Autonomous Robots, 4, 1–23 (1997)

Cooperative Mobile Robotics: Antecedents and Directions*

Y. UNY CAO
yu@cs.ucla.edu
Computer Science Department, University of California, Los Angeles, CA 90024-1596

ALEX S. FUKUNAGA
alex.fukunaga@jpl.nasa.gov
Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109-8099

ANDREW B. KAHNG
abk@cs.ucla.edu
Computer Science Department, University of California, Los Angeles, CA 90024-1596

Editors: R.C. Arkin and G.A. Bekey

Abstract. There has been increased research interest in systems composed of multiple autonomous mobile robots exhibiting cooperative 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 cooperative 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.

Keywords: cooperative robotics, swarm intelligence, distributed robotics, artificial intelligence, mobile robots, multiagent systems

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:

- tasks may be inherently too complex (or impossible) for a single robot to accomplish, or performance benefits can be gained from using multiple robots;
- 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
- the constructive, synthetic approach inherent in cooperative mobile robotics can possibly yield insights into fundamental problems in the social sciences (organization theory, economics, cognitive psychology), and life sciences (theoretical biology, animal ethology).

The study of multiple-robot systems 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

* This is an expanded version of a paper which originally appeared in the proceedings of the 1995 IEEE/RSJ IROS conference.

the subject of the present survey, is a subclass of collective behavior that is characterized by cooperation. Webster's dictionary [118] 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:

1. "joint collaborative behavior that is directed toward some goal in which there is a common interest or reward" [22];
2. "a form of interaction, usually based on communication" [108]; and
3. "[joining] together for doing something that creates a progressive result such as increasing performance or saving time" [137].

These definitions show the wide range of possible motivating perspectives. For example, definitions such as (1) 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 (2) 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. Finally, definition (3) 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.

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. In particular, we concentrated on fundamental *theoretical* issues that impinge on cooperative robotics. Thus, the following related subjects were outside the scope of this work:

- coordination of multiple manipulators, articulated arms, or multi-fingered hands, etc.
- *human-robot cooperative systems*, and user-interface issues that arise with multiple-robot systems [184] [8] [124] [1].
- the *competitive* subclass of collective behavior, which includes pursuit-evasion [139], [120] and one-on-one competitive games [12]. Note that a cooperative team strategy for, e.g., work on the robot soccer league recently started in Japan [87] would lie within our present scope.
- emerging technologies such as nanotechnology [48] and Micro Electro-Mechanical Systems [117] that are likely to be very important to cooperative robotics are beyond the scope of this paper.

Even with these restrictions, we find that over the past 8 years (1987-1995) alone, 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.

We are aware of two previous works that have surveyed or taxonomized the literature. [13] is a broad, relatively succinct survey whose scope encompasses *distributed autonomous robotic systems* (i.e., not restricted to mobile robots). [50] focuses on several well-known "swarm" architectures (e.g., SWARM and Mataric's Behavior-based architecture – see Section 2.1) and proposes a taxonomy to characterize these architectures. The scope and intent of our work differs significantly from these, in that (1) we extensively survey the field of cooperative mobile robotics, and (2) we provide a taxonomical organization of the literature based on *problems* and *solutions* that have arisen in the field (as opposed to a selected group of architectures). In addition, we survey much new material that has appeared since these earlier works were published.

Towards 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 [46]. Coordination and interactions of multiple intelligent agents have been actively studied in the field of distributed artificial intelligence (DAI) since the early 1970's [28], 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 CEBOT [59], SWARM [25], ACTRESS [16], GOFER [35], and the work at Brussels [151]. These early projects were done primarily in simulation, and, while the early work on CEBOT, ACTRESS and GOFER have all had physical implementations (with ≤ 3 robots), in some sense these implementations were presented by way of proving the simulation results. Thus, several more recent works (cf. [91], [111], [131]) are significant for establishing an emphasis on the actual physical implementation of cooperative robotic systems. Many of the recent cooperative robotic systems, in contrast to the earlier works, are based on a behavior-based approach (cf. [30]). Various perspectives on autonomy and on the connection between intelligence and environment are strongly associated with the behavior-based approach [31], but are not intrinsic to multiple-robot systems and thus lie beyond our present scope. Also note that a recent incarnation of CEBOT, which has been implemented on physical robots, is based on a behavior-based control architecture [34].

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. For example, traffic rules are commonly used to reduce planning cost for avoiding collision and deadlock in a real-world environment, such as a network of roads. (Interestingly, behavior-based approaches identify collision avoidance as one of the most basic behaviors [30], and achieving a collision-avoidance behavior is the natural solution to collision avoidance among multiple robots. However, in reported experiments that use the behavior-based approach, robots are never restricted to road networks.)
- **Box-Pushing/Cooperative Manipulation.** Many works have addressed the box-pushing (or couch-pushing) problem, for widely varying reasons. The focus in [134] is on task allocation, fault-tolerance and (reinforcement) *learning*. By contrast, [45] studies two box-pushing protocols in terms of their intrinsic communication and hardware requirements, via the concept of information invariants. Cooperative manipulation of large objects is particularly interesting in that cooperation can be achieved without the robots even knowing of each others' existence [147], [159]. Other works in the class of box-pushing/object manipulation include [175] [153] [82] [33] [91] [94] [92] [114] [145] [72] [146].
- **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 [32] [151] [10] [67] [102] [49] [108] [9] [24]. The task is interesting because (1) it can be performed by each robot independently (i.e., the issue is whether multiple robots achieve a performance gain), and (2) as discussed in Section 3.2, the task is also interesting due to motivations related to the biological inspirations behind cooperative robot systems. There are some conceptual overlaps with the related task of materials handling in a manufacturing workcell [47]. A wide variety of techniques have been applied, ranging from simple stigmergy (essentially random movements that result in the fortuitous collection of objects [24] to more complex algorithms in which robots form chains along which objects are passed to the goal [49]. [24] 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. On the other hand, that stigmergy may not satisfy our definition of cooperation given above, since there is no performance improvement over the "naive algorithm" – in this particular case, the proposed stigmergic algorithm *is* the naive algorithm. Again, group architecture and learning are major research themes in addressing this problem.

Other interesting task domains that have received attention in the literature include multi-robot security systems [53], landmine detection and clearance [54], robotic structural support systems (i.e., keeping structures stable in case of, say, an earthquake) [107], map making [149], and assembly of objects using multiple robots [175].

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.

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?* 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. (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 ([39], [45] and many others) or flocking [140], [108]. 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).

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 motivated and achieved. Here, 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. 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. 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 exceptions of [126], [106], [134]), (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). Note that there is also a related, geometric problem of optimizing the allocation of tasks spatially. This has been recently studied in the context of the division of the search of a work area by multiple robots [97]. 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" [165]. 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 the ability to model other agents. We then describe several representative systems that have addressed these specific problems.

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.

The behavior of decentralized systems is often described using such terms as "emergence" and "self-organization." It is widely claimed that decentralized architectures (e.g., [24], [10], [152], [108]) have several inherent advantages over centralized architectures, including fault tolerance, natural exploitation of parallelism, reliability, and scalability. However, we are not aware of any published empirical or theoretical comparison that supports these claims directly. Such a comparison would be interesting, particularly in scenarios where the team of robots is relatively small (e.g., two robots pushing a box), and it is not clear whether the scaling properties of decentralization offset the coordinative advantage of centralized systems.

