17] An exception is [HB92], where two different types of robots were used.

[18] Mechanisms for self-organization in SWARM are studied in [HB92, BW91, BH92, HB91, JLB94].

times, broadcast communication is performed as in [WB88], which uses a signboard mechanism.

CEBOT

CEBOT (CEllular RoBOTics System) is a decentralized, hierarchical architecture inspired by the cellular organization of biological entities [FN87, FK93, UF93]. The system is dynamically reconfigurable in that basic autonomous "cells" (robots), which can be physically coupled to other cells, dynamically reconfigure their structure to an "optimal" configuration in response to changing environments. In the CEBOT hierarchy there are "master cells" that coordinate subtasks and communicate with other master cells.[19] Communications requirements have been studied extensively with respect to the CEBOT architecture, and various methods have been proposed that seek to reduce communication requirements by making individual cells more intelligent (e.g., enabling them to model the behavior of other cells).[20]

ALLIANCE/L-ALLIANCE

The ALLIANCE architecture was developed by Parker [Par94] in order to study cooperation in a heterogeneous, small-to-medium-sized team of largely independent, loosely coupled robots. Robots are assumed able to, with some probability, sense the effects of their own actions and the actions of other agents through perception and explicit broadcast communications. Individual robots are based on a behavior-based controller with an extension for activating "behavior sets" that accomplish certain tasks. These sets are activated by motivational behaviors whose activations are in turn determined by the robots' awareness of their teammates. L-ALLIANCE [Par94] is an extension to the architecture that uses reinforcement learning to adjust the parameters controlling behavior set activation. The ALLIANCE/L-ALLIANCE architecture has been implemented both on real robots and in simulation, and has been successfully demonstrated for tasks including box-pushing, puck-gathering, marching in formation, and simulations of hazardous waste cleanup and janitorial service.

Behavior-Based Cooperative Behavior

Mataric [Mat94a] proposes a behavior-based architecture for the synthesis of collective behaviors such as flocking, foraging, and docking based on the direct and temporal composition of primitive basic behaviors (safe-wandering, following, aggregation, dispersion, homing). A method for automatically constructing composite behaviors based on reinforcement learning is also proposed. The architecture has been implemented both on groups of up to 20 real robots (the largest group reported in the works we surveyed), and in simulation.

GOFER

The GOFER architecture [CCL+90, LeP90] was used to study distributed problem solving by multiple mobile robots in an indoor environment using traditional AI techniques. In GOFER, a central task planning and scheduling system (CTPS) communicates with all robots and has a global view of both the tasks to be performed and the availability of robots to perform the tasks. The CTPS generates a plan structure (template for an instance of a plan) and informs all available robots of the pending goals and plan structures. Robots use a task allocation algorithm like the Contract Net Protocol [Smi80] to determine their roles. Given the goals assigned during the task allocation process, they attempt to achieve their goals using fairly standard AI planning techniques. The GOFER architecture was successfully used with three physical robots for tasks such as following, box-pushing, and wall tracking in a corridor.

2.2 Resource Conflict

When a single indivisible resource is requested by multiple robots, *resource conflict* arises. This issue has been studied in many guises, notably the mutual exclusion problem in distributed algorithms and the multiaccess problem in computer networks. With multiple robots, resource conflict occurs when there is a need to share space, manipulable objects or communication media. Few works have dealt specifically with object sharing or sharing of communication media. We therefore center on the space sharing problem, which has been studied primarily via multiple-robot path planning (the "traffic control" formulation from above) and the collision and deadlock avoidance problems.

In a multi-robot system, each robot can conceivably plan a path that accounts for other robots and the global environment via configuration space-time, explicit models of other agents, or other techniques (cf. [FS94, Rud94]; also see Section 2.5). However, researchers considering real-world multi-robot systems typically conclude that planning paths in advance is impossible. Thus, robots are often restricted to prescribed paths or roads, with rules (much like traffic laws in the human world) and communications used to avoid collision and deadlock [CCL+90, AHE+91].

Grossman [Gro88] classifies instances of the traffic control problem into three types: (i) restricted roads, (ii) multiple possible roads with robots selecting autonomously between them, and (iii) multiple possible roads with centralized traffic control. When individual robots possess unique roads from one point to another, no conflict is possible; when there is global knowledge and centralized control, it is easy to prevent conflict. Thus, the interesting case is (ii), where robots are allowed to autonomously select roads[21] (cf. "modest cooperation" [PY90], where robots are assumed to be benevolent for the common good of the system). Solutions to the traffic control problem range from rule-based solutions to approaches with antecedents in distributed processing.[22]

2.3 The Origin of Cooperation

In almost all of the work in collective robotics so far, it has been assumed that cooperation is explicitly designed into the system. An interesting research problem is to study how cooperation can arise without explicit human motivation among possibly selfish agents.

McFarland [McF94] distinguishes between two significantly different types of group behaviors that are found in

[19] A solution to the problem of electing these master cells was discussed in [UFA+93b]. Formation of structured cellular modules from a population of initially separated cells was studied in [UF93].

[20] [FS94] studies the problem of modeling the behavior of other cells, while [KIF93a, KIF93b] present a control method that calculates the goal of a cell based on its previous goal and on its master's goal. [FKA90] gives a means of estimating the amount of information exchanged between cells, and [UFA+93a] gives a heuristic for finding master cells for a binary communication tree.

[21] Analysis in [Gro88] shows that restricted roads are highly suboptimal, and that the autonomous road choice coupled with a greedy policy for escaping blocked situations is far more effective.

