In *USC Center for Robotics and Embedded Systems Technical Report CRES-04-001, Jan 2004.*

Communication in Multi-Robot Coordination

Chris Jones and Maja J Matarić
Computer Science Department
University of Southern California
941 West 37th Place, Mailcode 0781
Los Angeles, CA 90089-0781 USA
Email: {cvjones|@robotics.usc.edu}

***Abstract*— To enable the successful deployment of multi-robot systems (MRS), effective coordination mechanisms must be designed in a principled manner. There is currently little formal work addressing how to design such coordination mechanisms nor is there a formal understanding of the capabilities and performance characteristics that can be expected of various coordination mechanisms in a given task domain. We address this problem by presenting a formal framework from which to discuss and reason about coordination in a MRS and use this formalism to develop a method by which to automatically construct individual robot controllers in a MRS for a given task where the robots have the capability for inter-robot communication but maintain no non-transient state. We present a graph coloring approach to determine the minimal number of unique communication messages necessary such that the task is correctly executed. Understanding the capabilities and limitations of a MRS composed of such robots contributes to the understanding of when and why inter-robot communication is necessary to achieve the desired coordination and when it is not sufficient. We validate our automatic controller construction method in a simulated multi-robot construction domain.**

I. Introduction and Motivation

The design of coordination mechanisms for multi-robot systems (MRS) has received considerable attention in recent years. This is not surprising, as continually improving technology has made it realistically possible to deploy and use MRS consisting of increasingly larger numbers of robots. Successfully controlling such an MRS to properly and efficiently execute a given task has proven to be a difficult problem. Coordination among the individuals in the MRS is seen as an essential characteristic of any task-achieving MRS and the design and understanding of appropriate coordination mechanisms is a central challenge facing the community.

Work on the analysis of MRS coordination mechanisms is only beginning Gerkey and Matarić [7], and no general framework exists to guide the designer toward the form of coordination most befitting a given task domain and performance requirements. Although some elegant coordination mechanisms have been demonstrated in specific task domains, both in simulation and on physical robots, a formal understanding of why one coordination mechanism is more appropriate than another or how various mechanisms are related is lacking.

The range of possible coordination mechanisms spans from *explicit coordination* to *implicit coordination* mechanisms. Explicit coordination mechanisms make use of internal state maintained by individual robots and explicit inter-robot communication, often involving centralized or hierarchical control. Implicit coordination mechanisms, on the other hand, make use of fortuitous structure in the environment and its synergistic relationship to the task definition and the robots' sensing, control, and mobility characteristics to produce a different class of coordination techniques, often categorized by terms such as *emergent, self-organized, or stigmergic*.

To provide a formal basis from which to understand the relationships between this spectrum of coordination mechanisms, we are actively involved in ongoing work to develop a coordination formalism which provides a framework for precisely defining and reasoning about the intertwined entities intrinsically involved in any task-achieving multi-robot system – the task environment, task definition, and the capabilities of the robots themselves. We hope this formalism will provide useful insight into the question of how to rationally choose the most appropriate coordination mechanism for a given task domain and performance requirements as well as lead to a more formal understanding of how various classes of coordination are related.

Our previous work investigated the use of non-transient internal state in robot controllers, specifically, when such state was necessary and when it alone was insufficient for the MRS to achieve correct task execution. We also demonstrated how the robot controllers for a given MRS could be automatically constructed for a given task using our coordination formalism [10]. The work presented in this paper builds upon this foundation and is centered on MRS composed of robots equipped with the capability for inter-robot communication but lacking the ability to maintain non-transient internal state. We present an algorithm which automatically constructs the controllers for a MRS composed of such communicating, stateless robots by reducing the problem to a graph coloring approach. A formal understanding of the capabilities and limitations of such a system contributes to the understanding of when and why inter-robot communication is sufficient to achieve the desired coordination and when the maintenance of internal state becomes necessary. We demonstrate our formalism in a simulated multi-robot construction domain where we are able to provide specific task instances accompanied by formal explanations of the suitability of the use of inter-robot communication.