In practice, many systems do not conform to a strict centralized/decentralized dichotomy, e.g., many largely decentralized architectures utilize "leader" agents. We are not aware of any instances of systems that are completely centralized, although there are some hybrid centralized/decentralized architectures wherein there is a central planner that exerts high-level control over mostly autonomous agents [126], [106], [3], [36].

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. [134] has introduced the concept of *task coverage*, which measures the ability of a given team member to achieve a given task. This parameter is an index of the demand for cooperation: when task coverage is high, tasks can be accomplished without much cooperation, but otherwise, cooperation is necessary. Task coverage is maximal in homogeneous groups, and decreases as groups become more heterogeneous (i.e., in the limit only one agent in the group can perform any given task).

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 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 are either known at design-time, or arise dynamically at run-time.

Communication Structures The communication structure of a group determines the possible modes of inter-agent interaction. We characterize three major types of interactions that can be supported. ([50] proposes a more detailed taxonomy of communication structures).

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. This modality has also been called "cooperation without communication" by some researchers. Systems that depend on this form of interaction include [67], [24], [10], [151], [159], [160], [147].

Interaction via sensing

Corresponding to arms-length relationships in organization theory [75], interaction via sensing refers to local interactions that occur between agents as a result of agents sensing one another, but without explicit communication. This type of interaction requires the ability of agents to distinguish between other agents in the group and other objects in the environment, which is called "kin recognition" in some literatures [108]. Interaction via sensing is indispensable for modeling of other agents (see Section 2.1.4 below). Because of hardware limitations, interaction via sensing has often been emulated using radio or infrared communications. However, several recent works attempt to implement true interaction via sensing, based on vision [95], [96], [154]. Collective behaviors that can use this kind of interaction include flocking and pattern formation (keeping in formation with nearest neighbors).

Interaction via communications

The third form of interaction involves explicit communication with other agents, by either directed or broadcast intentional messages (i.e. the recipient(s) of the message may be either known or unknown). Because architectures that enable this form of communication are similar to communication networks, many standard issues from the field of networks arise, including the design of network topologies and communications protocols. For example, in [168] a media access protocol (similar to that of Ethernet) is used for inter-robot communication. In [78], robots with limited communication range communicate to each other using the "hello-call" protocol, by which they establish "chains" in order to extend their effective communication ranges. [61] describes methods for communicating to many ("zillions") robots, including a variety of schemes ranging from broadcast channels (where a message is sent to all other robots in the system) to modulated retroreflection (where a master sends out a laser signal to slaves and interprets the response by the nature of the reflection). [174] describes and simulates a wireless CSMA/CD (Carrier Sense Multiple Access with Collision Detection) protocol for the distributed robotic systems.

There are also communication mechanisms designed specially for multiple-robot systems. For example, [171] proposes the "sign-board" as a communication mechanism for distributed robotic systems. [7] gives a communication protocol modeled after diffusion, wherein local communication similar to chemical communication mechanisms in animals is used. The communication is engineered to decay away at a preset rate. Similar communications mechanisms are studied in [102], [49], [67]. Additional work on communication can be found in [185], which analyzes optimal group sizes for local communications and communication delays. In a related vein, [186], [187] analyzes optimal local communication ranges in broadcast communication.

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.

The second of two box-pushing protocols in [45] 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. [153] [73]). In a two-robot bar carrying task, Fukuda and Sekiyama's agents [60] each uses a probabilistic model of the other agent. When a risk threshold is exceeded, an agent communicates with its partner to maintain coordination. In [43], [44], 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 [147], where the robots achieve box pushing but are not aware of each other at all. For a more complex task involving the placement of five desks in [154], a homogeneous group of four robots share a ceiling camera to get positional information, but do not communicate with each other. Each robot relies on modeling of other agents to detect conflicts of paths and placements of desks, and to change plans accordingly.

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.

CEBOT

CEBOT (CEllular roBOTics System) is a decentralized, hierarchical architecture inspired by the cellular organization of biological entities (cf. [59] [57], [162] [161] [56]). 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. A solution to the problem of electing these master cells was discussed in [164]. Formation of structured cellular modules from a population of initially separated cells was studied in [162]. 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). [60] studies the problem of modeling the behavior of other cells, while [85], [86] present a control method that calculates the goal of a cell based on its previous goal and on its master's goal. [58] gives a means of estimating the amount of information exchanged between cells, and [163] gives a heuristic for finding master cells for a binary communication tree. A new behavior selection mechanism is introduced in [34], based on two matrices, the priority matrix and the interest relation matrix, with a learning algorithm used to adjust the priority matrix. Recently, a Micro Autonomous Robotic System (MARS) has been built consisting of robots of 20 cubic mm and equipped with infrared communications [121].

ACTRESS

The ACTRESS (ACTor-based Robot and Equipments Synthetic System) project [16], [80], [15] is inspired by the Universal Modular ACTOR Formalism [76]. In the ACTRESS system, "robotors", including 3 robots and 3 workstations (one as interface to human operator, one as image processor and one as global environment manager), form a heterogeneous group trying to perform tasks such as object pushing [14] that cannot be accomplished by any of the individual robotors alone [79], [156]. Communication protocols at different abstraction levels [115] provide a means upon which "group cast" and negotiation mechanisms based on Contract Net [150] and multistage negotiation protocols are built [18]. Various issues are studied, such as efficient communications between robots and environment managers [17], collision avoidance [19].

SWARM

A *SWARM* is a distributed system with a large number of autonomous robots [81]. (Note that the work on SWARM systems began as work on Cellular Robotic Systems [25], where many simple agents occupied one- or two-dimensional environments and were able to perform tasks such as pattern generation and self-organization). *SWARM intelligence* is "a property of systems of non-intelligent robots exhibiting collectively intelligent behavior" [69]. *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 (an exception is [70], where two different types of robots were used). Interaction takes place by each cell reacting to the state of its nearest neighbors. Mechanisms for self-organization in SWARM are studied in [70] [27] [26] [69] [103] [81]. Examples for possible applications include large-scale displays and distributed sensing [71]. Communication primitives have been an important part of research in SWARM [170], [173] (see Section 3.2 below for more details).

GOFER

The GOFER architecture [35], [100] 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 [150] 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.

Other architectures that make significant use of concepts studied within the classical distributed paradigm are described in [106] [123] [126] [127] [125] [4] [5].

ALLIANCE/L-ALLIANCE

The ALLIANCE architecture was developed by Parker [134], [133] 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 [134] is an extension to ALLIANCE 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 [112], [110], [113], [108] 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.

Similar behavior-based architectures include the work by Kube et al, which is based on an Adaptive Logic Network, a neural network [91], [92], [94], [93], the Tropism-Based Cognitive Architecture [2], and an architecture based on "instinctive behaviors" [40].

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 (i.e., sharing of communication media is usually achieved using very basic techniques – wireless LAN, straightforward time-division multiplexing, or broadcast over an RF channel; recently, [172], [187], [186] have considered some problems in sharing communications channels). 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 For example, [60] proposes a "hierarchical prediction model" which essentially uses simulation to achieve collision avoidance. [144] considers the problem of crossing an intersection: event transforms into the local space-time coordinate frame of a robot are applied, and each robot (i) iteratively updates the local frame and its objects, (ii) evaluates collision risk, and (iii) generates a modified path depending on the collision risk. (See also 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 [35], [14].