[22] In [KNT92], robots follow preplanned paths and use rules for collision avoidance. Example rules include "keep-right", "stop at intersection", "keep sufficient space to the robot in front of you", etc. [AOI+91] solves collision avoidance using two simple rules and a communication protocol that resolves conflict by transmitting individual priorities based on the task requirement, the environment, and robot performance. In [YP92], the robots stop at an intersection and indicate both the number of robots at the intersection and the directions in which they are traveling. If deadlock is possible, each robot performs "shunting" (trying to obtain high priority) and proceeds according to the agreed-upon priorities. [Wan91, WB90] adapt solutions (mutual exclusion and deadlock detection) from distributed processing to solve the traffic control problem.

nature: *eusocial behavior* and *cooperative behavior*. Eusocial behavior is found in many insect species (e.g., colonies of ants or bees), and is the result of genetically determined individual behavior. In eusocial societies, individual agents are not very capable, but seemingly intelligent behavior arises out of their interactions. This "cooperative" behavior is necessary for the survival of the individuals in the colonies.[23]

On the other hand, [McF94] defines cooperative behavior as the social behavior observed in higher animals (vertebrates); cooperation is the result of interactions between selfish agents. Unlike eusocial behavior, cooperative behavior is not motivated by innate behavior, but by an intentional desire to cooperate in order to maximize individual utility.

Inspired by economics and game-theoretic approaches, [GGR86, RG85, BG88, RZ94] and others have studied the emergence of cooperation in selfish rational agents in the field of distributed artificial intelligence (DAI).

2.4 Learning

Finding the correct values for control parameters that lead to a desired cooperative behavior can be a difficult, time-consuming task for a human designer. Therefore, it is highly desirable for multiple-robot systems to be able to *learn* control parameter values in order to optimize their task performance, and to adapt to changes in the environment. Reinforcement learning [BSW83, Kae93] has often been used in cooperative robotics.[24] In addition, techniques inspired by biological evolution have also been used in cooperative robotics.[25]

2.5 Geometric Problems

Because mobile robots can move about in the physical world and must interact with each other physically, geometric problems are inherent to multiple-robot systems.[26] Geometric problems that have been studied in the cooperative robotics literature include multiple-robot path planning, moving to (and maintaining) formation, and pattern generation.

(Multiple-Robot) Path Planning

Recall that multiple-robot path planning requires agents to plan routes that do not intersect. This is a case of resource conflict, since the agents and their goals are embedded in a finite amount of space. However, we note path planning separately because of its intrinsic geometric flavor as well as its historical importance in the literature.

Detailed reviews of path planning are found in [Fuj91, Lat91, AO92]. Fujimura [Fuj91] views path planning as either centralized (with a universal path-planner making decisions) or distributed (with individual agents planning and adjusting their paths). Arai and Ota [AO92] make a similar distinction in the nature of the planner, and also allow hybrid systems that can be combinations of on-line, off-line, centralized, or decentralized. Latombe [Lat91] gives a somewhat different taxonomy: his "centralized planning" is planning that takes into account all robots, while "decoupled" planning entails planning the path of each robot independently. For centralized planning, several methods used for single-robot systems can be applied. For decoupled planning, two approaches are given: (i) *prioritized planning* considers one robot at a time according to a global priority, while (ii) the *path coordination method* essentially plans paths by scheduling the configuration space-time resource.[27]

The Formation and Marching Problems

The Formation and Marching problems respectively require multiple robots to form up and move in a specified pattern. Solving these problems is quite interesting in terms of distributed algorithms [SS90], balancing between global and local knowledge [Par94],[28] and intrinsic information requirements for a given task. Solutions to Formation and Marching are also useful primitives for larger tasks, e.g., moving a large object by a group of robots [SB93][29] or distributed sensing [WB88].

The Formation problem seems very difficult, e.g., no published work has yet given a distributed "circle-forming" algorithm that guarantees the robots will actually end up in a circle.[30] We observe that the circle-forming problem, while quite simple, reveals several pitfalls in formulating distributed geometric tasks. For example, the ability of an individual agent to sense attributes of the formation must be carefully considered: too much information makes the problem trivial, but too little information (e.g., returns from localized sensors) may prevent a solution (e.g., robots may never find each other). Information lower bounds (e.g., for robots to be able to realize that they have achieved the prescribed formation) are also largely unexplored in the literature.

Related to the Formation problem is the pattern generation problem in Cellular Robotic Systems. A Cellular Robotic System (CRS) is a multiple-robot system which can "encode information as patterns of its own structural units" [Ben88]. Typically, one- or two-dimensional grids constitute the workspace, and sensing of neighboring cells is the only input. Within these constraints, a set of rules is devised and applied to all agents; a standard result is to show in simulation that convergence to some spatial pattern is guaranteed. The meaningful aspect of this work lies in providing a system with the capability of spatial self-organization: without in-

---

[23] [WD92] studies the evolution of herding behavior in "prey" agents in a simulated ecology, where there is no *a priori* drive for cooperation. Recently, [McF94, Ste94] have laid the initial groundwork to address the problem of emergent cooperation in an ecological system inhabited by actual mobile robots. In this ecosystem, individual robots are selfish, utility-driven agents that must cooperate in order to survive (i.e., maintain some minimal energy level).

[24] [Mat94a, Mat94b] propose a reformulation of the reinforcement learning paradigm using higher levels of abstraction (conditions, behaviors, and heterogeneous reward functions and progress estimators instead of states, actions, and reinforcement) to enable robots to learn a composite foraging behavior. [Par94] uses standard reinforcement algorithms to improve the performance of cooperating agents in the L-ALLIANCE architecture by having the agents learn how to better estimate the performance of other agents. [SSH94] uses reinforcement learning in a two-robot box-pushing system, and [YS92] applies reinforcement learning to learn a simple, artificial robot language. Other relevant works in multiagent reinforcement learning (done in simulation, as opposed to the works above which were implemented on actual robots) include [Whi91, Tan93, Lit94].

[25] [WD92] uses a genetic algorithm [Gol89] to evolve neural network controllers for simulated "prey" creatures that learn a herding behavior to help avoid predators. [Rey92] uses genetic programming [Koz90] to evolve flocking behavior in simulated "boids."

[26] This distinguishes robots from traditional distributed computer systems in which individual nodes are stationary.

[27] The work of [ELP86] is a typical decoupled approach where every robot is prioritized and robots plan global paths with respect to only higher-priority robots. On the other hand, [YB87] presents a distributed approach.

[28] [Par93] studies the problem of keeping four marching robots in a side-by-side formation; this increases in difficulty when the leader has to perform obstacle avoidance or other maneuvers. Parker also defines the concepts of global goals and global/local knowledge. To study the effects of different distributions of global goals and global knowledge, four strategies are compared both in simulation and on mobile robots.

