

# Collective Robotic Intelligence

C. Ronald Kube and Hong Zhang

Department of Computing Science

University of Alberta

Edmonton, Alberta CANADA T6G 2J9

kube@cs.ualberta.ca zhang@cs.ualberta.ca

September 1, 1992

## Abstract

In this paper, we examine the problem of controlling multiple behaviour-based autonomous robots. Based on observations made from the study of social insects, we propose five simple mechanisms used to invoke group behaviour in simple sensor-based mobile robots. The proposed mechanisms allow populations of behaviour-based robots to perform tasks without centralized control or use of explicit communication. We have verified our collective control strategies by designing a robot population simulator called *SimbotCity*. We have also constructed a system of five homogeneous sensor-based mobile robots, capable of achieving simple collective tasks, to demonstrate the feasibility of some of the control mechanisms.

## 1 Introduction

Can simple behaviour-based mobile robots achieve tasks collectively? The behaviour-based approach, characterized by a direct coupling of perception to action, has been demonstrated by a number of researchers on situated and embodied mobile robots (see[16] for several papers). The majority of research projects have concentrated on designing *single* autonomous robots, capable of functioning in a dynamic world. Taking behaviour-based robotics towards more sophisticated capabilities has resulted in the interest, by a number of researchers, in collective task-achieving behaviour [5, 7, 30]. By organizing multiple robots into collections of task-achieving populations, we conjecture that useful tasks can be accomplished with simple behaviour-based control mechanisms. In this paper, we describe our approach to real-time multiple robot control motivated by several observations of collective behaviour made from the study of social insects.

Recent interest in multiple robot systems has produced several interesting ideas [2, 10, 17]. Brooks and Flynn [7] propose sending a colony of small robots to explore the

surface of the Moon. Dario *et al.* [8] propose a social organization of *societies* of cellular mobile robots where useful tasks are carried out through collaboration rather than individual effort. Additionally, Yuta and Premvuti [30] describe an approach to cooperation of multiple mobile robots using environmental resources while working toward a common goal.

Brooks [5] has outlined several issues raised in controlling multiple autonomous mobile robots to generate a global behaviour. Our own efforts to address the problem of real-time multiple robot control have led to a simple strategy involving *group behaviours* and five mechanisms with which to invoke them. The mechanisms proposed do not represent a comprehensive set, but rather strategies we have found based on observations from the study of social insects. The proposed mechanisms may be used individually or in combination to achieve a desired collective behaviour. Our method is fundamentally different from the approach taken in Distributed Robotic Systems[3] in that we do not use any centralized coordination; subsequently, the complexity of the system does not increase with the number of robots. In addition, no explicit communication is used to control the robots, but rather a form of implicit communication is used through passive sensing. Although the proposed mechanisms are not restricted to behaviour-based control architectures, we have found the approach is better suited for this style of implementation.

The five mechanisms presented here represent our initial exploration into multiple robot control. The first mechanism makes use of a common task and a simple cooperation strategy of non-interference. The second mechanism produces a *follow* behaviour that keeps a group of robots together in herds. The third mechanism uses environmental cues to invoke the group behaviour. The fourth mechanism allows the robot to invoke its group behaviour once it senses it is within a group. The fifth mechanism invokes group behaviour through *autostimulation*, a method in which an individual robot may invoke a behaviour within the group.

In order to test our proposed control mechanisms for collective behaviour, we have constructed a simulator. Simulation allows one to test the feasibility of a given control mechanism. We have tested the first two control mechanisms in simulation. However, given the importance of *situatedness* in this style of work and that we ultimately build these robots, it is very important to simulate only what we can build. To do otherwise would leave us open to the simplifications so often criticized of simulation work. A series of experiments in simulation, using our Robot Population Simulator *SimbotCity*, led to the refinement of our strategy which we then implemented in a system of physical robots consisting of five homogeneous behaviour-based mobile robots capable of achieving simple collective tasks.