Grossman [68] 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. Analysis in [68] 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 (cf. "modest cooperation" [137], 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. In [84], robots follow pre-planned paths and use rules for collision avoidance. Example rules include "keep-right", "stop at intersection", and "keep sufficient space to the robot in front of you". [19] 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 the robot performance. In [188], 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. [166] takes a distributed computing approach to traffic control, where the particular problem solved is to keep the number of robots traveling on any given path below a threshold value. Robots use a mutual exclusion protocol to compete for the right to travel on each path. Wang and Beni [171] adapt two distributed algorithms to solve two problems in their CRS/SWARM architecture. The "n-way intersection problem" is solved using an algorithm similar to mutual exclusion, and the "knot detection problem" is solved using an algorithm similar to distributed deadlock detection.

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 [116] distinguishes between two significantly different types of group behaviors that are found in 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. [177] studies the evolution of herding behavior in "prey" agents in a simulated ecology, where there is no *a priori* drive for cooperation. Recently, [116], [152] have laid the initial groundwork to address the problem of emergent cooperation in an ecological system inhabited by actual mobile robots. In their ecosystem, individual robots are selfish, utility-driven agents that must cooperate in order to survive (i.e., maintain some minimal energy level).

On the other hand, [116] defines cooperative behavior as the social behavior observed in higher animals (vertebrates); i.e., 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, [28] [62] [142] [143] and others have studied the emergence of cooperation in selfish rational agents in the field of distributed artificial intelligence (DAI). A recent work in the robotics literature that adopts this game-theoretic approach is [128].

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 [23], [83] has often been used in cooperative robotics.

Mataric [109], [108] proposes 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. [134] 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. [147] uses reinforcement learning in a two-robot box-pushing system, and [181] applies reinforcement learning to learn a simple, artificial robot language. Other relevant works in multiagent reinforcement learning (done in simulation, in contrast to the above works which were implemented on actual robots) include [178], [157], [105].

In addition, techniques inspired by biological evolution have also been used in cooperative robotics. [177] uses a genetic algorithm [66] to evolve neural network controllers for simulated "prey" creatures that learn a herding behavior to help avoid predators. [138] uses genetic programming [90] to evolve flocking behavior in simulated "boids."

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. This is a fundamental property that distinguishes multiple-robot systems from traditional distributed computer systems in which individual nodes are stationary. 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 [55], [99], [6]. Fujimura [55] 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 [6] 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 [99] 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 originally 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. The work of [52] is a typical decoupled approach where every robot is prioritized and robots plan global paths with respect to only higher-priority robots (e.g., the highest-priority robot plans only around the obstacles in the environment). Note that this is still a centralized method according to the terminology of [55], [6]. On the other hand, [183] presents a distributed approach (per Fujimura's taxonomy) where each robot initially attempts a straight-line path to the goal; if an interfering obstacle is seen, then the robot will scan the visible vertices of the obstacle and move toward the closest one. In general, this continues until the goal is reached. Dynamically varying priorities are given to each robot (based on current need) to resolve path intersection conflicts, and conflicting robots can either negotiate among themselves or allow a global blackboard manager to perform this function.

Some recent works have addressed some nontraditional motion planning problems. For example, [74] proposes an algorithm for path planning in tethered robots, and [129] consider the problem of moving while grasping large objects.

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 [155], balancing between global and local knowledge [134], 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 [153] [37] [38] or distributed sensing [170].

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. For this problem, the best known solution is the distributed algorithm of [155], which guarantees only that the robots will end up in a shape of constant diameter (e.g., a Reuleaux triangle can be the result). It is assumed that the i^{th} mobile robot knows the distances D_i and d_i to its farthest and nearest neighbors, respectively; the algorithm attempts to match the ratios D_i/d_i to a prescribed constant. No method of detecting termination was given. [37], [38] extend the method of [155] to incorporate collision avoidance when the robots are moving. [180] approaches the shape-generation problem using systems of linear equations; starting at some initial location, each robot changes its (x, y) position according to a linear function of its neighbors' positions and some fixed constant. Simulations of the method show that a group of initially collinear robots will converge into the shape of an arc.

We observe that the circle-forming problem, while quite simple to state, 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. Interestingly, we note that the algorithm of [155] can be slightly modified: rather than each robot seeking to achieve a prescribed ratio D/d, each robot could seek to achieve a prescribed angle (close to 90 degrees) subtended by its farthest neighbor and its closest neighbor to the right. This uses very similar sensing capabilities but guarantees the desired circular shape.

For Marching, [37] employs positional constraint conditions in a group of robots that makes turns while maintaining an array pattern. In [38] a leader-follower approach is used to solve a similar task. [132] 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. Simplified instances of the Marching problem require robots to reliably follow each other and to move in a group (without tight constraints on their relative positions). Some works that address this problem (sometimes referred to as the herding/flocking problem) include [140], [108], [77], [29], [119]. A somewhat related problem is the problem of cooperative positioning (determining the locations of the robots in a group using limited information) [98].

Related to the Formation problem is the pattern generation problem in Cellular Robotic Systems, multiple-robot systems which can "encode information as patterns of its own structural units" [25]. 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: a CRS will reconfigure itself without intervention in certain situations or under certain conditions.

In [170], [103], 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, e.g., alternating robots and spaces in a one-dimensional grid; covering the top row of a two-dimensional grid by robots; or covering the boundary of a two-dimensional grid by robots. Egecioglu and Zimmermann [51] pose the "Random Pairing" problem, and seek a set of rules by which for any given number, a CRS will converge to a pattern such that there is a group of two robots with that number of vacant spaces between them (see also [26]). An analogous cellular approach is adopted by Genovese et al. [63], who describe the simulation of a system of pollutant-seeking mobile robots. The simulation uses a potential field mechanism to attract robots to the pollutant and to repulse robots from each other. The combined effect of these two forces yields a gradient pattern that "points" toward the source of the pollutant.

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. In this section, 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 on "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. For example, in cooperative robotics systems where modeling of other agents (see Section 2.1) is used, the lack of an effective sensor array can render the system unimplementable in practice. 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-robotic systems, exacerbated by the multiplicity of agents.

Due to the difficulties (such as those outlined above) encountered when working with real robots, much of collective robotics has been studied exclusively in simulation. Some researchers have argued (cf. [31]) 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. (Conversely, mobile research robots can also come to "look like the simulator", i.e., circular footprint, sonar ring, synchro-drive is a common configuration.) 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. Thus, cooperative mobile robot researchers have used a variety of techniques to simulate perception, while using physical robots. For instance, the use of a global positioning system can in part compensate for the lack of vision, but can place severe environmental constraints under which robots can operate (because many objects and acoustic features of the environment can interfere with the GPS). For the basic problem of differentiating between other agents and all other objects in the environment. some researchers [134] use radio communication to solve this problem. In other works [43], [134] interaction via sensing is done by explicit radio communication. There are recent attempts to perform recognition via vision [95], [96].

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 [108], [134]. 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 [24] [116] [152]. In studies of locomotion in large herds of (upto 100) one-legged robots and simulated human cyclists, [77] [29] take an alternate approach of 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.

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 subsection, 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 (there are some exceptions, such as [134], which presents a formalization of motivational behavior in the ALLIANCE architecture). (2) Formal metrics for cooperation and system performance, as well as for 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. Although [108] has suggested parameters such as agent density for estimating interference in a multi-robot system, much more work in this area is necessary. (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. For example, we have noted 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. However, we note that recently, researchers have begun to empirically study quantitative measures of cooperation, trying to identify conditions under which mechanisms for cooperation are beneficial [11], [9], [135].