[29] [CL94a] uses positional constraint conditions in a group of robots that makes turns while maintaining an array pattern. In [CL94b] a leader-follower approach is used to solve a similar task.

[30] For the circle-forming problem, the best known solution is the distributed algorithm of [SS90], which guarantees only that the robots will end up in a shape of constant diameter (e.g., a Reuleaux triangle can be the result). [CL94a, CL94b] extend the method of [SS90] to incorporate collision avoidance when the robots are moving. [YA94] approaches the shape-generation problem using systems of linear equations .

tervention, a CRS will reconfigure itself in certain situations or under certain conditions.[31]

3 Perspectives

As an integrative engineering discipline, robotics has always had to confront technological constraints that limit the domains that can be studied. Cooperative robotics has been subject to these same constraints, but the constraints tend to be more severe because of the need to cope with multiple robots. At the same time, cooperative robotics is a highly interdisciplinary field that offers the opportunity to draw influences from many other domains. We first outline some of the technological constraints that face the field. We then mention some directions in which *cooperative robotics* might progress, and describe related fields that have provided and will continue to provide influences.

3.1 Technological Constraints

It is clear that technological constraints have limited the scope of implementations and task domains attempted in multiple-robot research systems.

One obvious problem that arises is the general problem of researchers having to solve various instances of the vision problem before being able to make progress in "higher-level" problems. Often, difficulties arising from having to solve difficult perceptual problems can limit the range of tasks that can be implemented on a multiple-robot platform.[32] In addition, robot hardware is also notoriously unreliable; as a result, it is extremely difficult to maintain a fleet of robots in working condition. Again, collective robotics must deal with all of the hardware problems of single-agent robotics systems, exacerbated by the multiplicity of agents.

Due to the difficulties encountered when working with real robots (such as those outlined above), much of collective robotics has been studied exclusively in simulation. Some researchers have argued (cf. [Bro91]) that by ignoring most of the difficulties associated with perception and actuation, simulations ignore the most difficult problems of robotics. By making overly simplistic assumptions, it is possible to generate "successful" systems in simulation that would be infeasible in the real world.[33] Nevertheless, simulation must inevitably play a role in multi-agent robotics at some level. Although it is currently possible for researchers to study groups of 10-20 robots, it is unlikely that truly large-scale collective behavior involving hundreds or thousands of real robots will be feasible at any time in the near future.

An approach taken by some researchers is to use simulations as prototypes for larger-scale studies, and small numbers of real robots as a proof-of-concept demonstration [Mat94a, Par94]. On the other hand, some researchers, citing the necessity of working in the real world domain, have chosen to eschew simulations altogether and implement their theories directly on actual robots [BHD94, McF94, Ste94].[34]

3.2 Towards a Science of Cooperative Robotics

The field of cooperative mobile robotics offers an incredibly rich application domain, integrating a huge number of distinct fields from the social sciences, life sciences, and engineering. That so many theories have been brought to bear on "cooperative robotics" clearly shows the energy and the allure of the field. Yet, cooperative robotics is still an emerging field, and many open directions remain. In this section, we point out some promising directions that have yet to be fully explored by the research community. By way of a preface, we also point out three "cultural" changes which may come as the field matures: (1) Because of the youth of the field, cooperative robotics research has been necessarily rather informal and "concept" oriented. However, the development of rigorous formalisms is desirable to clarify various assumptions about the systems being discussed, and to obtain a more precise language for discussion of elusive concepts such as cooperation.[35] (2) Formal metrics for cooperation, system performance, as well as grades of cooperation are noticeably missing from the literature. While the notion of cooperation is difficult to formalize, such metrics will be very useful in characterizing various systems, and would improve our understanding of the nature of agent interactions.[36] (3) Experimental studies might become more rigorous and thorough, e.g., via standard benchmark problems and algorithms. This is challenging in mobile robotics, given the noisy, system-specific nature of the field. Nevertheless, it is necessary for claims about "robustness" and "near-optimality" to be appropriately quantified, and for dependencies on various control parameters to be better understood.[37]

Finally, several basic analogies remain incomplete, and must be revisited and resynthesized as the field matures. For instance, many multi-robot problems are "canonical" for distributed computation and are interesting primarily when viewed in this light.[38] More generally, it is likely that more structural and less superficial analogies with other disciplines will be needed in order to obtain "principled" theories of cooperation among (mobile) robots; integration of formalisms and methodologies developed in these more mature disciplines is likely to be an important step in the development of cooperative robotics. Disciplines most critical to the growth of cooperative robotics are: distributed artificial intelligence, biology, and distributed systems.

Distributed Artificial Intelligence

The field of distributed artificial intelligence (DAI) concerns itself with the study of distributed systems of intelligent agents. As such, this field is highly relevant to cooperative robotics. Bond and Gasser [BG88] define DAI as "the subfield of artificial intelligence (AI) concerned with concurrency in AI computations, at many levels." Grounded

[31] In [WB88], a CRS is characterized as an arbitrary number of robots in a one- or two-dimensional grid. The robots are able to sense neighboring cells and communicate with other robots via a signboard mechanism. Protocols are presented for creating different patterns (see also [EZ88]). An analogous cellular approach is adopted in [GDMO92].

[32] For example, in cooperative robotics systems where modeling of other agents (see 2.1.4) is used, the lack of an effective sensor array can render the system unimplementable in practice.

[33] Conversely, mobile research robots can also come to "look like the simulator", i.e., circular footprint, sonar ring, synchro-drive is a common configuration.

[34] An alternate approach adopted in [HB94], a study of locomotion in large herds of (up to 100) one-legged robots, is to design a very physically realistic simulation. While this approach brings realism to actuation, the issue of perception is still simulated away; it is still unclear whether it will be feasible to realistically model sophisticated agents in more complex environments, or whether the effort will outweigh the benefits.

[35] [Par94], which presents a formalization of motivational behavior in the ALLIANCE architecture, is a notable exception.