This paper is organized as follows. In Section II we provide the relevant related work. In Section III we provide formal definitions of the task environment, task definition, and robot characteristics used in the formalism. In Section IV we present an algorithm to construct the controller of each robot in a MRS such that a given task is correctly executed. In Section V we demonstrate and validate the formalism in a simulated multi-robot construction domain. In Section VI we draw conclusions and discuss future work.

II. Related Work

This section summarizes some of the most relevant related work involving characterization and analysis of coordination in multi-robot systems. The work of Parker [16] discusses the trade-offs of local versus global information for coordination in multi-robot systems. Beckers et al. [2] present a coordination mechanism in a multi-robot object clustering domain. Matarić [13] presents work on group coordination in multi-robot systems using a collection of simple basis behaviors. The information invariants work of Donald [5] addresses the problem of determining the information requirements to perform robot tasks and means in which this information may be acquired. Dudek et al. [6] present a taxonomy which classifies multi-robot systems based on communication and computational capabilities. In a clustering domain, Martinoli et al. [12] demonstrate how the collective behavior of a group of mobile robots can be accurately studied using a simple probabilistic model. They show how the results of the model are descriptive of the results obtained through experiments with real robots and in sensor-based simulations. Balch [1] presents hierarchical social entropy, an information theoretic method of analysis used to determine the extent of diversity among robots in multi-robot systems. Goldberg and Matarić [8] precisely define the foraging task for multi-robot systems and provide a collection of general distributed behavior-based coordination algorithms and their empirical evaluation. Gerkey and Matarić [7] present a formalism for the analysis of task allocation in multi-robot systems with an emphasis on explicit coordination mechanisms. Lerman and Galstyan [11] present a mathematical model of the dynamics of collective behavior in a multi-robot adaptive task allocation domain. The work presented in this paper builds upon our previous work which presented a coordination formalism that was used to automatically construct robot controllers which maintain internal state but do not have inter-robot communication capabilities [10].

We validate our approach in a simulated multi-robot construction domain. Related work in this area includes the work of Bonabeau et al. [4] which uses a rule-based model in the construction of biologically-plausible nest structures similar to those of some wasp species. Bonabeau et al. [3] investigate the use of genetic algorithms to generate such rules used in the construction of biologically- plausible structures and explores the relationship between the space of rules and resulting structures. In the area of construction by physical robots, Melhuish et al. [14] demonstrate how a group of minimalist robots can construct defensive walls using biologically-inspired templates. Wawerla et al. [17] present work on the comparison of different coordination strategies in the construction of simple 2D structures using a group of mobile robots. Jones and Matarić [9] present a method by which to automatically generate controllers for rule-based agents using local sensing and control in an intelligent self-assembly domain.

III. Definitions and Notation

In this section we formally define the intertwined entities intrinsically involved in any task-achieving multi-robot system – the task environment, task definition, and the capabilities of the robots themselves, including control, sensing, and inter-robot communication.

A. Task Environment

The task environment is the world in which the multi-robot system is expected to perform a defined task. The environment state, *s*, at any given time is an element of the finite set *S* of all possible states. An action, *a*, performed in the environment by a robot is drawn from the finite set *A* of all possible actions. An environment is defined by a state transition function $s_j = F(s_i, a)$, which states that when action $a \in A$ is executed in state $s_i \in S$, the next state will be $s_j \in S$. In this work, we assume the state can transition only as the result of an action performed by a robot.

B. Task Definition

We define a task, *T*, assumed to be Markovian, as a set of *n* ordered environment states, $T_s = \{s_1, ..., s_n\}$ which must be progressed through in sequence. From here on, the use of the word *state* refers to task state. An action *a* is called a *task action* for state s_i, denoted by $A_t(s_i) \in A$, if $s_{i+1} = F(s_i, a)$. A task *T* is said to be executed correctly if and only if for each task state $s_i \in T_s$ any executed action *a* falls into one of the two following categories: $a = A_t(s_i)$ or $s_i = F(s_i, a)$. This means that all performed actions are either task actions or are actions which do not result in a task state transition.