The ultimate test of the validity of the control strategy requires the physical construction of the robots. The robots are each equipped with two photovoltaic goal sensors, two near-infrared robot-avoidance sensors and one stagnation sensor used to provide positive achievement feedback. As a demonstration, the robots form a homogeneous group of task-achieving autonomous agents capable of collectively locating and moving a box otherwise unmovable by a single robot. They accomplish this simple task without any central coordination or explicit communication between the robots. All communication is implicit and in the form of passive sensing and avoidance of other robots. The experimental result demonstrates the feasibility of constructing a simple homogeneous group of autonomous mobile robots to achieve tasks collectively.

The remainder of this paper is organized as follows. The next section describes several interesting examples of collective behaviour found in the literature on social insects. Based on these examples, we then present in section 3 our approach to controlling multiple mobile robots using *group behaviours* and the five mechanisms used to invoke them. Section 4 discusses our initial exploration of this biologically-inspired control strategy for autonomous behaviour-based mobile robots which we call the Collective Robotic Intelligence Project (CRIP). We then discuss our simulation results used to verify the first two mechanisms. Section 5 discusses a system of five homogeneous sensor-based mobile robots, capable of achieving simple collective tasks, which we constructed in order to test the feasibility of the first mechanism. Finally, section 6 provides a summary and discusses the advantages and disadvantages of this approach to controlling multiple mobile robots and suggests some directions for future research.

## 2 Collective Behaviour in Social Insects

Life provides us with countless examples of collective task achieving societies. Bees, ants and termites all function collectively in groups, efficiently accomplishing tasks with, seemingly simple, insect intelligence. Can the study of collective behaviour in social insects lead to a viable control strategy for autonomous robots? [9, 23, 27, 1] have proposed various cooperative strategies for autonomous robots inspired by social insects. By examining several specific examples of collective behaviour, we have found that the vast array of task achieving behaviours displayed by social insects are a result of the diverse assortment of sensing capabilities. Moreover, it is this sensor array that serves to trigger the behavioural patterns which result in the emergent collective behaviour [28, 18]. It would appear that evolution has made up for a lack of higher level reasoning, in insects, by increasing the number of behaviour producing sensors. Whereas man has made up for a lack of sensors by developing higher level reasoning.

For behavioural biologists one of the main problems in understanding how an insect society functions, is to be able to deduce collective activity from individual behaviour. Some believe that collective behaviour is not simply the sum of each participant's behaviour, as others emerge at the society level[20]. This they claim creates a paradox—how can individual ants appear so inefficient and disorganized, for example in their nest building activity, while at the same time build highly elaborate nest structures?

To answer such a question, Pasteels *et al.* [20] feel researchers adopt one of several attitudes. The first is to consider that their behaviour is far less random than it appears. This is the viewpoint taken by the majority of communication or division of labour studies. Ants are not random particles, they do communicate and subtle forms of division of labour are often observed.

A second attitude is to consider the behavioural variance as being irrelevant to the society's functioning. Descriptions of behavioural sequences are reported in deterministic terms, with only the functional acts being reported as the observations of the society are filtered.

Another view is to admit randomness at the individual level could be part of the society's functioning[20]. Due to their great number, Oster and Wilson [19] have suggested that social insects can well afford behavioural variance. This variance, they claim, could increase the probability that a social activity will eventually be performed. Their collective reliability more than compensates for the individual inefficiency.

In the absence of a stochastic theory, our attention has been focused on several specific examples of collec-

tive behaviour[14]. Using a common task, [29] described nest construction by weaver ants. Weaver ants construct the walls of their nest by folding leaves. They begin by spreading over its surface and randomly tug at any edge they can grasp. One part is turned more easily than the others, and the initial success causes other ants to aid the effort and abandon their own.