[36] [Mat94a] has suggested parameters such as agent density for estimating interference in a multi-robot system. However, much more work in this area is necessary.

[37] For example, we note that despite a number of claims that various decentralized approaches are superior to centralized approaches, we have not seen any thorough, published experimental comparisons between the major competing paradigms on a particular task.

[38] A typical example is moving to formation, which has been solved optimally in the computational geometry literature (it is the "geometric matching under isometry" problem [PH92]), but which is difficult in the distributed context due to issues like synchronization, fault-tolerance, leader election, etc. However, the distributed context can be selectively ignored, e.g., [SS90] use "human intervention" to perform what is essentially leader election (breaking symmetry in a circle of robots to choose vertices of the desired polygonal formation). The introduction of such devices runs counter to the implicit assumption that it is the *distributed* problem that holds research interest.

in traditional symbolic AI and the social sciences, DAI is composed of two major areas of study: Distributed Problem Solving (DPS) and Multiagent Systems (MAS).

Research in DPS is concerned with the problem of solving a single problem using many agents. Agents can cooperate by independently solving subproblems (task-sharing), and by periodically communicating partial solutions to each other (result-sharing). DPS involves three possibly overlapping phases: (i) problem decomposition (task allocation), (ii) subproblem solution, and (iii) solution synthesis.[39] The ACTOR formalism [Hew77] is a significant computational model, developed in DAI, which maps naturally to parallel computation.[40] One important assumption in DPS is that the agents are predisposed to cooperate. Research in DPS is thus concerned with developing frameworks for cooperative behavior between willing agents, rather than developing frameworks to enforce cooperation between potentially incompatible agents, as is the case with multiagent systems and distributed processing.

Multiagent Systems (MAS) research is the study of the collective behavior of a group of possibly heterogeneous agents with potentially conflicting goals. In other words, researchers in MAS discard the "benevolent agent" assumption of DPS [GGR86]. Genesereth et al. [GGR86] state the central problem of MAS research as follows: "in a world in which we get to design only our own intelligent agent, how should it interact with other intelligent agents?" Therefore, areas of interest in MAS research include game-theoretic analysis of multi-agent interactions (cf. [GGR86, RG85, RZ94]), reasoning about other agents' goals, beliefs, and actions (cf. [Geo83, Geo84, Ros82]), and analysis of the complexity of social interactions [ST92].

The influence of DAI on cooperative robotics has been limited. This is in part because researchers in DAI have mostly concentrated on domains where uncertainty is not as much of an issue as it is in the physical world. Work in MAS has tended to be theoretical and in very abstract domains where perfect sensing is usually assumed; typical DPS domains are in disembodied, knowledge-based systems. Another is that although agents may be selfish, they are rational and highly deliberative. However, achieving strict criteria of rationality and deliberativeness can often be prohibitively expensive in current robotic systems. Thus, it has been argued that DAI, while suited for unsituated, knowledge-based systems, will not succeed in the domain of cooperative robotics [Par94, Mat94a].[41]

Biology

Biological analogies and influences abound in the field of cooperative robotics. The majority of existing work in the field has cited biological systems as inspiration or justification. Well-known collective behaviors of ants, bees, and other eusocial insects [Wil71] provide striking existence proofs that systems composed of simple agents can accomplish sophisticated tasks in the real world. It is widely held that the cognitive capabilities of these insects are very limited, and that complex behaviors *emerge* out of interactions between the agents, which are individually obeying simple rules. Thus, rather than following the AI tradition of modeling robots as rational, deliberative agents, some researchers in cooperative robotics have chosen to take a more "bottom-up" approach in which individual agents are more like ants – they follow simple rules, and are highly reactive.ootnoteThis is the approach taken in the field of Artificial Life. Works based on this insect-colony analogy include [Mat94a, BHD94, SB93, DA93, JB94]. The pattern generation of CRS's can also be considered as bottom-up (see Section 2.5), since each robot is designed as a very simple agent which follows a set of prespecified rules.

A more general, biological[42] metaphor that is often used in cooperative robotics is the concept of a *self-organizing system* [NP77, Yat87].[43] Representative work that is based on this concept includes [WB88, Ste90, HB91, HB92, BH92]. Self-organization in multi-cellular biological systems has been an inspiration for [Ben88, EZ88, HW88, GDMO92]. Hierarchical organization of biological multi-cellular organisms (i.e., from cellular to tissue to organism level) has been used as a guiding metaphor for cellular robotics in the CEBOT project [FN87]. Biological analogies have also influenced the choice of task domains studied in cooperative robotics (note the large body of work on foraging and flocking/herding tasks). Finally, as we noted in Section 2.4, there have been some biological influences on the learning and optimization algorithms used to tune control parameters in multiple-robot systems.

Distributed Systems

A multiple-robot system is in fact a special case of a distributed computing system. Thus, the field of distributed systems is a natural source of ideas and solutions. [Ben88] describes cellular robotics as belonging to the general field of distributed computing. It is noted, however, that distributed computing can only contribute general theoretical foundations and that further progress needs to be made concerning the application of such methods to collective robotics. [WB90] states, "a distributed computing system contains a collection of computing devices which may reside in graphically separated locations called sites." By noting the similarities with distributed computing, theories pertaining to deadlock [WB88, WB90], message passing [WB90] and resource allocation [Wan91] have been applied to collective robotics in a number of works.[44] See also the discussion in Section 2.1.1 and Section 2.5.

Broadcast communication, which is widely assumed in cooperative robotics, exhibits poor scaling properties. As robots become more numerous and widely distributed, techniques and issues from the field of networking become relevant. A rich body of research on algorithms, protocols, performance modeling and analysis in computer networks can be applied to cooperative robotics.[45]

Finally, distributed control is a promising framework for the coordination of multiple robots. Due to difficulty of sens-

[39] Perhaps the best known scheme for task allocation problem is the Contract Net Protocol [Smi80], which has been used in the ACTRESS [IAT+94, AOI+94, OAI+93] and GOFER [CCL+90] projects.

[40] This work was the basis for the ACTRESS system [AMI89].