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. 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 [136]), 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., [155] 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.

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 [28] define DAI as "the subfield of artificial intelligence (AI) concerned with concurrency in AI computations, at many levels." Grounded 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 issue 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. Of these, problem decomposition has attracted the greatest interest among DAI researchers. The critical issue in task sharing is finding the appropriate agent to assign to a subproblem. This is nontrivial, since if the most appropriate agent for a subtask is not obvious, then the system must try to determine which of the many eligible agents should be assigned the task, and often there are too many eligible agents to attempt an exhaustive search. Perhaps the best known scheme for task allocation is the Contract Net Protocol [150], which has been used in the ACTRESS [79], [18], [130] and GOFER [35] projects.

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 [62]. [62] states 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. [62], [142], [143]), reasoning about other agents' goals, beliefs, and actions (cf. [141], [64], [65]), and analysis of the complexity of social interactions [148].

Work in MAS has tended to be theoretical and in very abstract domains. A common underlying assumption is that although the agents may be selfish, they are rational and highly deliberative. This is in stark contrast with research in swarm intelligence (see Section 2.5), in which individual agents are assumed to be relatively unintelligent.

However, 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 assumption of DAI that has prevented its application in cooperative robotics is the assumption 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 [134], [108]. However, we observe 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.

Distributed Systems

A multiple-robot system is in fact a special case of a distributed system. Thus, the field of distributed systems is a natural source of ideas and solutions. [25] 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. [171] states, "a distributed computing system contains a collection of computing devices which may reside in geographically separated locations called sites." By noting the similarities with distributed computing, theories pertaining to deadlock [170], [171], [104], message passing [171] and resource allocation [166], and the combination of the above as primitives [169], [173], have been applied to collective robotics in a number of works.

In work done on multiple AGV systems and Distributed Robotic Systems, deadlock detection and resource allocation methods are applied to allow many robots to share the limited resource of path space [171], [166]. Pattern generation in a CRS may also rely on distributed computing to resolve conflicts [166], [171]. Finally, [167], [168] describe a task allocation algorithm where the robots vie for the right to participate in a task. See also the discussion in Section 2.1 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 computer networks become relevant. A rich body of research on algorithms, protocols, performance modeling and analysis in computer networks can be applied to cooperative robotics. There is currently a great amount of effort being put into studying networking issues related to mobile/nomadic/ubiquitous computing (cf. [20], [176], [88], [21]). Results from this field could be applied in a straightforward way to multi-robot systems.

Finally, distributed control is a promising framework for the coordination of multiple robots. Due to difficulty of sensing 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. A distributed control scheme (known as the GUR game) developed originally by [158] and recently studied in [159], [160] provides a framework in which groups of agents with minimal sensing capability and no communication are controlled by simple finite state automata and converge to optimal behaviors. [160] describes possible cooperative robotics applications in moving platform control and perimeter guarding.

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 [179] 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 (this is the approach taken in the field of Artificial Life). Works based on this insect-colony analogy include [108], [24], [153], [47], [82], [41], [42]. 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 metaphor that is often used in cooperative robotics is the concept of a *self-organizing system* [122], [182]. (Note that researchers from many fields have studied self-organization; it is by no means an exclusively biological concept. However, in the field of cooperative robotics, references to self-organization have often been made in a biological context.) The behavior of insect colonies described above can be characterized more generally as that of self-organizing systems. Representative work that is based on this concept includes [170], [151], [69], [70], [26]. Self-organization in multi-cellular biological systems has been an inspiration for [71], [25], [51], [63]. 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 [59].

Biological analogies have also influenced the choice of task domains studied in cooperative robotics. While foraging is a natural abstraction of some practical applications such as waste retrieval and search and rescue, one major reason that it has become identified as the canonical cooperative robotic task is that it is a natural task, given the group architectures resulting from analogies to insect colonies. Another example of this phenomenon is the flocking/herding task. It seems no accident that biological inspirations led to “natural” models of group motion, as opposed to more structured models of coordinated motion (such as moving in some arbitrary formation).

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.

4. Conclusions

We have 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). The works of [45], [145], [33] have begun to address this issue, in the context of object manipulation tasks; interestingly, [33] observes that given a robot system, some tasks are *strongly cooperative* – the robots must act in concert to achieve the goal, and the strategy for the task is not trivially serializable.
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). An interesting open problem is how well solutions that have been developed in discretized abstractions of these domains (cf. [101], [89]) translate to the physical world.

Acknowledgements

Partial support for this work was provided by NSF Young Investigator Award MIP-9257982; the UCLA Commotion Laboratory is supported by NSF CDA-9303148. Portions of this work were performed by the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration. The authors would like to thank B. Donald, T. Fukuda, M. Anthony Lewis, M. Mataric, J. Wang, the anonymous reviewers, and members of the UCLA Commotion Lab for helpful comments, suggestions, and discussions. Frank Meng assisted in the preparation of a previous version of this paper.

References

1. J.A. Adams, R. Basjcsy, J. Kosecka, V. Kumar, R. Mandelbaum, M. Mintz, R. Paul, C. Wang, Y. Yamamoto, and X. Yun. Cooperative material handling by human and robotic agents: Module development and system synthesis. In *IEEE/RSJ IROS*, pages 200–205, 1995.
2. A. Agah and G. A. Bekey. Autonomous mobile robot teams. In *Conf. on Intelligent Robotics in Filed, Factory, Service and Space (CIRFFSS94)*, 1994.
3. L. Aguilar, R. Alami, S. Fleury, M. Herrb, F. Ingrand, and F. Robert. Ten autonomous mobile robots (and even more). In *IEEE/RSJ IROS*, pages 260–267, 1995.
4. R. Alani, F. Robert, F. Ingrand, and S. Suzuki. Multi-robot cooperation through incremental plan-merging. In *IEEE ICRA*, pages 2573–2579, 1995.
5. J. S. Albus. *A Control Architecture for Cooperative Intelligent Robots*, pages 713–743. 1993.
6. T. Arai and J. Ota. Motion planning of multiple robots. In *IEEE/RSJ IROS*, pages 1761–1768, 1992.
7. T. Arai, E. Yoshida, and J. Ota. Information diffusion by local communication of multiple mobile robots. In *IEEE Conference on Systems, Man and Cybernetics*, pages 535–540, 1993.
8. R. Arkin and K. Ali. Integration of reactive and telerobotic control in multi-agent robotic systems. In *Proc. Simulation of Adaptive Behavior*, 1994.
9. R. Arkin and J. Hobbs. Dimensions of communication and social organization in multi-agent robotic systems. In *Proc. Simulation of Adaptive Behavior*, 1993.
10. R. C. Arkin. Cooperation without communication: Multiagent schema-based robot navigation. *Journal of Robotic Systems*, 9(3):351–364, 1992.
11. R. C. Arkin, T. Balch, and E. Nitz. Communication of behavioral state in multi-agent retrieval tasks. In *IEEE ICRA*, volume 3, pages 588–594, 1993.