A second example of collective behaviour, we found interesting, involves invoking a *follow behaviour* used in such activities as *swarming*. Several different sensor modalities are used to invoke the behaviour. Tandem running, which uses tactile sensing, rapid running, which uses visual sensing and odour trails are all examples of a *follow behaviour* which cause the insects to move as a group[28].

A third example involves invoking collective behaviour by a stimulus in the environment. These *environmental cues* invoke the same behaviour in the insects almost simultaneously. For example, bees use light intensity to govern the amount of time used for food collecting behaviours [22]. Similarly, some ants begin their daily foraging activity upon the detection of the light of dawn[21]. Thus, in these cases events in the environment (i.e. dawn or dusk) serve to activate the collective behaviour.

Another example suggests that ants alter their behaviour when found in large groups. For example, worker ants were found to excavate the soil and attend larvae at a higher rate when found in large groups[11]. Wilson [28] found that when workers of *Acanthomyops claviger* are kept in solitude, they are nearly insensitive to the natural alarm substances of the species. In contrast, those placed in the same nest with a few hundred nestmates respond normally to the alarm substances.

The last example, which motivated our mechanisms, is called *autostimulation* and involves an individual insect invoking the collective behaviour using an alarm substance. In Stuart's nest reparation experiments [25, 26], termites used chemical odour to attract additional workers to the site of repair activity. Foraging workers were found to communicate with themselves when they dispense orientation pheromones in their odour trails and then follow the traces during the return journey to the target areas[12]. These examples demonstrate how an insect can invoke a collective behaviour (i.e. nest-building or trail-following) within its own colony.

Although the field of behavioral biology lacks a common theory that adequately explains the collective behaviour of social insects, it does provide many well researched observations that can be used in our quest for a theory to control multiple autonomous mobile robots. In the next section we outline our proposed mechanisms for invoking group behaviour in mobile robots.

### 3 Collective Behaviour in Autonomous Robots

The collective behaviour which we seek for our autonomous robots is in the same vein as emergent behaviour, a key concept, in Artificial Life [15]. Collective behaviour can be viewed as an emergent property<sup>1</sup> of a self-organizing system with a few simple rules of interaction. This emergent property results from the system interacting with its dynamic environment. Collective tasks must therefore be designed with an interaction loop comprising of the system and the environment, ultimately converging towards the desired performance. In order to achieve collective behaviour<sup>2</sup> we must first learn to control groups of autonomous robots. Our approach is to design robots with group behaviours. From the examples of the previous section we propose five mechanisms with which to invoke a group behaviour, resulting in the collective behaviour of the system.

The first mechanism used to invoke the group behaviour makes use of a *common goal* and *non-interference* as a simple form of cooperation. For example, if we design a group of robots to locate and converge upon a single object and while doing so the robots do not interfere with one another, we have then successfully controlled the movement of the group as a whole. Non-interference as suggested by Yuta and Premvuti [30] is a simple form of cooperation in which the robots do not interfere with the operation of other robots in the group. In our implementation (described in section 5) non-interference simply becomes *robot-avoidance*. Collective tasks can be designed in which the common task invokes the group behaviour. For example, robots equipped with heat and fire sensors roam an area until a fire causes them to converge and extinguish the flame.

The second mechanism creates a 'herd' of robots by using a *follow* behaviour. The collective task may require that groups of robots remain together. Consider grass cutting by a group of small robots who travel in herds. The control of the group is accomplished using a behaviour designed to keep them together; the task is accomplished by having all the robots execute a function (i.e. cutting) while the group is moved through its environment.

The third mechanism uses *environmental cues* to invoke group behaviour. Both dawn and dusk provide ants with a visual cue to begin or end food collecting behaviours and are an example of an environmental cue. Collective tasks can be designed to allow the environment, in which the robots work, to provide the cue that invokes the group behaviour. For example, a group of building cleaning robots, designed to keep the outside surface of

<sup>1</sup> See[24] for an excellent discussion of emergent properties.