C. Observations

An observation made by an individual robot consists of accessible information external to the robot (e.g., camera image, sonar readings, etc.) and formally represents a subset of the task state. The content and properties of an observation are dependent on the specific sensing properties of the robot. The finite set of all possible observations is denoted as *X*. Since a given observation may occur in multiple states, for notational convenience we use x_s to mean the observation *x* as made in state *s*. An observation *x* with no sub-script, unless otherwise noted, refers to the observation *x* as made in any state. The distinction between observations as made in a specific state and a given observation as made in any state is made only for use in the formal framework to be described below, the robots themselves just make observations without any explicit knowledge of underlying state.

In state *s*, the function $G(s) \subseteq X$ returns the set of all observations which can be made in *s*. An observation *x* is called *unique* if and only if there exists only one state *s* for

```
procedure Build_Satisficing_Controller()
  E ← {}
  B ← {}
  for each v ∈ V do
    E ← E ⋃ {C_v(v) = E(O_v(v))}
  endfor
  for each s ∈ S do
    B ← B ⋃ {A_t(s_i) = B(Y(s_i), C_v(v) ∀v s.t. O_v(v) ∈ o_s}
  endfor
end procedure Build_Satisficing_Controller
```

Fig. 1. Procedure for building a satisficing controller.

which $x \in G(s)$. The observation at the physical location where the task action of state *s* is to be executed is denoted by $Y(s) \in G(s)$.

D. Robot Characterization

A robot's observation, *x*, at any given time is an element of *X*. The set of all possible communication messages a robot may send and receive is denoted by the set *C*. A robot may only send one communication message at a time and we assume all other robots instantaneously receive the communication. The communication message a robot is currently sending is denoted as $S_c \in C$. We assume a robot may receive any number of communication messages simultaneously. The set of communication messages a robot is currently receiving is denoted as $R_c \subseteq C$.

Two functions define a robot's action in the environment, known collectively as the robot's *controller*. The deterministic action function $a = B(x, R_c)$ specifies the robot's action, *a*, given its current observation is *x* and it is currently receiving the communication messages R_c. The communication function $c = E(x)$ is a deterministic function specifying robot will send communication message *c* given that its current observation is *x*.

A value of "$\cdot$" for any parameter in the *B* function signifies the set of all possible values for that parameter (i.e., "don't care"). For example, the rule $a = B(x, \cdot)$ specifies that if a robot makes observation *x*, it will perform action *a* regardless of what communications it is currently receiving.

IV. Building a Satisficing Controller

We now describe a method by which a controller using inter-robot communication can be automatically constructed such that it executes a given task correctly. We call such a controller *satisficing*. Correctly executing a task in a multi-robot system has additional challenges from doing so in a single robot system. It can never be assumed that a particular robot will or will not make a certain observation during a given time interval, as it could be the case that a robot is completely unaware of the progress of the task that is resulting from the actions of other robots. Formally, from the perspective of an individual robot, the task environment is highly non-stationary.

In constructing such a satisficing controller, it is trivial to see that if for a task state s_i for which $Y(s_i)$ is unique, no communication is required as the action rule $A_t(s_i) = B(Y(s_i), \cdot)$ is sufficient.

A more interesting case is if $Y(s_i)$ is *not* unique because $Y(s_i) \in G(s_j)$ for some state $s_j \neq s_i$. In such a situation, for a robot to execute the action $A_t(s_i)$ at the observation of $Y(s_i)$ it must be sure the state is indeed s_i and not s_j. To accomplish this with communication, we carefully define the robots' communication function such that a communication message is sent on the observation of designated observations. This set of designated observations, denoted by O_c, is chosen such that for every state s_i there is some set of observations, denoted by o_{s_i}, such that $o_{s_i} \subseteq O_c$ and $(o_{s_i} \bigcup Y(s_i)) \nsubseteq G(s_j)$, $s_i \neq s_j$.

Selecting such a set of observations O_c to communicate for each state is not addressed in this paper. Obviously, the set of observations chosen influences the number of unique communication message necessary. We assume O_c is given and then determine the minimal number of unique communication messages necessary for this set to ensure correct task execution.