12. 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.
13. H. Asama. Distributed autonomous robotic system configurated with multiple agents and its cooperative behaviors. *Journal of Robotics and Mechatronics*, 4(3), 1992.
14. 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*, pages 1921–6, 1991.
15. H. Asama, Y. Ishida, K. Ozaki, M. K. Habib, A. Matsumoto, H. Kaetsu, and I. Endo. A communication system between multiple robotic agents. In M. Leu, editor, *Proc. the Japan U.S.A. Symposium on Flexible Automation*, pages 647–654, 1992.
16. H. Asama, A. Matsumoto, and Y. Ishida. Design of an autonomous and distributed robot system: ACTRESS. In *IEEE/RSJ IROS*, pages 283–290, 1989.
17. H. Asama, K. Ozaki, Y. Ishida, M. K. Habib, A. Matsumoto, and I. Endo. Negotiation between multiple mobile robots and an environment manager. In *IEEE ICRA*, pages 533–538, 1991.
18. 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.
19. 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*, pages 1215–1220, 1991.
20. B. Awerbuch and D. Peleg. Concurrent online tracking of mobile users. *Computer Communication Review*, 21(4):221–233, 1991.
21. 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.
22. D. Barnes and J. Gray. Behaviour synthesis for cooperant mobile robot control. In *International Conference on Control*, pages 1135–1140, 1991.
23. 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*, pages 539–603. MIT Press, 1983.
24. 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.
25. G. Beni. The concept of cellular robotic system. In *IEEE International Symposium on Intelligent Control*, pages 57–62, 1988.
26. G. Beni and S. Hackwood. Stationary waves in cyclic swarms. In *IEEE International Symposium on Intelligent Control*, pages 234–242, 1992.
27. G. Beni and J. Wang. Theoretical problems for the realization of distributed robotic systems. In *IEEE ICRA*, pages 1914–1920, 1991.
28. A. H. Bond and L. Gasser. *Readings in Distributed Artificial Intelligence*. Morgan Kaufmann Publishers, 1988.
29. D.C. Brogan and J.C. Hodgins. Group behaviors for systems with significant dynamics. In *IEEE/RSJ IROS*, pages 528–534, 1995.
30. R. A. Brooks. A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, RA-2(1):14–23, 1986.
31. R. A. Brooks. Intelligence without reason. In *Proc. Intl. Joint Conf. Artificial Intelligence*, pages 569–595, 1991.
32. R. A. Brooks, P. Maes, M. J. Mataric, and G. More. Lunar base construction robots. In *IEEE/RSJ IROS*. IEEE, July 1990.
33. R.G. Brown and J.S. Jennings. A pusher/steerer model for strongly cooperative mobile robot cooperation. In *IEEE/RSJ IROS*, pages 562–568, 1995.
34. A.-H. Cai, T. Fukuda, F. Arai, T. Ueyama, and A. Sakai. Hierarchical control architecture for cellular robotic system. In *IEEE ICRA*, pages 1191–1196, 1995.
35. P. Caloud, W. Choi, J.-C. Latombe, C. Le Pape, and M. Yin. Indoor automation with many mobile robots. In *IEEE/RSJ IROS*, pages 67–72, 1990.
36. O. Causse and L.H. Pampagnin. Management of a multi-robot system in a public environment. In *IEEE/RSJ IROS*, pages 246–252, 1995.
37. Q. Chen and J. Y. S. Luh. Coordination and control of a group of small mobile robots. In *IEEE ICRA*, pages 2315–2320, 1994.
38. Q. Chen and J. Y. S. Luh. Distributed motion coordination of multiple robots. In *IEEE/RSJ IROS*, pages 1493–1500, 1994.
39. J. Connell. Creature design with the subsumption architecture. In *Proc. AAAI*, pages 1124–1126, 1987.
40. P. Dario, F. Ribechini, V. Genovese, and G. Sandini. Instinctive behaviors and personalities in societies of cellular robots. In *IEEE ICRA*, pages 1927–1932, 1991.
41. J. Deneubourg, S. Goss, N. Franks, A. Sendova-Franks, C. Detrain, and L. Chretien. The dynamics of collective sorting: Robot-like ants and ant-like robots. In *Proc. Simulation of Adaptive Behavior*, 1991.
42. J-L. Deneubourg, G. Theraulaz, and R. Beckers. Swarm-made architectures. In Verela and Bourgine, editors, *Proc. European Conference on Artificial Life*, pages 123–133. MIT Press, 1991.
43. B. R. Donald. Information invariants in robotics: I. state, communication, and side-effects. In *IEEE ICRA*, pages 276–283, 1993.
44. B. R. Donald. Information invariants in robotics: II. sensors and computation. In *IEEE ICRA*, volume 3, pages 284–90, 1993.
45. B. R. Donald, J. Jennings, and D. Rus. Analyzing teams of cooperating mobile robots. In *IEEE ICRA*, pages 1896–1903, 1994.
46. R. Dorf. *Concise International Encyclopedia of Robotics: Applications and Automation*. Wiley-Interscience, 1990.
47. K. L. Doty and R. E. Van Aken. Swarm robot materials handling paradigm for a manufacturing workcell. In *IEEE ICRA*, volume 1, pages 778–782, 1993.