<sup>2</sup> A precursor to Collective Intelligence.

buildings clean, would begin their activity at dusk by depositing photosensitive chemicals on the outside surface. Cleaning action would begin the next day when the chemical reacted with sunlight.

The fourth mechanism relies on the ability of the robot to detect if it is within a group before invoking its group behaviour. Sensors would be placed around the periphery of the robot, enabling it to detect other robots both in front and behind as well as to the left and right. The group behaviour is then executed once the robot finds itself surrounded by other robots. For example, bulldozer robots designed to level the ground would only be effective once a large group was formed travelling in the same direction. The robots in the center of the group would execute the group behaviour while those on the periphery would be responsible for navigating the group over the surface to be leveled. This mechanism to control multiple robots would be useful if the group behaviour required the robots to be in a certain formation, for example, when cleaning a hockey rink of snow.

The fifth mechanism invokes group behaviour through autostimulation. For example, suppose we are using a group of robots to search an area for a particular substance; once a single robot finds the substance it broadcasts a signal which in turn invokes a group behaviour in all robots receiving the signal. This method of control is different from the above mechanisms because it is a form of *self-facilitation*<sup>3</sup>.

The above mechanisms, taken from examples of collective behaviour in social insects, are designed to control groups of mobile robots. In doing so we hope to achieve tasks unsuitable or impossible for a single robot. In this respect, this work shares in the goals of the Distributed Robotic Systems [3] which are suited for such tasks as distributed actuation or distributed perception. The mechanisms proposed do not represent a comprehensive set, but rather strategies we have found based on observations from the study of social insects. In the next section we discuss the verification of the first two mechanisms by simulation.

## 4 Verification by Simulation

The ultimate goal of this work is to design and build a number of real physical robots capable of achieving simple tasks collectively. In this section we present our robot population simulator *SimbotCity* and discuss its use as a tool for investigating control mechanisms used to control populations of mobile robots. We have verified by simulation the first two control mechanisms of the previous

section. We have found it is possible to control a group of mobile robots using a common task and a simple non-interference cooperation strategy and that groups of mobile robots could be kept together in herds, and therefore controlled, using a simple *follow* behaviour. Other control mechanisms can be similarly examined with simple modifications.

### 4.1 Simulation Objectives

When multiple robots start to interact a whole series of new issues begin to surface. Brooks[6] outlined several of these issues, a subset of which we shall consider here.

**Emergence** Each robot's control system consists of a set of behaviours, we would like to see what the collective behaviour of a group of homogeneous robots will be. Further, if an incremental modification to the individual robot is made, we would like to determine its effect on the collective behaviour of the group.

**Cooperation** In achieving collective tasks, some form of cooperation will be necessary. This may simply take the form of not interfering with other robots, as they progress towards a common goal, as [30] have suggested or may involve some other form of cooperation requiring more explicit communication.

**Allocation** We would like to know the minimum number of robots necessary to accomplish a collective task. Brooks[6] refers to this as 'density-dependence'. Also, should we decide to use more than the minimum number, at what point does the system cease to be functional due to the glut of robots? For example, a task requiring the spatial distribution of robots along one side of an object can eventually have too many robots to perform the task effectively.

**Herding** There are advantages in keeping a group of robots together. Collectively they can respond much more quickly to a given stimulus than if they are more spatially distributed. An example might be a group of fire fighting robots whose *extinguish* behaviour activates upon fire detection. A group would respond quicker to the blaze and gain control easier than just one robot. Given the limited local perceptive abilities of the robots, what are the suitable behaviours needed for herding?

These are the issues we wish to investigate with simulation; and the lessons learned in the process, have served as a guide in building the physical system.

<sup>3</sup>A term taken from the psychological literature meaning communication that promotes rather than inhibits activity.




Figure 1: Initial robot configuration. Robots must locate and collectively push the black box.




Figure 2: Robot configuration after 378 simulation steps; box is being pushed upwards.

