

## Synthesis and Analysis of Non-Reactive Controllers for Multi-Robot Sequential Task Domains

Chris Jones and Maja J Matarić

Computer Science Department  
University of Southern California  
Los Angeles, CA 90089-0781, USA  
{cvjones|maja}@robotics.usc.edu

**Abstract.** In this paper we present a macroscopic model for the analysis of homogeneous task-directed multi-robot systems (MRS). The model is used to compute the probability that a given MRS will correctly execute a given sequential Markovian task. We consider distributed MRS composed of non-communicative robots that maintain a limited amount of non-transient internal state. The model shares a common formal framework with our past work on the development of systematic methods for the synthesis of MRS. As such, it can be used to improve design procedures by providing an analytical approach to the evaluation of design decisions. The unified nature of the modeling and synthesis methods are part of our ongoing work toward a general, comprehensive and principled MRS design methodology. We apply the model to the analysis of system performance in a multi-robot construction task domain. Our past work on synthesis methods provides robot controllers for this task domain; the model provides quantitative predictions of system performance. We discuss the assumptions and limitations inherent in the model and discuss how the complementary synthesis and analysis methods may be more fully integrated.

### 1 Introduction

A critical undertaking in the design of task-directed multi-robot systems (MRS) is managing the complexity introduced by the coordination of multiple, interacting robots. In order for such systems to be effective, the robots' actions must be carried out in a coordinated fashion and directed toward task achievement. The system overall effectiveness typically strongly depends on the underlying coordination mechanism(s) used to mediate the interactions among the robots and between the robots and the task environment. Such coordination mechanisms, in turn, typically depend on the use of inter-robot communication and the maintenance of some form of internal state or representation by each of the robots.

Our work is focused on the design and analysis of homogeneous distributed MRS, in which each robot operates independently, under local sensing and control. The design of such systems, especially at a large-scale, is challenging and expensive. Unexpected collective behaviors may emerge due to unanticipated ramifications of the robots' local interactions, often resulting from uncertainty in sensing and action. There exist few principled tools to aid in the design process; resource-intensive design through trial and error is frequently the only available option.

To address the lack of principled tools for MRS design, our previous work presented a formal framework and an accompanying suite of principled, task-*nonspecific* methods for the synthesis of coordinated MRS. Each method represents a systematic approach using a variety of known coordination tools, including inter-robot communication [6] and the maintenance of internal state [7]. In the case of the latter, one is confronted with many options as to how the robots should utilize internal state to achieve the desired system-level coordination. System performance can frequently be improved by using more internal state. However, there is a point in the design space at which the cost of additional internal state outweighs improvements in performance. Analogously, a stateless system can usually be used if resulting task performance or efficiency are not a critical design considerations. Intelligently evaluating such trade-offs through direct experimentation is inefficient and costly; our work is aimed toward a formal framework that can address these issues.

In this paper, we present a macroscopic model for calculating the probability that a given homogeneous task-directed MRS, comprised of non-communicative robots that maintain a limited amount of non-transient internal state, will correctly execute a given sequential Markovian task. This model can be used to improve the design process by aiding in decisions related to the appropriate use of internal state and its trade-offs with resulting system task performance. The model shares a common formal framework with our synthesis methods. The unified nature of the synthesis methods and the modeling technique are part of our ongoing work to develop a complete and systematic MRS design methodology. The use of the unified synthesis and analysis, in conjunction with indispensable real-world experiments, should allow for a more systematic, efficient, and complete exploration of the design space, thereby leading to more effective design.

We apply the model to the analysis of coordinated MRS in a multi-robot construction task domain. Our controller synthesis method, presented in [7], provides robot controllers for this domain. Analysis of the synthesized MRS is then performed using the analytical model presented in this paper. The model is shown to provide quantitatively correct predictions of system performance. Extensive experiments were conducted in physically-realistic simulation to verify the correctness of the model's predictions. A limited number of real robot experiments were also performed to validate the simulation results [7].

## 2 Related Work

Work related to the analysis of coordinated MRS includes the work of Donald [2], who presented the derivation of information invariants aimed at defining the information requirements of a given task and ways in which those requirements can be satisfied in a robot controller. Donald et al. [1] applied information invariants to the study of distributed manipulation. Parker [13] extended the idea of information invariants by defining equivalence classes among task definitions and robot capabilities to assist in the choice of appropriate controller class. Sugawara and Sano [14] studied a macroscopic modeling technique applied to a multi-robot foraging task.

Martinoli et al. [12] presented a macroscopic probabilistic modeling methodology for the study of collective robot behavior. Martinoli and Easton [11] discussed the effectiveness and accuracy of microscopic and macroscopic modeling techniques compared to real robot experiments and embodied simulations. Gerkey and Mataric [3] presented a principled framework and analysis methodology, based on theories from economics and operations research, toward the formal study of multi-robot task allocation. Lerman et al. [10] presented a macroscopic analytical model of the dynamics of collective behavior in a collaborative stick-pulling domain using a series of coupled differential equations. Lerman and Galstyan [9] described a general macroscopic model for the study of adaptation in multi-agent systems and applied it to analysis of a multi-robot adaptive task allocation domain that our own work had previously addressed experimentally [5].