48. K.E. Drexler. *Nanosystems: Molecular Machinery, Manufacturing, and Computation.* John Wiley and Sons, Inc., 1992.
49. A. Drgoul and J. Ferber. From tom thumb to the dockers: Some experiments with foraging robots. In *Proc. Simulation of Adaptive Behavior*, 1993.
50. G. Dudek, M. Jenkin, E. Milios, and D. Wilkes. A taxonomy for swarm robots. In *IEEE/RSJ IROS*, pages 441–447, 1993.
51. O. Egecioglu and B. Zimmermann. The one dimensional random pairing problem in a cellular robotic system. In *IEEE International Symposium on Intelligent Control*, pages 76–80, 1988.
52. M. Erdmann and T. Lozano-Perez. On multiple moving objects. In *IEEE ICRA*, pages 1419–1424, 1986.
53. H. Everett, G.A. Gilbreath, T.A. Heath-Pastore, and R.T. Laird. Coordinated control of multiple security robots. *Mobile Robots VIII*, 2058:292–305, 1993.
54. D.E. Franklin, A.B. Kahng, and M.A. Lewis. Distributed sensing and probing with multiple search agents: toward system-level landmine detection solutions. In *Detection Technologies for Mines and Minelike Targets, Proceedings of SPIE, Vol.2496*, pages 698–709, 1995.
55. K. Fujimura. *Motion Planning in Dynamic Environments.* Springer-Verlag, New York, NY, 1991.
56. T. Fukuda and G. Iritani. Construction mechanism of group behavior with cooperation. In *IEEE/RSJ IROS*, pages 535–542, 1995.
57. T. Fukuda and Y. Kawauchi. *Cellular Robotics*, pages 745–782. Springer-Verlag, 1993.
58. 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*, pages 827–834, 1990.
59. T. Fukuda and S. Nakagawa. A dynamically reconfigurable robotic system (concept of a system and optimal configurations). In *International Conference on Industrial Electronics, Control, and Instrumentation*, pages 588–95, 1987.
60. T. Fukuda and K. Sekiyama. Communication reduction with risk estimate for multiple robotic system. In *IEEE ICRA*, pages 2864–2869, 1994.
61. D. Gage. How to communicate to zillions of robots. In *Mobile Robots VIII, SPIE*, pages 250–257, 1993.
62. M. R. Genesereth, M. L. Ginsberg, and J. S. Rosenschein. Cooperation without communication. In *Proc. AAAI*, pages 51–57, 1986.
63. 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*, pages 1575–1582, 1992.
64. M. Georgeff. Communication and interaction in multi-agent planning. In *Proc. AAAI*, pages 125–129, 1983.
65. M. Georgeff. A theory of action for multi-agent planning. In *Proc. AAAI*, pages 121–125, 1984.
66. D. Goldberg. *Genetic Algorithms in search, optimization, and machine learning.* Addison Wesley, 1989.
67. S. Goss and J. Deneubourg. Harvesting by a group of robots. In *Proc. European Conference on Artificial Life*, 1992.
68. D. Grossman. Traffic control of multiple robot vehicles. *IEEE Journal of Robotics and Automation*, 4:491–497, 1988.
69. S. Hackwood and G. Beni. Self-organizing sensors by deterministic annealing. In *IEEE/RSJ IROS*, pages 1177–1183, 1991.
70. S. Hackwood and G. Beni. Self-organization of sensors for swarm intelligence. In *IEEE ICRA*, pages 819–829, 1992.
71. S. Hackwood and J. Wang. The engineering of cellular robotic systems. In *IEEE International Symposium on Intelligent Control*, pages 70–75, 1988.
72. F. Hara, Y. Yasui, and T. Aritake. A kinematic analysis of locomotive cooperation for two mobile robots along a genereal wavy road. In *IEEE ICRA*, pages 1197–1204, 1995.
73. M. Hashimoto and F. Oba. Dynamic control approach for motion coordination of multiple wheeled mobile robots transporting a single object. In *IEEE/RSJ IROS*, pages 1944–1951, 1993.
74. S. Hert and V. Lumelsky. Moving multiple tethered robots between arbitrary configurations. In *IEEE/RSJ IROS*, pages 280–285, 1995.
75. C. Hewitt. Toward an open systems architecture. In *Information Processing 89. Proceedings of the IFIP 11th World Computer Congress*, pages 389–92, 1993.
76. C. Hewitt, P. Bishop, I. Greif, B. Smith, T. Matson, and R. Steiger. A universal modular actor formalism for artificial intelligence. In *Proc. Intl. Joint Conf. Artificial Intelligence*, pages 235–245, 1973.
77. J. Hodgins and D. Brogan. Robot herds: Group behaviors for systems with significant dynamics. In *Proc. A-Life IV*, 1994.
78. S. Ichikawa, F. Hara, and H. Hosokai. Cooperative route-searching behavior of multi-robot system using hello-call communiction. In *IEEE/RSJ IROS*, pages 1149–1156, 1993.
79. Y. Ishida, H. Asama, S. Tomita, K. Ozaki, A. Matsumoto, and I. Endo. Functional complement by cooperation of multiple autonomous robots. In *IEEE ICRA*, pages 2476–2481, 1994.
80. Y. Ishida, I. Endo, and A. Matsumoto. Communication and cooperation in an autonomous and decentralized robot system. In *IFAC int. Symp. on Distributed Intelligent Systems*, pages 299–304, 1991.
81. K. Jin, P. Liang, and G. Beni. Stability of synchronized distributed control of discrete swarm structures. In *IEEE ICRA*, pages 1033–1038, 1994.
82. P. J. Johnson and J. S. Bay. Distributed control of autonomous mobile robot collectives in payload transportation. Technical report, Virginia Polytechnic Institute and State University, Bradley Dept. of Elec. Engr., 1994.
83. L. P. Kaelbling. *Learning in Embedded Systems.* MIT Press, 1993.
84. S. Kato, S. Nishiyama, and J. Takeno. Coordinating mobile robots by applying traffic rules. In *IEEE/RSJ IROS*, pages 1535–1541, 1992.
85. Y. Kawauchi, M. Inaba, and T. Fukuda. A principle of distributed decision making of cellular robotic system (CEBOT). In *IEEE ICRA*, volume 3, pages 833–838, 1993.

86. Y. Kawauchi, M. Inaba, and T. Fukuda. A relation between resource amount and system performance of the cellular robotic system. In *IEEE/RSJ IROS*, pages 454–459, 1993.
87. H. Kitano. personal communication, 1994.
88. L. Kleinrock. Nomadic computing - an opportunity. *Computer Communications Review*, January 1995.
89. R. Korf. A simple solution to pursuit games. In *Proc. 11th International Workshop on Distributed Artificial Intelligence*, 1992.
90. J. Koza. *Genetic Programming: On the Programming of Computers By the Means of Natural Selection.* MIT Press, 1990.
91. C. R. Kube and H. Zhang. Collective robotic intelligence. In *Proc. Simulation of Adaptive Behavior*, pages 460–468, 1992.
92. C. R. Kube and H. Zhang. Collective robotics: From social insects to robots. *Adaptive Behavior*, 2(2):189–219, 1993.
93. C. R. Kube and H. Zhang. Stagnation recovery behaviours for collective robotics. In *IEEE/RSJ IROS*, pages 1883–1890, 1994.
94. C. R. Kube, H. Zhang, and X. Wang. Controlling collective tasks with an ALN. In *IEEE/RSJ IROS*, pages 289–293, 1993.
95. Y. Kuniyoshi, N. Kita, S. Rougeaux, S. Sakane, M. Ishii, and M. Kakikura. Cooperation by observation - the framework and basic task patterns. In *IEEE ICRA*, pages 767–774, 1994.
96. Y. Kuniyoshi, J. Riekki, M. Ishii, S. Rougeaux, N. Kita, S. Sakane, and M. Kakikura. Vision-based behaviors for multi-robot cooperation. In *IEEE/RSJ IROS*, pages 925–931, 1994.
97. D. Kurabayashi, J. Ota, T. Arai, and E. Yoshida. An algorithm of dividing a work area to multiple mobile robots. In *IEEE/RSJ IROS*, pages 286–291, 1995.
98. R. Kurazume and S. Nagata. Cooperative positioning with multiple robots. In *IEEE ICRA*, pages 1250–1257, 1994.
99. J. Latombe. *Robot Motion Planning.* Kluwer Academic, Boston, MA, 1991.
100. C. LePape. A combination of centralized and distributed methods for multi-agent planning and scheduling. In *IEEE ICRA*, pages 488–493, 1990.
101. R. Levy and J.S. Rosenschein. A game theoretic approach to distributed artificial intelligence and the pursuit problem. In *European Workshop on Modelling Autonomous Agents in a Multi-Agent World*, pages 129–146, 1991.
102. M.A. Lewis and G.A. Bekey. The behavioral self-organization of nanorobots using local rules. In *IEEE/RSJ IROS*, pages 1333–1338, 1992.
103. P. Liang and G. Beni. Robotic morphogenesis. In *IEEE ICRA*, pages 2175–2180, 1995.
104. F.-C. Lin and J. Y.-J. Hsu. Cooperation and deadlock-handling for an object-sorting task in a multi-agent robotic system. In *IEEE ICRA*, pages 2580–2585, 1995.
105. M. Littman. Markov games as a framework for multi-agent reinforcement learning. In *Proceedings of the International Machine Learning Conference*, pages 157–163, 1994.
106. T.C. Lueth and T. Laengle. Task description, decomposition and allocation in a distributed autonomous multi-agent robot system. In *IEEE/RSJ IROS*, pages 1516–1523, 1994.
107. S. Ma, S. Hackwood, and G. Beni. Multi-agent supporting systems (MASS): Control with centralized estimator of disturbance. In *IEEE/RSJ IROS*, pages 679–686, 1994.
108. M. Mataric. *Interaction and Intelligent Behavior.* PhD thesis, MIT, EECS, May 1994.
109. M. Mataric. Reward functions for accelerated learning. In *Proceedings of the International Machine Learning Conference*, pages 181–189, 1994.
110. M. J. Mataric. Designing emergent behaviors: From local interactions to collective intelligence. In J.-A. Meyer, H. Roitblat, and S. Wilson, editors, *From Animals to Animats 2, Second International Conference on Simulation of Adaptive Behavior (SAB-92)*, pages 432–441. MIT Press, 1992.
111. M. J. Mataric. Distributed approaches to behavior control. In *SPIE - Sensor Fusion V*, volume 1828, pages 373–382, 1992.
112. M. J. Mataric. Minimizing complexity in controlling a mobile robot population. In *IEEE ICRA*, pages 830–835, May 1992.
113. M. J. Mataric. Kin recognition, similarity, and group behavior. In *Fifteenth Annual Cognitive Science Society Conference*, pages 705–710. Lawrence Erlbaum Associates, June 1993.
114. M.J. Mataric, M. Nilsson, and K.T. Simsarian. Cooperative multi-robot box-pushing. In *IEEE/RSJ IROS*, pages 556–561, 1995.
115. A. Matsumoto, H. Asama, Y. Ishida, K. Ozaki, and I. Endo. Communication in the autonomous and decentralized robot system: ACTRESS. In *IEEE/RSJ IROS*, pages 835–840, 1990.
116. D. McFarland. Towards robot cooperation. In *Proc. Simulation of Adaptive Behavior*, 1994.
117. M. Mehregany, K.J. Gabriel, and W.S. Trimmer. Integrated fabrication of polysilicon mechanisms. *IEEE Trans. Electron Devices*, 35(6):719–723, 1988.
118. Merriam-Webster. *Webster's 7th Collegiate Dictionary.* Merriam-Webster, Inc., 1963.
119. G. Dudek M. Jenkin E. Milios and D. Wilkes. Experiments in sensing and communication for robot convoy navigation. In *IEEE/RSJ IROS*, pages 268–273, 1995.
120. 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.
121. N. Mitsumoto, T. Fukuda, K. Shimojina, and A. Ogawa. Micro autonomous robotic system and biologically inspired immune swarm strategy as a multi agent robotic system. In *IEEE ICRA*, pages 2187–2192, 1995.
122. G. Nicolis and I. Prigogine. *Self-Organization in Nonequilibrium Systems.* Wiley-Interscience, 1977.
123. F. Noreils. Integrating multirobot coordination in a mobile robot control system. In *IEEE/RSJ IROS*, pages 43–49, 1990.