## 4.2 SimbotCity: A Robot Population Simulator

*SimbotCity* is our simulator used to model robot populations (see Figure 1). A robot model consists of a set of sensor and actuator resources as well as a set of behaviour modules that map sensor inputs to actuator outputs. Behavior arbitration is currently handled with a modified fixed priority *subsumption* architecture[4].

### 4.3 Control Mechanism 1: *common task*

In order to verify the feasibility of the first control mechanism, we created a simple common task requiring the collective efforts of at least two robots to accomplish. The task involved locating and pushing an object in the robots' environment, a task otherwise unaccomplishable by a single robot due to the weight of the object.

A simple form of cooperation called *non-interference*[30] was employed which successfully kept the robots from colliding with one another as they converged on the goal (see Figure 1 and Figure 2).

Since the task required the robots to spatially distribute themselves along the sides of the object, we found that any number of robots greater than the number that could fit on the surface of the object was redundant for this particular task.

Once an object was located, the robots collectively pushed the object off an edge of their world. Should a robot find itself pushing on the opposite side of a group, a progress behaviour which monitors progress towards the goal invokes an avoidance response allowing the robot to assume a new position on the object. Behaviours used to implement an individual robot's control are explained in[14].

By combining a common task and non-interference as a cooperation strategy a group of simulated robots could be successfully controlled to execute a collective task otherwise unaccomplishable by a single robot. The first mechanism is further verified in Section 5. Next we consider the second mechanism for group control.

### 4.4 Control Mechanism 2: *follow behaviours*

To test *follow* behaviours as a means of controlling a group of mobile robots, we equipped each simulated robot with sensors to detect other robots. The idea is based on equipping our physical robots with near-infrared transmitter-receiver pairs for robot detection.

A follow behaviour was designed which caused the robot to move towards a detected neighbouring robot. Robots begin by executing a random walk until a robot is detected, at which point the follow behaviour is invoked. Avoidance behaviours keep the forward moving robots from colliding. Herds begin to form with robots executing a random walk leading the herd (see Figure 3). Once formed, the herd remains intact by adapting the follow behaviour using a *behaviour preference*.

Behavior preferences are methods by which a behaviour module narrows its focus on sensor inputs. For example, the follow behaviour receives sensor input with an initial *sensor-view-angle* of 180 degrees, allowing the robot to 'see' in a very wide forward-looking direction. As the robot joins a herd by following a neighbouring robot its behaviour preference narrows the sensor's view angle to 45 degrees, ensuring the robot is not distracted from its group by another herd passing in close proximity. In this way the robot is able to adapt its behaviour to suit the task at hand. This method is similar to the way humans adjust their focus-of-attention in visual tasks.

The follow behaviour coupled with behaviour prefer-




Figure 3: Initial robot configuration. Robots must locate and follow another robot.




Figure 4: Robot configuration after 828 simulation steps; robots have formed a herd using a follow behaviour.

ences successfully controls a group of simulated robots and allows them to form a variety of ‘herd’ configurations (see Figure 4). This control mechanism can be combined with the first mechanism to allow groups of robots to roam in herds searching for objects to push. By keeping robots together in a herd the system of robots responds more quickly to the task at hand due to the simultaneous sensing of an object by several robots. We found this increases the likelihood that the distribution of robots around the perimeter of the object will be asymmetrical, causing the object to move in one direction quicker than if a symmetrical distribution occurred simultaneously. In the next section we discuss a system of five homogeneous mobile robots designed and constructed to further verify the first control mechanism.

## 5 Verification by Implementation

In order to test the *common task* control mechanism, verified by simulation, in the real world we have constructed

a system of five identical behaviour-based mobile robots capable of achieving simple collective tasks without centralized coordination or use of explicit communication. Control of the group of robots is accomplished by having each robot work toward a common goal. A simple form of cooperation among the robots is achieved through non-interference as they progress towards the common task. This is accomplished with a simple robot-avoidance behaviour. Each robot is autonomous and equipped with sensors for detecting both the goal (a brightly lit object) and obstacles in their environment.