The behavior function is then defined so that the task action $A_t(s_i)$ for state s_i will only be performed if a specific set of communication messages are currently being received. This set of communication messages corresponds to the set of observations o_{s_i}, the union of which can only be made in state s_i. When such communications are received by a robot observing $Y(s_i)$, the state will be disambiguated and the action $A_t(s_i)$ can be executed without compromising the correctness of the task execution.

We now explain how to map the given set of observations o_s, $\forall s \in S$ to a *minimal* number of unique communication messages using a graph (vertex) coloring approach. We begin by creating a graph G consisting of a set of verticies V and a set of edges E, both initially the empty set. For each unique observation in O_c, a vertex is added to V. We denote the observation represented by vertex x as $O_v(x)$. Next, we add edges to E between each pair of verticies in V for which the associated observations interfere with each other. The test for interference between two verticies v_i and v_j is given by the function $I(v_i, v_j)$ shown in Equation 1. We now apply a standard graph coloring algorithm to G in which unique communication messages are assigned to each vertex in V. We denote the communication message assigned through graph coloring to vertex v as $C_v(v)$.

Graph coloring is generally known to be NP-Complete; however, there are a number of well studied heuristics which provide understood bounds on resulting solution quality [15].

$$I(v_i, v_j) = \begin{cases} 1, & \text{if } \exists s \in S\text{: } v_i, v_j \in O_s, \\ 1, & \text{if } \exists s, t \in S\text{: } v_i \in O_s \text{ and } Y(s), v_j \in G(t), \\ 1, & \text{if } \exists s, t \in S\text{: } v_j \in O_s \text{ and } Y(s), v_i \in G(t), \\ 0, & \text{otherwise} \end{cases} \tag{1}$$

Given the mapping from the observations in O_c to unique communication messages provided by graph coloring, we

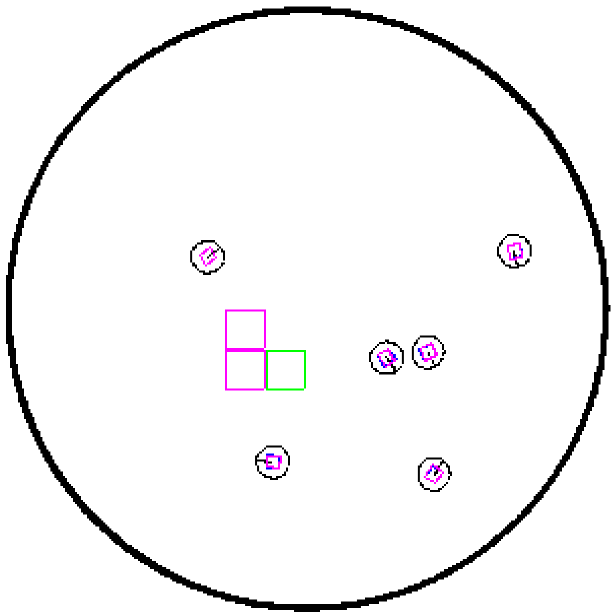

Fig. 2. A screenshot from the Stage simulation environment populated 3 bricks and 6 robots performing the construction task shown in Table I.

can now construct the action and communication functions constituting a satisficing robot controller. Figure 1 provides pseudocode by which the satisficing controller is constructed.

We note that inter-robot communication alone is not sufficient to guarantee correct task execution in a task environment for which there does not exist a unique set of observations o_s for some state s.

V. Validation: Coordination in Multi-Robot Construction

We experimentally demonstrate and validate our approach int the design of a satisficing controller in a simulated multi-robot construction task. The construction task requires the placement of a series of square colored bricks, 0.5 meters on a side, into a desired 2D planar structure in a specified sequence. For all examples used in this section, a brick's color is denoted by the letters `R`, `G`, `B`, and `Y` which stand for `Red`, `Green`, `Blue`, and `Yellow`, respectively. The construction task starts with a seed structure, which is a small number of initially placed bricks forming a core structure.