[41] It must be noted that direct comparisons of DAI and alternative paradigms are notably missing from the literature; such comparisons are needed to evaluate the true utility of DAI techniques in cooperative robotics. Also, as lower-level processes (perception and actuation) are better understood and implemented, and as computational power increases, the high-level results of DAI research may become increasingly applicable to collective mobile robotics.

[42] Researchers from many fields have studied self-organization; it is by no means an exclusively biological concept. However, in cooperative robotics references to self-organization have often been made in a biological context.

[43] Note that the behavior of insect colonies described above can be characterized more generally as that of self-organizing systems.

[44] In work done on multiple AGV systems, deadlock detection and resource allocation methods are applied to allow many robots to share the limited resource of path space [Wan91]. Pattern generation in a CRS may also rely on distributed computing to resolve conflicts [Wan91, WB90]. Finally, [Wan93, Wan94] describe a task allocation algorithm where the robots vie for the right to participate in a task.

[45] There is currently a great amount of effort being put into studying networking issues related to mobile/nomadic/ubiquitous computing (cf.[AP91, BAI94, Wei93]).

ing and communication, a parsimonious formulation which can coordinate robots having minimal sensing and communication capabilities is desirable. In an ideal scenario, maximal fault tolerance is possible, modeling of other agents is unnecessary, and each agent is controlled by a very simple mechanism.[46]

4 Conclusions

In this paper we synthesized a view of the theoretical bases for research in cooperative mobile robotics. Key research axes in the field were identified, particularly with respect to achieving a "mechanism of cooperation", and existing works were surveyed in this framework. We then discussed technological constraints and interdisciplinary influences that have shaped the field, and offered some general precepts for future growth of the field. Finally, we identified distributed artificial intelligence, biology, and distributed systems as disciplines that are most relevant to cooperative robotics, and which are most likely to continue to provide valuable influences. Based on our synthesis, a number of open research areas become apparent. We believe that the following are among the major, yet tractable, challenges for the near future: (1) robust definitions and metrics for various forms of cooperation, (2) achieving a more complete theory of information requirements for task-solving in spatial domains, perhaps for the canonical tasks of pattern formation or distributed sensing (e.g., measures of pattern complexity, information lower bounds for pattern recognition and maintenance, abstraction of sensor models from the solution approach), (3) principled transfer of the concepts of fault-tolerance and reliability from the field of distributed and fault-tolerant computing, (4) incorporation of recent ideas in distributed control to achieve oblivious cooperation or cooperation without communication (e.g., when robots have minimal sensing and communication capabilities), (5) achieving cooperation within competitive situations (e.g., for robot soccer, or pursuit-evasion with multiple pursuers and evaders), and (6) deepening ethological analogies, e.g., integrating known data regarding information structures (cognitive maps) of animals that together solve various cooperative tasks.

5 Acknowledgements

We thank B. Donald, T. Fukuda, M. Mataric, J. Wang, M. Anthony Lewis, and members of the UCLA Commotion Lab for helpful comments, suggestions, and discussions. We also thank L. E. Parker for her kind invitation to submit this paper.

REFERENCES

[AA94] R. Arkin and K. Ali. Integration of reactive and telerobotic control in multi-agent robotic systems. In *Proc. Simulation of Adaptive Behavior*, 1994.

[AH93] R. Arkin and J. Hobbs. Dimensions of communication and social organization in multi-agent robotic systems. In *Proc. Simulation of Adaptive Behavior*, 1993.

[AHE+91] H. Asama, M. K. Habib, I. Endo, K. Ozaki, A. Matsumoto, and Y. Ishida. Functional distribution among multiple mobile robots in an autonomous and decentralized robot system. In *IEEE ICRA*, 1921–6, 1991.

[AMI89] H. Asama, A. Matsumoto, and Y. Ishida. Design of an autonomous and distributed robot system: ACTRESS. In *IEEE/RSJ IROS*, 283–290, 1989.

[AO92] T. Arai and J. Ota. Motion planning of multiple robots. In *IEEE/RSJ IROS*, 1761–1768, 1992.

[AOI+91] 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 *IEEE/RSJ IROS*, 1215–1220, 1991.

[46] A distributed control scheme (known as the GUR game) developed originally by [Tse64] and recently studied in [TK93, Tun94] provides a framework in which groups of agents with minimal sensing capability and no communication can be controlled by simple finite state automata and converge to optimal behaviors. [Tun94] describes possible cooperative robotics applications in moving platform control and perimeter guarding.

[AOI+94] H. Asama, K. Ozaki, Y. Ishida, K. Yokita, A. Matsumoto, H. Kaetsu, and I. Endo. Collaborative team organization using communication in a decentralized robotic system. In *IEEE/RSJ IROS*, 1994.

[Ark92] R. C. Arkin. Cooperation without communication: Multiagent schema-based robot navigation. *Journal of Robotic Systems*, 9(3):351–364, 1992.

[AUN+94] M. Asada, E. Uchibe, S. Noda, S. Tawaratsumida, and K. Hosoda. Coordination of multiple behaviors acquired by a vision-based reinforcement learning. In *IEEE/RSJ IROS*, 1994.

[AYO93] T. Arai, E. Yoshida, and J. Ota. Information diffusion by local communication of multiple mobile robots. In *IEEE Conf. on Systems, Man and Cybernetics*, 535–540, 1993.

[AP91] B. Awerbuch and D. Peleg. Concurrent online tracking of mobile users. *Computer Communication Review*, 21(4):221–233, 1991.

[BAI94] B.R. Badrinath, A. Acharya, and T. Imielinski. Structuring distributed algorithms for mobile hosts. In *Proceedings of the 14th International Conference on Distributed Computing Systems*, pages 21–24, June 1994.

[Ben88] G. Beni. The concept of cellular robotic system. In *IEEE Int. Symp. on Intelligent Control*, 57–62, 1988.

[BG88] A. H. Bond and L. Gasser. *Readings in Distributed Artificial Intelligence*. Morgan Kaufmann Publishers, 1988.

[BG91] D. Barnes and J. Gray. Behaviour synthesis for co-operant mobile robot control. In *Int. Conf. on Control*, 1135–1140, 1991.