In this section we present a simple collective task implementation and discuss its use in verifying the first control mechanism. We found it was possible to control a group of robots without the need for explicit communication or centralized coordination using simple reflexive behaviours and a common task control mechanism. We begin by outlining our objectives in constructing robots and present the collective task we have demonstrated.

### 5.1 Implementation Objectives

Building physical robots and testing their performance and dynamic interaction in the real world is the ultimate test for any *proposed* system. In doing so, we hoped to discover how well the proposed control mechanism did or did not work. By having already tested the control strategy in simulation we have lessened the burden and increased the likelihood for success.

The collective task we chose to implement was *box-pushing*. The task is such that it cannot be accomplished by one robot due to the weight of the box; therefore, the collective effort of several robots is necessary. To accomplish this task the robots must locate the box; move toward it while avoiding collisions with other robots; distribute themselves along a side and push. In doing so, we hoped to test the first control mechanism, namely, control of the group can result by executing a common goal oriented collective task. We also hoped to test the simple cooperative strategy of non-interference thereby verifying its usefulness in collective behaviour.

### 5.2 A Simple Collective Task

The box-pushing robot’s architecture consists of GOAL and AVOID behaviour modules implemented in simple combinational logic circuits. The GOAL behaviour functions to locate and guide the robot toward the box, while the AVOID behaviour handles obstacle avoidance including other robots. The control model is illustrated in Figure 5.

Each robot is equipped with five sensors. Sensors  $S1$  and  $S2$  are used by the GOAL behaviour to locate the




Figure 5: The box-pushing robot's control model used to control two wheel motors.

box. Sensors  $S3$  and  $S4$  provide forward looking obstacle detection and sensor  $S5$  is used to detect lack of progress toward the goal. A simple fixed priority between the behaviours controls arbitration with the AVOID behaviour having the highest priority.

All demonstrations designed to test the system were video recorded. The robots were placed in an initial configuration, much the same way as in simulation. By adjusting the goal sensors to respond to the light on the box only, the system could be started and stopped by turning the box light on and off. Video recordings were useful in that it allowed us to review test runs, making corrective adjustments in between.

The system was tested with a variety of initial configurations. The robots converged on the goal and pushed the box in a number of directions depending on how many robots were on a given side (see Figure 6, Figure 7, and Figure 8). In the event that a robot began pushing on a side opposite a group of robots, and subsequently was pushed backwards along with the box, a progress sensor invoked the avoidance behaviour allowing the robot to find a new spot on the box. As the robots progressed toward the box the AVOID behaviour kept the robots from colliding for the majority of the time. Collisions occurred whenever the sensors missed an oncoming robot. Reliability of the AVOID behaviour could be increased by adding additional obstacle sensors.

We also observed that stagnation could occur in the system if the robots approached the box from all sides resulting in an equal distribution and force applied to the sides of the box. In this case the robots did not move the box. This could be solved with a *stagnation detection* behaviour, to be implemented, which would have as input a sensor monitoring constant wheel motion.

The system demonstrated that the common task control mechanism was a feasible approach to controlling a small group of robots using a non-interference cooperation strategy. In the next section we present our conclusions and discuss future work.



Figure 6: Initial configuration of three box-pushing robots before sensing the box at the top of the picture.




Figure 7: Robots 1 and 2 about to collide before avoid behaviour activated.




Figure 8: Robot 1 overtakes robot 2 to avoid a collision, while progressing towards the box. Robots 1 and 3 pushing the box forward

## 6 Conclusions and Future Work

Research in behaviour-based robotics has led to radically different architectures for controlling autonomous robots. These new architectures emphasize a more direct coupling of perception to action and a dynamic interaction with the environment resulting in systems with emergent properties. Systems that choose to employ this methodology must be designed in a way that makes use of an interaction loop between the system and the environment which ultimately converges towards the desired performance. Most research projects have concentrated on designing *single* autonomous robots capable of achieving a simple insect-like intelligence. Useful tasks may be accomplished with these simple behaviour-based control mechanisms provided multiple robots are organized into collections of task achieving populations.