Experimental demonstration was performed using Player and the Stage simulation environment. Player (Gerkey et al. 2001) is a server that connects robots, sensors, and control programs over the network. Stage (Vaughan 2000) simulates a set of Player devices. Together, the two represent a high-fidelity simulation tool for individual robots and robot teams which has been validated on a collection of real-world robot experiments using Player control programs transferred directly to physical Pioneer 2DX mobile robots.

Our construction task is conducted in a circular arena of approximately 315 square meters using 6 robots. The robots are realistic models of the ActivMedia Pioneer 2DX mobile robot.

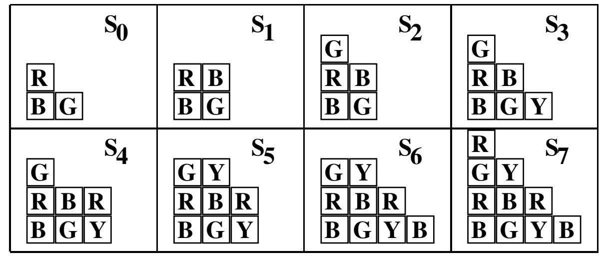

TABLE I

A sequence of environment states that define a construction task. The structure images are taken from a birds-eye-view. Each brick is labeled with its color: R=Red, B=Blue, G=Green, Y=Yellow.

Observations in $G(s_2)$
<FLUSH G B>
*<FLUSH B G>
<CORNER G B>
<FLUSH R G>
<FLUSH B R>

TABLE II

All observations in $G(s_2)$ from Table I. The observation $Y(s_2)$ is marked by a "*".

Each robot, approximately 30 cm in diameter, is equipped with a differential drive, a forward-facing 180 degree scanning laser rangefinder, and a forward-looking color camera with a 60-degree field-of-view and a color blob detection system. The bricks are taller than the robot's sensors, so the robots can only sense the bricks on the periphery of the structure. Figure 2 shows a screenshot of the simulation environment while the construction task is being performed.

A. Definition of the Construction Task

We define the environment state for the construction task as being a specific spatial configuration of bricks; therefore, a construction task is defined as a desired sequence of brick configurations – a specific construction sequence. The actions we are interested in are the placement of individual bricks to the growing structure; we do not consider construction tasks in which robots may remove bricks from the structure nor those in which sub-structures consisting of multiple bricks may be connected together. Other actions performed by the robots, such as moving through the environment, do not affect task state. Table I shows a set of environment states defining the example construction task that we will use throughout this section.

B. Observations in the Construction Task

Since the content of an observation is dependent on a robot's sensing capabilities, an observation in the construction domain is the spatial configuration and color of bricks in the field-of-view of the robot's laser rangefinder and color camera.

O_c	Observation
o_{s_2}	`<CORNER G B>`
o_{s_4}	`<FLUSH R Y>`
o_{s_5}	`<CORNER Y R>`
o_{s_6}	`<CORNER R B>`

TABLE III

SET OF OBSERVATIONS O_c FOR THE CONSTRUCTION TASK IN TABLE I.

There are two general categories of observations that can be made. The first is two adjacent, aligned bricks. Such an observation would be made, for example, if in state s_1 in Table I, a robot were positioned above and oriented toward the surface of the structure made up by the Red and Blue bricks. Such an observation is denoted as <FLUSH R B>.

The second observation category consists of two bricks forming a corner. Such an observation would be made if in state s_0 in Table I a robot were positioned in the upper right-hand corner and oriented toward the corner formed by the Red and Green bricks. Such an observation is denoted as <CORNER R B>.

The observations <FLUSH R B> and <FLUSH B R> constitute two different observations in which the spatial relationship between the Red and Blue bricks are switched. A similar point holds for the observations <CORNER R B> and <CORNER B R>. Given the task state s_2 from Table I, Table II lists all observations in the set $G(s_2)$ with the observation $Y(s_2)$ highlighted.

C. Brick Placement Actions

The only actions in this construction domain that can transition the task state are brick placement actions. There are three such actions, with the first being the placement of a brick on the right side (from the perspective of the acting robot) of a pair of adjacent, aligned bricks. An example action of this type can be seen in the placement of the Green brick which transitions the state in the task in Table I from s_1 to s_2. Such an action is denoted as <G RIGHT FLUSH R B>.