[BH92] G. Beni and S. Hackwood. Stationary waves in cyclic swarms. In *IEEE Int. Symp. on Intelligent Control*, 234–242, 1992.

[BHD94] R. Beckers, O. E. Holland, and J. L. Deneubourg. From local actions to global tasks: Stigmergy and collective robotics. In *Proc. A-Life IV*. MIT Press, 1994.

[Bro86] R. A. Brooks. A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, RA-2(1):14–23, 1986.

[Bro91] R. A. Brooks. Intelligence without reason. In *Proc. Intl. Joint Conf. Artificial Intelligence*, 569–595, 1991.

[BSW83] A.G. Barto, R.S. Sutton, and C.J.C.H. Watkins. Learning and sequential decision making. In M. Gabriel and J. Moore, editors, *Learning and Computational Neuroscience: Foundations of Adaptive Networks*, 539–603. MIT Press, 1983.

[BW91] G. Beni and J. Wang. Theoretical problems for the realization of distributed robotic systems. In *IEEE ICRA*, 1914–1920, 1991.

[CCL+90] P. Caloud, W. Choi, J.-C. Latombe, C. Le Pape, and M. Yin. Indoor automation with many mobile robots. In *IEEE/RSJ IROS*, 67–72, 1990.

[CL94a] Q. Chen and J. Y. S. Luh. Coordination and control of a group of small mobile robots. In *IEEE ICRA*, 2315–2320, 1994.

[CL94b] Q. Chen and J. Y. S. Luh. Distributed motion coordination of multiple robots. In *IEEE/RSJ IROS*, 1493–1500, 1994.

[Con87] J. Connell. Creature design with the subsumption architecture. In *Proc. AAAI*, 1124–1126, 1987.

[DA93] K. L. Doty and R. E. Van Aken. Swarm robot materials handling paradigm for a manufacturing workcell. In *IEEE ICRA*, 778–782, 1993.

[DF93] A. Drgoul and J. Ferber. From tom thumb to the dockers: Some experiments with foraging robots. In *Proc. Simulation of Adaptive Behavior*, 1993.

[DJR94] B. R. Donald, J. Jennings, and D. Rus. Analyzing teams of cooperating mobile robots. In *IEEE ICRA*, 1896–1903, 1994.

[Don93a] B. R. Donald. Information invariants in robotics: I. state, communication, and side-effects. In *IEEE ICRA*, 276–283, 1993.

[Don93b] B. R. Donald. Information invariants in robotics: II. sensors and computation. In *IEEE ICRA*, vol. 3, 284–90, 1993.

[Dor90] R. Dorf. *Concise Int. Encyclopedia of Robotics: Applications and Automation*. Wiley-Interscience, 1990.

[Dre92] K.E. Drexler. *Nanosystems: Molecular Machinery, Manufacturing, and Computation*. John Wiley and Sons, Inc., 1992.

[ELP86] M. Erdmann and T. Lozano-Perez. On multiple moving objects. In *IEEE ICRA*, 1419–1424, 1986.

[EZ88] O. Egecioglu and B. Zimmermann. The one dimensional random pairing problem in a cellular robotic system. In *IEEE Int. Symp. on Intelligent Control*, 76–80, 1988.

[FK93] T. Fukuda and Y. Kawauchi. *Cellular Robotics*, 745–782. Springer-Verlag, 1993.

[FKA90] T. Fukuda, Y. Kawauchi, and H. Asama. Analysis and evaluation of cellular robotics (CEBOT) as a distributed intelligent system by communication amount. In *IEEE/RSJ IROS*, 827–834, 1990.

[FN87] T. Fukuda and S. Nakagawa. A dynamically reconfigurable robotic system (concept of a system and optimal configurations). In *Int. Conf. on Industrial Electronics, Control, and Instrumentation*, 588–95, 1987.

[FS94] T. Fukuda and K. Sekiyama. Communication reduction with risk estimate for multiple robotic system. In *IEEE ICRA*, 2864–2869, 1994.

[Fuj91] K. Fujimura. *Motion Planning in Dynamic Environments*. Springer-Verlag, New York, NY, 1991.

[Gag93] D. Gage. How to communicate to zillions of robots. In *Mobile Robots VIII, SPIE*, 250–257, 1993.

[GD92] S. Goss and J. Deneubourg. Harvesting by a group of robots. In *Proc. European Conf. on Artificial Life*, 1992.

[GDMO92] V. Genovese, P. Dario, R. Magni, and L. Odetti. Self organizing behavior and swarm inteligence in a pack of mobile miniature robots in search of pollutants. In *IEEE/RSJ IROS*, 1575–1582, 1992.

[Geo83] M. Georgeff. Communication and interaction in multi-agent planning. In *Proc. AAAI*, 125–129, 1983.

[Geo84] M. Georgeff. A theory of action for multi-agent planning. In *Proc. AAAI*, 121–125, 1984.

[GGR86] M. R. Genesereth, M. L. Ginsberg, and J. S. Rosenschein. Cooperation without communication. In *Proc. AAAI*, 51–57, 1986.

[Gol89] D. Goldberg. *Genetic Algorithms in search, optimization, and machine learning*. Addison Wesley, 1989.

[Gro88] D. Grossman. Traffic control of multiple robot vehicles. *IEEE Journal of Robotics and Automation*, 4:491–497, 1988.

[HB91] S. Hackwood and G. Beni. Self-organizing sensors by deterministic annealing. In *IEEE/RSJ IROS*, 1177–1183, 1991.

[HB92] S. Hackwood and G. Beni. Self-organization of sensors for swarm intelligence. In *IEEE ICRA*, 819–829, 1992.

[HB94] J. Hodgins and D. Brogan. Robot herds: Group behaviors for systems with significant dynamics. In *Proc. A-Life IV*, 1994.

[Hew77] C. Hewitt. Viewing control structures as patterns of passing messages. *Artificial Intelligence*, 8(3):323–364, 1977.

[Hew93] C. Hewitt. Toward an open systems architecture. In *Information Processing 89. Proceedings of the IFIP 11th World Computer Congress*, 389–92, 1993.