The research described in this paper attempts as a first step to propose five control mechanisms suitable for controlling populations of behaviour-based robots. Our approach to controlling multiple robots involves the use of group behaviours which may be invoked using several sensory-based mechanisms. The mechanisms proposed have resulted from the study of social insects which exhibit collective task achieving behaviours. To test our control mechanisms we created a simulator, called *SimbotCity*, which allowed us to create configurations of multiple



robots designed to achieve simple collective tasks. Once satisfied the control strategies were feasible, we then constructed a system of five physical robots designed to accomplish a simple collective task, without any centralized coordination or use of explicit communication. The approach, employed to control the group of five robots, involved having the robots work toward a common goal. Using non-interference as a simple form of cooperation, the robots were able to collectively locate and push a brightly lit box in their environment.

An important feature of the system is that simple reflexive behaviours can be used to control the individual robot in a goal directed manner using equally simple binary sensors. The behaviours and their arbitration mechanism are constructed using simple combinational logic. An important implication of this simplicity is that the control architecture could be scaled down to fit on a small silicon chip. This would allow for the creation of a large number of small cost effective robots to be used in areas too small for more traditional robots.

What is lacking in this approach to collective behaviour is a formal mathematics on which to base control models designed to solve tasks requiring multiple robots. Kiss[13] suggests using process dynamics, or Chaos Theory, with its well established mathematical theory, as a tool with which to study the behaviour of autonomous intelligent systems. If successfully applied, process dynamics could aid in the understanding of collective behaviour and its control models. A structured methodology for designing collective tasks is needed. Currently this is more of an 'art' than science. How should a control strategy be synthesized in terms of simple mechanisms? And what is their relationship in both time and space? Some of these questions could be answered if a formal theory can be established to analyze and design the control strategy.

The remaining three control mechanisms need to be verified in simulation as well as implementation. The research described is intended to be an initial exploration into achieving tasks collectively using a system of multiple robots. As such, its primary goal was to examine the feasibility of the approach we have outlined in this paper. Achieving collective behaviour with multiple robots is the first step in achieving Collective Robotic Intelligence.