The second action type is similar to the first except the brick is placed on the left side of a pair of adjacent, aligned bricks. An example of this action type can be found in Figure I in the placement of the Yellow brick which transitions the state from s_2 to s_3. Such an action is denoted as <Y RIGHT FLUSH B G>.

The third action type is the placement of a brick in the corner formed by two other bricks. An example of this action type can be seen found in Figure I in the placement of the Blue brick which transitions the state from s_0 to s_1. Such an action is denoted as <B CORNER R G>.

D. Satisficing Robot Controller

Table III shows the observations for each o_s in O_c and Figure 3 shows the graph generated from the observations in O_c. The observations in this graph were individually assigned a communication message through the use of a graph coloring algorithm.

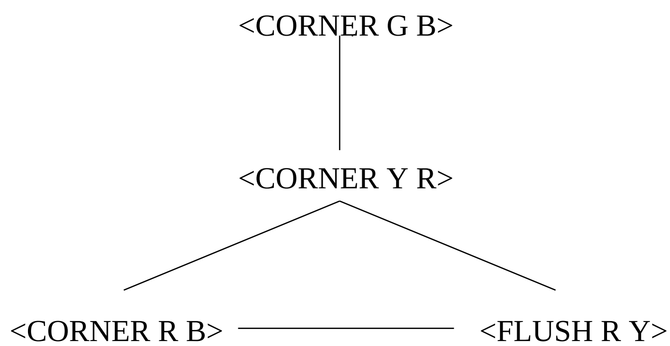

Fig. 3. Graph of observations shown in Table III used by graph coloring. Observations connected by an edge interfere with each other and will therefore be required to use different communication messages.

Observation	$\rightarrow$	S_c
`<CORNER G B>`	$\rightarrow$	c_1
`<FLUSH R Y>`	$\rightarrow$	c_2
`<CORNER Y R>`	$\rightarrow$	c_3
`<CORNER R B>`	$\rightarrow$	c_1

TABLE IV

COMMUNICATION FUNCTION CONSTITUTING A SATISFICING CONTROLLER FOR THE CONSTRUCTION TASK IN TABLE I. $c_1, c_2, c_3 \in C$.

The communication and action functions for the satisficing controller for the construction task shown in Table I are shown in Tables IV and V, respectively. The communication function specifies what, if any, communication message the robot should send given the current observation. The action function specifies what action the robot should execute given the current observation and some set of communication messages being received. A matched rule in the action function causes the robot to execute the corresponding brick placement action by visual servoing toward the location where the placement action is to be performed. Once the robot is within range to perform the action, the brick of appropriate color is placed on the structure. If no rule in the action function is matched, the robot performs a random walk, makes another observation, and the process repeats.

As can be seen, the satisficing robot controller for the construction task shown in Table I requires 3 unique inter-robot communication messages in order to guarantee correct task execution. A controller using fewer communication messages would result in a ambiguous situation and could lead to incorrect task execution.

VI. CONCLUSIONS AND FUTURE WORK

We have presented a method for automated multi-robot controller construction for correct task execution. The individual robots in the system execute controllers capable of inter-robot communication but do not maintain non-transient internal state. Given these individual robot capabilities, we have shown characteristics the task must exhibit such that these capabilities are sufficient for correct task execution. Understanding the

Observation	Req_Com	→	Action
<CORNER R G>	·	→	<B CORNER R G>
<FLUSH R B>	·	→	<G RIGHT FLUSH R B>
<FLUSH B G>	c_1	→	<Y LEFT FLUSH B G>
<CORNER B Y>	·	→	<R CORNER B Y>
<CORNER G B>	c_2	→	<Y CORNER B Y>
<FLUSH R Y>	c_3	→	<B LEFT FLUSH R Y>
<FLUSH G Y>	c_1	→	<R RIGHT FLUSH G Y>

TABLE V

ACTION FUNCTION CONSTITUTING A SATISFICING CONTROLLER FOR THE CONSTRUCTION TASK IN TABLE I.

capabilities and limitations of a multi-robot system composed of robots capable of inter-robot communication but lacking the use of internal state contributes insight into the larger question of understanding the necessary characteristics of a coordination mechanism in a multi-robot system required to correctly execute a given task. Furthermore, such understanding can aid the designer in making modifications to the task environment or definition or the robot capabilities in order to transform a situation in which inter-robot communication alone is not sufficient to one in which it is sufficient to achieve correct task execution.