[HO93] M. Hashimoto and F. Oba. Dynamic control approach for motion coordination of multiple wheeled mobile robots transporting a single object. In *IEEE/RSJ IROS*, 1944–1951, 1993.

[HW88] S. Hackwood and J. Wang. The engineering of cellular robotic systems. In *IEEE Int. Symp. on Intelligent Control*, 70–75, 1988.

[IAT+94] Y. Ishida, H. Asama, S. Tomita, K. Ozaki, A. Matsumoto, and I. Endo. Functional complement by cooperation of multiple autonomous robots. In *IEEE ICRA*, 2476–2481, 1994.

[IHH93] S. Ichikawa, F. Hara, and H. Hosokai. Cooperative route-searching behavior of multi-robot system using hello-call communiction. In *IEEE/RSJ IROS*, 1149–1156, 1993.

[JB94] P. J. Johnson and J. S. Bay. Distributed control of autonomous mobile robot collectives in payload transportation. Technical report, Virginia Polytechnic Institute and State Univ., Bradley Dept. of Elec. Engr., 1994.

[JLB94] K. Jin, P. Liang, and G. Beni. Stability of synchronized distributed control of discrete swarm structures. In *IEEE ICRA*, 1033–1038, 1994.

[Kae93] L. P. Kaelbling. *Learning in Embedded Systems*. MIT Press, 1993.

[KIF93a] Y. Kawauchi, M. Inaba, and T. Fukuda. A principle of distributed decision making of cellular robotic system (CEBOT). In *IEEE ICRA*, vol. 3, 833–838, 1993.

[KIF93b] Y. Kawauchi, M. Inaba, and T. Fukuda. A relation between resource amount and system performance of the cellular robotic system. In *IEEE/RSJ IROS*, 454–459, 1993.

[Kit94] H. Kitano. personal communication, 1994.

[KNT92] S. Kato, S. Nishiyama, and J. Takeno. Coordinating mobile robots by applying traffic rules. In *IEEE/RSJ IROS*, 1535–1541, 1992.

[Koz90] J. Koza. *Genetic Programming: On the Programming of Computers By the Means of Natural Selection*. MIT Press, 1990.

[KZ92] C. R. Kube and H. Zhang. Collective robotic intelligence. In *Proc. Simulation of Adaptive Behavior*, 460–468, 1992.

[Lat91] J. Latombe. *Robot Motion Planning*. Kluwer Academic, Boston, MA, 1991.

[LB92] M.A. Lewis and G.A. Bekey. The behavioral self-organization of nanorobots using local rules. In *IEEE/RSJ IROS*, 1333–1338, 1992.

[LeP90] C. LePape. A combination of centralized and distributed methods for multi-agent planning and scheduling. In *IEEE ICRA*, 488–493, 1990.

[Lit94] M. Littman. Markov games as a framework for multi-agent reinforcement learning. In *Proceedings of the Int. Machine Learning Conf.*, 157–163, 1994.

[Mat92] M. J. Mataric. Distributed approaches to behavior control. In *SPIE - Sensor Fusion V*, vol. 1828, 373–382, 1992.

[Mat94a] M. Mataric. *Interaction and Intelligent Behavior*. MIT AI Lab Technical Report, AI-TR-1495, August 1994.

[Mat94b] M. Mataric. Reward functions for accelerated learning. In *Proceedings of the Int. Machine Learning Conf.*, 181–189, 1994.

[MC94] G. F. Miller and D. Cliff. Protean behavior in dynamic games: Arguments for the co-evolution of pursuit-evasion tactics. In D. Cliff, P. Husbands, J.-A. Meyer, and S. W. Wilson, editors, *Proc. Simulation of Adaptive Behavior*, 1994.

[McF94] D. McFarland. Towards robot cooperation. In *Proc. Simulation of Adaptive Behavior*, 1994.

[MGT88] M. Mehregany, K.J. Gabriel, and W.S. Trimmer. Integrated fabrication of polysilicon mechanisms. *IEEE Trans. Electron Devices*, 35(6):719–723, 1988.

[MHB94] S. Ma, S. Hackwood, and G. Beni. Multi-agent supporting systems (MASS): Control with centralized estimator of disturbance. In *IEEE/RSJ IROS*, 679–686, 1994.

[MW63] Merriam-Webster. *Webster's 7th Collegiate Dictionary*. 1963.

[NP77] G. Nicolis and I. Prigogine. *Self-Organization in Nonequilibrium Systems*. Wiley-Interscience, 1977.

[NR91] F. Noreils and A. Recherche. Adding a man/machine interface to an architecture for mobile robots. In *IEEE/RSJ IROS*, 1991.

[OAI+93] K. Ozaki, H. Asama, Y. Ishida, A. Matsumoto, K. Yokota, H. Kaetsu, and I. Endo. Synchronized motion by multiple mobile robots using communication. In *IEEE/RSJ IROS*, 1164–1169, July 1993.

[Par92] L. E. Parker. Adaptive action selection for cooperative agent teams. In *Proc. Simulation of Adaptive Behavior*, 442–450. MIT Press, December 1992.

[Par93] L. E. Parker. Designing control laws for cooperative agent teams. In *IEEE ICRA*, 582–587, 1993.

[Par94] L. Parker. *Heterogeneous Multi-Robot Cooperation*. PhD thesis, MIT EECS Dept., February 1994.

[PH92] S. Schirra P.J. Heffernan. Approximate decision algorithms for point set congruence. In *8th Annual Compuational Geometry*, 93–101, 1992.

[PY90] S. Premvuti and S. Yuta. Consideration on the cooperation of multiple autonomous mobile robots. In *IEEE/RSJ IROS*, 59–63, 1990.

[Rey87] C. W. Reynolds. Flocks, herds and schools: a distributed behavioural model. *Computer Graphics*, 21(4):71–87, 1987.

[Rey92] C. Reynolds. An evolved, vision-based behavioral model of coordinated group motion. In *Proc. Simulation of Adaptive Behavior*, 1992.

[Rey94] C. Reynolds. Competition, coevolution and the game of tag. In *Proc. A-Life IV*, 1994.