124. F. Noreils and A. Recherche. Adding a man/machine interface to an architecture for mobile robots. In *IEEE/RSJ IROS*, 1991.
125. F. R. Noreils. Multi-robot coordination for battlefield strategies. In *IEEE/RSJ IROS*, pages 1777–1784, July 1992.
126. F. R. Noreils. Toward a robot architecture integrating cooperation between mobile robots: Application to indoor environment. *The International Journal of Robotics Research*, 12(1), February 1993.
127. F. R. Noreils and R. de Nozay. Coordinated protocols: An approach to formalize coordination between mobile robots. In *IEEE/RSJ IROS*, pages 717–724, July 1992.
128. C. Numaoka. Collective alteration of strategic types with delayed global information. In *IEEE/RSJ IROS*, pages 1077–1084, 1993.
129. J. Ota, N. Miyata, T. Arai, E. Yoshida, D. Kurabayashi, and J. Sasaki. Transferring and regrasping a large object by cooperation of multiple mobile robots. In *IEEE/RSJ IROS*, pages 543–548, 1995.
130. 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*, pages 1164–1169, July 1993.
131. L. E. Parker. Adaptive action selection for cooperative agent teams. In *Second Annual International Conference on Simulation of Adaptive Behavior*, pages 442–450. MIT Press, December 1992.
132. L. E. Parker. Designing control laws for cooperative agent teams. In *IEEE ICRA*, volume 3, pages 582–587, 1993.
133. L. E. Parker. ALLIANCE: an architecture for fault tolerant, cooperative control of heterogeneous mobile robots. In *IEEE/RSJ IROS*, pages 776–783, 1994.
134. L. E. Parker. *Heterogeneous Multi-Robot Cooperation*. PhD thesis, MIT EECS Dept., February 1994.
135. L. E. Parker. The effect of action recognition and robot awareness in cooperative robotic teams. In *IEEE/RSJ IROS*, pages 212–219, 1995.
136. S. Schirra P.J. Heffernan. Approximate decision algorithms for point set congruence. In *8th Annual Compuational Geometry*, pages 93–101, 1992.
137. S. Premvuti and S. Yuta. Consideration on the cooperation of multiple autonomous mobile robots. In *IEEE/RSJ IROS*, pages 59–63, 1990.
138. C. Reynolds. An evolved, vision-based behavioral model of coordinated group motion. In *Proc. Simulation of Adaptive Behavior*, 1992.
139. C. Reynolds. Competition, coevolution and the game of tag. In *Proc. A-Life IV*, 1994.
140. C. W. Reynolds. Flocks, herds and schools: a distributed behavioural model. *Computer Graphics*, 21(4):71–87, 1987.
141. J.S. Rosenschein. Synchronization of multi-agent plans. In *Proc. AAAI*, pages 115–119, 1982.
142. J.S. Rosenschein and M.R. Genesereth. Deals among rational agents. In *Proc. Intl. Joint Conf. Artificial Intelligence*, pages 91–99, 1985.
143. J.S. Rosenschein and G. Zlotkin. *Rules of Encounter:designing conventions for automated negotiation among computers*. MIT Press, 1994.
144. M. Rude. Cooperation of mobile robots by event transforms into local space-time. In *IEEE/RSJ IROS*, pages 1501–1507, 1994.
145. D. Rus, B. Donald, and J. Jennings. Moving furniture with teams of autonomous robots. In *IEEE/RSJ IROS*, pages 235–242, 1995.
146. J. Sasaki, J. Ota, E. Yoshida, D. Kurabayashi, and T. Arai. Cooperating grasping of a large object by multiple mobile robots. In *IEEE ICRA*, pages 1205–1210, 1995.
147. S. Sen, M. Sekaran, and J. Hale. Learning to coordinate without sharing information. In *Proc. AAAI*, pages 426–431, 1994.
148. Y. Shoham and M. Tennenholtz. On the synthesis of useful social laws for artificial agent societies (preliminary report). In *Proc. AAAI*, pages 276–281, 1992.
149. K. Singh and K. Fujimura. Map making by cooperating mobile robots. In *IEEE ICRA*, volume 2, pages 254–259, 1993.
150. R. Smith. The contract net protocol: high-level communication and control in a distributed problem solver. *IEEE Trans. Computers*, pages 1104–1113, 1980.
151. L. Steels. Cooperation between distributed agents through self-organization. In *European Workshop on Modelling Autonomous Agents in a Multi-Agent World*, pages 175–195, 1990.
152. L. Steels. A case study in the behavior-oriented design of autonomous agents. In *Proc. Simulation of Adaptive Behavior*, 1994.
153. D. J. Stilwell and J. S. Bay. Toward the development of a material transport system using swarms of ant-like robots. In *IEEE ICRA*, volume 6, pages 766–771, 1993.
154. H. Sugie, Y. Inagaki, S. Ono, H. Aisu, and T. Unemi. Placing objects with multiple mobile robots – mutual help using intention inference. In *IEEE ICRA*, pages 2181–2186, 1995.
155. K. Sugihara and I. Suzuki. Distributed motion coordination of multiple mobile robots. In *Proc. IEEE International Symposium on Intelligent Control*, 1990.
156. T. Suzuki, H. Asama, A. Uegaki, S. Kotosaka, T. Fujita, A. Matsumoto, H. Kaetsu, and I. Endo. An infra-red sensory system with local communication for cooperative multiple mobile robots. In *IEEE/RSJ IROS*, volume 1, pages 220–225, 1995.
157. M. Tan. Multi-agent reinforcement learning: independent vs. cooperative agents. In *Proceedings of the International Machine Learning Conference*, 1993.
158. M.L. Tsetlin. *Finite Automata and Modeling the Simplest Forms of Behavior*. PhD thesis, V.A. Steklov Mathematical Institute, 1964.
159. B. Tung and L. Kleinrock. Distributed control methods. In *Proceedings of the 2nd International Symposium on High Performance Distr ibuted Computing*, pages 206–215, 1993.
160. Y-C. Tung. *Distributed Control Using Finite State Automata*. PhD thesis, UCLA Computer Science Department, 1994.