Our future work includes investigations into formally defining when inter-robot communication and internal state may be interchanged and how a robot controller utilizing both may be automatically constructed.

VII. ACKNOWLEDGMENTS

This work is supported in part by Defense Advanced Research Projects Agency (DARPA) Grant F30602-00-2-0573 and in part by National Science Foundation Grant EDA-0121141.

REFERENCES

[1] T. Balch. Measuring robot group diversity. In T. Balch and L. E. Parker, editors, *Robot Teams: From Diversity to Polymorphism*, pages 93–135. AK Peters, 2002.

[2] R. Beckers, O. Holland, and J. Deneubourg. From local actions to global tasks: Stigmergy and collective robotics. In *Proc. of the 4th Intl. Workshop on the Synthesis and Simulation of Living Systems*, pages 181–189, 1994.

[3] E. Bonabeau, M. Dorigo, and G. Théraulaz. *Swarm Intelligence: From Natural to Artificial Systems*. Oxford University Press, 1999.

[4] E. Bonabeau, G. Théraulaz, E. Arpin, and E. Sardet. The building behavior of lattice swarms. In R. Brooks and P. Maes, editors, *Artificial Life IV*, pages 307–312, 1994.

[5] B. R. Donald. Information invariants in robotics. *Artifical Intelligence*, 72(1–2):217–304, 1995.

[6] G. Dudek, M. Jenkin, E. Milios, and D. Wilkes. A taxonomy for multi-agent robotics. *Autonomous Robots 3*, pages 375–397, 1996.

[7] B. Gerkey and M. Matarić. Multi-robot task allocation: Analyzing the complexity and optimality of key architectures. In *Proc. of the IEEE Intl. Conf. on Robotics and Automation*, pages 3862–3867, Taipei, Taiwan, 2003.

[8] Dani Goldberg and Maja J. Matarić. Design and evaluation of robust behavior-based controllers. In T. Balch and L. E. Parker, editors, *Robot Teams: From Diversity to Polymorphism*, pages 315–344. AK Peters, 2002.

[9] C. Jones and M. Matarić. From local to global behavior in intelligent self-assembly. In *Proc. of the IEEE International Conference on Robotics and Automation*, pages 721–726, Taipei, Taiwan, 2003.

[10] C. Jones and M. Matarić. Towards a multi-robot coordinaton formalism. In *Proceedings of the Workshop on the Mathematics and Algorithms of Social Insects*, Atlanta, Georgia, 2003. To Appear.

[11] K. Lerman and A. Galstyan. Macroscopic analysis of adaptive task allocation in robots. In *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems*, pages 1951–1956, Las Vegas, Nevada, Oct 2003.

[12] A. Martinoli, A. Ijspeert, and F. Mondada. Understanding collective aggregation mechanisms: From probabilistic modeling to experiments with real robots. *Robotics and Autonomous Systems*, 29:51–63, 1999.

[13] M. Matarić. Designing and understanding adaptive group behavior. *Adaptive Behavior*, 4:51–80, December 1995.

[14] C. Melhuish, J. Welsby, and C. Edwards. Using templates for defensive wall building with autonomous mobile ant-like robots. In *Towards Intelligent Mobile Robots (TIMR)*, Bristol, UK, 1999.

[15] R. Nelson and R. Wilson. *Graph Colourings*. Longman Scientific and Technical, 1990.

[16] L. E. Parker. Designing control laws for cooperative agent teams. In *Proc. of the IEEE Intl. Conf. on Robotics and Automation*, pages 582–587, 1993.

[17] J. Wawerla, G. S. Sukhatme, and Maja J Matarić. Collective construction with multiple robots. In *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS-02)*, pages 2696–2701, Lausanne, Switzerland, Sep 2002.