[RG85] J.S. Rosenschein and M.R. Genesereth. Deals among rational agents. In *Proc. Intl. Joint Conf. Artificial Intelligence*, 91–99, 1985.

[Ros82] J.S. Rosenschein. Synchronization of multi-agent plans. In *Proc. AAAI*, 115–119, 1982.

[Rud94] M. Rude. Cooperation of mobile robots by event transforms into local space-time. In *IEEE/RSJ IROS*, 1501–1507, 1994.

[RZ94] J.S. Rosenschein and G. Zlotkin. *Rules of Encounter:designing conventions for automated negotiation among computers*. MIT Press, 1994.

[SB93] D. J. Stilwell and J. S. Bay. Toward the development of a material transport system using swarms of ant-like robots. In *IEEE ICRA*, 766–771, 1993.

[Smi80] R. Smith. The contract net protocol: high-level communication and control in a distributed problem solver. *IEEE Trans. Computers*, 1104–1113, 1980.

[SS90] K. Sugihara and I. Suzuki. Distributed motion coordination of multiple mobile robots. In *Proc. IEEE Int. Symp. on Intelligent Control*, 1990.

[SSH94] S. Sen, M. Sekaran, and J. Hale. Learning to coordinate without sharing information. In *Proc. AAAI*, 426–431, 1994.

[ST92] Y. Shoham and M. Tennenholtz. On the synthesis of useful social laws for artificial agent societies (preliminary report). In *Proc. AAAI*, 276–281, 1992.

[Ste90] L. Steels. Cooperation between distributed agents through self-organization. In *European Workshop on Modelling Autonomous Agents in a Multi-Agent World*, 175–195, 1990.

[Ste94] L. Steels. A case study in the behavior-oriented design of autonomous agents. In *Proc. Simulation of Adaptive Behavior*, 1994.

[Tan93] M. Tan. Multi-agent reinforcement learning: independent vs. cooperative agents. In *Proceedings of the Int. Machine Learning Conf.*, 1993.

[TK93] B. Tung and L. Kleinrock. Distributed control methods. In *Proceedings of the 2nd Int. Symp. on High Performance Distr ibuted Computing*, 206–215, 1993.

[Tse64] M.L. Tsetlin. *Finite Automata and Modeling the Simplest Forms of Behavior*. PhD thesis, V.A. Steklov Mathematical Institute, 1964.

[Tun94] Y-C. Tung. *Distributed Control Using Finite State Automata*. PhD thesis, UCLA Computer Science Dept., 1994.

[UF93] T. Ueyama and T. Fukuda. Self-organization of cellular robots using random walk with simple rules. In *IEEE ICRA*, 595–600, 1993.

[UFA+93a] T. Ueyama, T. Fukuda, F. Arai, Y. Kawauchi, Y. Katou, S. Matsumura, and T. Uesugi. Communication architecture for cellular robotic system. *JSME Int. Journal, Series C*, 36:353–360, 1993.

[UFA+93b] T. Ueyama, T. Fukuda, F. Arai, T. Sugiura, A. Sakai, and T. Uesugi. Distributed sensing, control and planning - cellular robotics approach. In *IMACS*, 433–438. Elsevier Science Publ. (North-Holland), 1993.

[Van91] K. VanLehn, editor. *Architectures for Intelligence: The 22nd Carnegie Mellon Symp. on Cognition*. Lawrence Erlbaum Associates, 1991.

[Wan91] J. Wang. Fully distributed traffic control strategies for many-AGV systems. In *IEEE/RSJ IROS*, 1199–1204, 1991.

[Wan93] J. Wang. DRS operating primitives based on distributed mutual exclusion. In *IEEE/RSJ IROS*, 1085–1090, 1993.

[Wan94] J. Wang. On sign-board based inter-robot communication in distributed robotic systems. In *IEEE ICRA*, 1045–1050, 1994.

[WB88] J. Wang and G. Beni. Pattern generation in cellular robotic systems. In *IEEE Int. Symp. on Intelligent Control*, 63–69, 1988.

[WB90] J. Wang and G. Beni. Distributed computing problems in cellular robotic systems. In *IEEE/RSJ IROS*, 819–826, 1990.

[Wei93] M. Weiser. Some computer science issues in ubiquitous computing. *Communications of the ACM*, 36(7):74–84, 1993.

[WD92] G. Werner and M. Dyer. Evolution of herding behavior in artificial animals. In *Proc. Simulation of Adaptive Behavior*, 1992.

[Whi91] S. Whitehead. A complexity analysis of cooperative mechanisms in reinforcement learning. In *Proc. AAAI*, 607–613, 1991.

[Wil71] E.O. Wilson. *The insect societies*. Harvard Univ. Press, 1971.

[WNM94] Z.-D. Wang, E. Nakano, and T. Matsukawa. Cooperating multiple behavior-based robots for object manipulation. In *IEEE/RSJ IROS*, 1524–1531, 1994.

[YA94] H. Yamaguchi and T. Arai. Distributed and autonomous control method for generating shape of multiple mobile robot group. In *IEEE/RSJ IROS*, 800–807, 1994.

[Yat87] F.E. Yates, editor. *Self-Organizing Systems: The Emergence of Order*. Plenum Press, 1987.

[YB87] D. Yeung and G. Bekey. A decentralized approach to the motion planning problem for multiple mobile robots. In *IEEE ICRA*, 1779–1784, 1987.

[YP92] S. Yuta and S. Premvuti. Coordinating autonomous and centralized decision making to achieve cooperative behaviors between multiple mobile robots. In *Proc. of the 1992 IEEE/RSJ Int. Conf. on Intelligent Robots and Systems, Raleigh, NC, July 7-10, 1992*, 1566–1574, July 1992.

[YS92] H. Yanco and L. Stein. An adaptive communication protocol for cooperating mobile robots. In *Proc. Simulation of Adaptive Behavior*, 478–485, 1992.

[YSA+94] K. Yokota, T. Suzuki, H. Asama, A. Masumoto, and I. Endo. A human interface system for the multi-agent robotic system. In *IEEE ICRA*, 1039–1044, 1994.