## References

- [1] Ronald C. Arkin. Cooperation without communication: Multiagent schema-based robot navigation. *Journal of Robotic Systems*, 9(3):351-364, 1992.
- [2] Hajime Asama, Maki K. Habib, Isao Endo, Koichi Ozaki, Akihiro Matsumoto, and Yoshiki Ishida. Functional distribution among multiple mobile robots in an autonomous and decentralized robot system. In *1991 IEEE International Conference on Robotics and Automation*, pages 1921-1926, 1991.
- [3] Gerardo Beni and Jing Wang. Theoretical problems for the realization of distributed robotic systems. In *1991 IEEE International Conference on Robotics and Automation*, pages 1914-1920, 1991.
- [4] Rodney A. Brooks. A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, RA-2(1):14-23, March 1986.
- [5] Rodney A. Brooks. Challenges for complete creature architectures. In *First International Conference on Simulation of Adaptive Behavior*, pages 434-443. MIT Press, 1990.
- [6] Rodney A. Brooks. Intelligence without reason. AI Memo 1293, MIT, 1991.
- [7] Rodney A. Brooks and Anita M. Flynn. Fast, cheap and out of control. AI Memo 1182, MIT, 1989.
- [8] P. Dario, F. Ribechini, V. Genovese, and G. Sandini. Instinctive behaviors and personalities in societies of cellular robots. In *1991 IEEE International Conference on Robotics and Automation*, pages 1927-1932, 1991.
- [9] J-L. 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 *First International Conference on Simulation of Adaptive Behavior*, pages 356-363. MIT Press, 1990.
- [10] Toshio Fukuda, Tsuyoshi Ueyama, and Fumihito Arai. Control strategy for a network of cellular robots. In *1991 IEEE International Conference on Robotics and Automation*, pages 1616-1621, 1991.
- [11] W. Hangartner. Carbon dioxide, a releaser for digging behavior in *solenopsis geminata*. *Psyche*, 76:58-67, 1969.
- [12] B. Holldobler, M. Moglich, and U. Maschwitz. Communication by tandem running in the ant *camponotus sericeus*. *Journal of Comparative Physiology*, 90:105-127, 1974.
- [13] George Kiss. Autonomous agents, AI and chaos theory. In *First International Conference on Simulation of Adaptive Behavior*, pages 518-524. MIT Press, 1990.
- [14] C. Ronald Kube. Collective robotic intelligence: A control theory for robot populations. Master's thesis, University of Alberta, 1992.
- [15] C. Langton. Artificial life. In *Proceeding of the First Conference on Artificial Life*, pages 1-47, 1987.
- [16] Pattie Maes, editor. *Designing Autonomous Agents: Theory and Practice from Biology to Engineering and Back*. MIT Press, Cambridge, MA, 1990.
- [17] Maja J. Mataric. Minimizing complexity in controlling a mobile robot population. In *1992 IEEE International Conference on Robotics and Automation*, pages 830-835, 1992.
- [18] J.C. Moser. Pheromones of social insects. In D. Wood, R. Silverstein, and M. Nakajima, editors, *Control of Insect Behavior by Natural Products*, pages 161-178. Academic Press, 1970.

- [19] G.F. Oster and E.O. Wilson. *Caste and Ecology in Social Insects*. Princeton University Press, 1978.
- [20] J.M. Pasteels, J. Deneubourg, and S. Gross. Self-organization mechanisms in ant societies: Trail recruitment to newly discovered food sources. In J.M. Pasteels and J. Deneubourg, editors, *From Individual to Collective Behavior in Social Insects*, pages 155–175, 1987.
- [21] T.C. Schneirla. The army ants. In *Report of The Smithsonian Institution for 1955*, pages 379–406, 1955.
- [22] B. Schricker. Die orientierung der honigbiene in der dammerung zugleich ein beitrag zur frage der ocellenfunktion bei bienen. *Zeitschrift fur Vergleichende Physiologie*, 49:420–458, 1965.
- [23] Luc Steels. Cooperation between distributed agents through self-organisation. In Yves Demazeau and Jean-Pierre Muller, editors, *Decentralized A.I.*, pages 175–196, Amsterdam, 1990. North-Holland.
- [24] Luc Steels. Towards a theory of emergent functionality. In *First International Conference on Simulation of Adaptive Behavior*, pages 451–461. MIT Press, 1990.
- [25] A.M. Stuart. Alarm, defence, and construction behavior relationships in termites (*isoptera*). *Science*, 156:1123–1125, 1967.
- [26] A.M. Stuart. Social behavior and communication. *Biology of Termites*, 1:193–232, 1969.
- [27] G. Theraulaz, S. Goss, J. Gervet, and J-L. Deneubourg. Task differentiation in polistes wasp colonies: a model for self-organizing groups of robots. In *First International Conference on Simulation of Adaptive Behavior*, pages 346–355. MIT Press, 1990.
- [28] E.O. Wilson. *The Insect Societies*. The Belkap Press of Harvard University Press, 1971.
- [29] E.O. Wilson and B. Holldobler. *The Ants*. The Belkap Press of Harvard University Press, 1990.
- [30] S. Yuta and S. Premvuti. Consideration on cooperation of multiple autonomous mobile robots. In *IEEE International Workshop on Intelligent Robots and Systems*, pages 545–549, 1991.