161. T. Ueyama and T. Fukuda. Knowledge acquisition and distributed decision making. In *IEEE ICRA*, volume 3, pages 167–172, 1993.
162. T. Ueyama and T. Fukuda. Self-organization of cellular robots using random walk with simple rules. In *IEEE ICRA*, volume 3, pages 595–600, 1993.
163. T. Ueyama, T. Fukuda, F. Arai, Y. Kawauchi, Y. Katou, S. Matsumura, and T. Uesugi. Communication architecture for cellular robotic system. *JSME International Journal, Series C*, 36:353–360, 1993.
164. T. Ueyama, T. Fukuda, F. Arai, T. Sugiura, A. Sakai, and T. Uesugi. Distributed sensing, control and planning - cellular robotics approach. In *IMACS*, pages 433–438. Elsevier Science Publ. (North-Holland), 1993.
165. K. VanLehn, editor. *Architectures for Intelligence: The 22nd Carnegie Mellon Symposium on Cognition.* Lawrence Erlbaum Associates, 1991.
166. J. Wang. Fully distributed traffic control strategies for many-AGV systems. In *IEEE/RSJ IROS*, pages 1199–1204, 1991.
167. J. Wang. DRS operating primitives based on distributed mutual exclusion. In *IEEE/RSJ IROS*, pages 1085–1090, 1993.
168. J. Wang. On sign-board based inter-robot communication in distributed robotic systems. In *IEEE ICRA*, pages 1045–1050, 1994.
169. J. Wang. Operating primitives supporting traffic regulation and control of mobile robots under distributed robotic systems. In *IEEE ICRA*, pages 1613–1618, 1995.
170. J. Wang and G. Beni. Pattern generation in cellular robotic systems. In *IEEE International Symposium on Intelligent Control*, pages 63–69, 1988.
171. J. Wang and G. Beni. Distributed computing problems in cellular robotic systems. In *IEEE/RSJ IROS*, pages 819–826, 1990.
172. J. Wang and S. Premvuti. Resource sharing in distributed robotic systems based on a wireless medium access protocol (CSMA/CD-W). In *IEEE/RSJ IROS*, pages 784–791, 1994.
173. J. Wang and S. Premvuti. Distributed traffic regulation and control for multiple autonomous mobile robots operating in discrete space. In *IEEE ICRA*, pages 1619–1624, 1995.
174. J. Wang, S. Premvuti, and A. Tabbara. A wireless medium access protocol (csma/cd-w) for mobile robot based distributed robotic system. In *IEEE ICRA*, pages 2561–2566, 1995.
175. Z.-D. Wang, E. Nakano, and T. Matsukawa. Cooperating multiple behavior-based robots for object manipulation. In *IEEE/RSJ IROS*, pages 1524–1531, 1994.
176. M. Weiser. Some computer science issues in ubiquitous computing. *Communications of the ACM*, 36(7):74–84, 1993.
177. G. Werner and M. Dyer. Evolution of herding behavior in artificial animals. In *Proc. Simulation of Adaptive Behavior*, 1992.
178. S. Whitehead. A complexity analysis of cooperative mechanisms in reinforcement learning. In *Proc. AAAI*, pages 607–613, 1991.
179. E.O. Wilson. *The insect societies.* Harvard University Press, 1971.
180. H. Yamaguchi and T. Arai. Distributed and autonomous control method for generating shape of multiple mobile robot group. In *IEEE/RSJ IROS*, pages 800–807, 1994.
181. H. Yanco and L. Stein. An adaptive communication protocol for cooperating mobile robots. In *Proc. Simulation of Adaptive Behavior*, pages 478–485, 1992.
182. F.E. Yates, editor. *Self-Organizing Systems: The Emergence of Order.* Plenum Press, 1987.
183. D. Yeung and G. Bekey. A decentralized approach to the motion planning problem for multiple mobile robots. In *IEEE ICRA*, pages 1779–1784, 1987.
184. K. Yokota, T. Suzuki, H. Asama, A. Masumoto, and I. Endo. A human interface system for the multi-agent robotic system. In *IEEE ICRA*, pages 1039–1044, 1994.
185. E. Yoshida, T. Arai, J. Ota, and T. Miki. Effect of grouping in local communication system of multiple mobile robots. In *IEEE/RSJ IROS*, pages 808–815, 1994.
186. E. Yoshida, M. Yamamota, T. Arai, J. Ota, and D. Kurabayashi. A design method of local communication area in multiple mobile robot system. In *IEEE ICRA*, pages 2567–2572, 1995.
187. E. Yoshida, M. Yamamoto, T. Arai, J. Ota, and D. Kurabayashi. A design method of local communication range in multiple mobile robot system. In *IEEE/RSJ IROS*, pages 274–279, 1995.
188. 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 International Conference on Intelligent Robots and Systems, Raleigh, NC, July 7-10, 1992*, pages 1566–1574, July 1992.

Y. Uny Cao is a Ph.D. student in the Computer Science Department at the University of California at Los Angeles. He received a B.E. in Electrical Engineering from Zhejiang University, Hangzhou, China, and a M.S. in Computer Science in 1993 from the University of Louisville, Kentucky. At UCLA, he has been working on several research projects projects related to the Internet and robotics .

Alex S. Fukunaga is a Member of the Technical Staff in the Artificial Intelligence Group, Information and Computing Technologies Research Section at the Jet Propulsion Laboratory, California Institute of Technology. He holds an A.B. in Computer Science from Harvard University, and a M.S. in Computer Science from the University of California at Los Angeles, where he is currently a Ph.D. student. His research interests include optimization, search, machine learning, and automated planning/scheduling,

Andrew B. Kahng is an Associate Professor in the Computer Science Department at the University of California at Los Angeles. He received the A.B. degree in Applied Mathematics/Physics from Harvard College in 1983. His M.S.(1986) and Ph.D.(1989) degrees in Computer Science are from the University of California at San Diego. Dr. Kahng has received the NSF Research Initiation Award and an NSF Young Investigator Award. Currently, he is co-Director of the UCLA VLSI CAD Laboratory. He is also Director of the UCLA Commotion Laboratory, studying cooperation and distributed task-solving using multiple mobile robots. His research areas include discrete algorithms for VLSI layout synthesis, computational geometry, and search/recognition tasks, as well as the theory of large-scale global optimization.