### 3 Definitions and Notation

The *world* is the domain in which the MRS is expected to perform a defined task. We assume the world is Markovian and the state is an element of the finite set  $S$  of all possible states. The set of all robots is denoted by the finite set  $R$ . We assume the robots are homogeneous. An action  $a_r$  performed in the world by a single robot  $r$  is drawn from the finite set  $A$  of all possible actions. An *observation*  $x$  made by robot  $r$ , drawn from the finite set of all observations  $X$ , consists of accessible information external to the robot and formally represents a subset of the world state. The world is defined by a probabilistic state transition function  $P : S \times X \times A \times S \rightarrow [0, 1]$ . That is, given a world state  $s$  at time  $t$ , a robot  $r$  making observation  $x$  and executing action  $a$ , and a world state  $s'$  at time  $t + 1$ ,  $P(s, x, a, s') = Pr(S^{t+1} = s' | S^t = s, X_r^t = x, A_r^t = a)$ . We note that the world state transition function involves an observation because the tasks we consider are spatial in nature and the physical location where an action is performed is just as important as the action itself. In this representation, an observation  $x$  is equated with the spatial location where the action  $a$  is performed. Therefore, an action  $a$  executed upon the observation of  $x_i$  will transition the world differently than the same action  $a$  performed upon the observation of  $x_j$ . We define a *task*, assumed to be sequential and Markovian, as a set of  $n$  ordered world states  $T_s = \{s_0, s_1, \dots, s_n\}$  which must be progressed through in sequence. We assume the initial state of the world is  $s_0$  and state  $s_n$  is terminal. We define *correct task execution* to be the case where, for all task states  $s_i \in T_s, i < n$  the only actions executed by any robot are those that transition the world state to  $s_{i+1}$ . Once the world state is  $s_n \in T_s$  the task is terminated. Therefore, we define an observation and action pair for a robot,  $x$  and  $a$ , to be correct for task state  $s_i$  if  $P(s_i, x, a, s_{i+1}) > 0$ . We assume that an observation  $x$  and action  $a$  cannot be correct for more than one task state. The probabilistic *observation function*  $O(s, x) = Pr(X_r^t = x | S^t = s)$  gives the probability observation  $x$  will be made in state  $s$  by a robot  $r$ . For each state  $s$ , the observation function is normalized as follows:  $\sum_{\forall x} O(s, x) = 1$ . A robot's internal state value  $m$  at any time is a member of the finite set  $M = \{m_0, \dots, m_p\}$ , with  $m_0$  being the initial internal state

value for all robots. Two probabilistic functions define a robot  $r$ 's behavior in the world, known collectively as the robot's *controller*. The controller is comprised of an *action function*  $A(x, m, a) = Pr(A_r^t = a | X_r^t = x, M_r^t = m)$  and an *internal state transition function*  $L(m, x, m') = Pr(M_r^{t+1} = m' | X_r^t = x, M_r^t = m)$ . We prohibit the internal state transition function from containing cycles. Although the controller is modeled with probabilistic functions for generality, in this paper these functions are assumed to be binary –  $A$  and  $L$  will always be either 0 or 1.

## 4 Macroscopic Model

Using the above formal framework, we next present a macroscopic analytical model for calculating the probability that a given MRS will correctly execute a particular sequential Markovian task. The model is specifically intended for a homogeneous MRS composed of non-communicative robots whose actions are based on their current local observation and internal state. A prime difficulty in constructing such a controller is in defining how available internal state values should be assigned and transitioned during the course of task execution so as to maximize resulting system performance.

During task execution, errors in sensing can inadvertently cause inappropriate internal state transitions and, consequently, the execution of incorrect actions. Therefore, the accuracy of our model relies on an understanding of how, given a controller, the robots' internal state values change during task execution. To address this issue, Equation 1 presents a method to calculate the probability a robot's internal state value will be  $m_j$  in task state  $s_i$ . This allows for the recursive calculation of the probability distribution of robot internal state values over task states. We note that  $Pr(m_j | s_i) = 0$ , if  $j < 0$ ,  $j \geq |M|$ ,  $i < 0$ , or  $i \geq |T_s|$ . Similarly,  $L(m_j, x, m_k) = 0$ , if  $j < 0$ .

$$Pr(m_j | s_i) = 1 - ((1 - Pr(m_{j-1}, s_i))(1 - \prod_{\forall x} (1 - O(s_i, x) \cdot L(m_{j-1}, x, m_j))))(1 - Pr(m_j | s_{i-1})(1 - Pr(m_{j+1} | s_{i-1}))) \quad (1)$$

Equation 2 provides a method to calculate the probability that a homogeneous MRS, composed of robots executing a given controller, will achieve correct task execution (CTE).

$$Pr(CTE) = \prod_{\forall s_i \forall m \forall x \forall a} (1 - Pr(m | s_i) O(s_i, x) A(x, m, a) \cdot (1 - (P(s_i, x, a, s_{i+1}) + P(s_i, x, a, s_i)))) \quad (2)$$

As with any modeling technique, the accuracy of the results depends on the accuracy and completeness of the underlying model. Constructing a model can be tedious and difficult. Fortunately, the definitions required by our model are task domain specific but are not specific to a particular task instance in that domain. Therefore, a single

model of a given task domain can be used unchanged for many specific task instances within that domain.

We assume there is no implicit “structure” in the world or the movements of the robots that correlates observation or action sequences. To satisfy this assumption, or at least minimize the likelihood it is not valid, in our investigations all robots execute random walks when not engaged in task-specific actions.

We are currently considering a number of extensions to our modeling approach. The approach and the MRS synthesis methods presented in our previous work are not fully integrated. Currently, the model is used for *post hoc* analysis of a previously synthesized MRS. However, as they share a common formal foundation, we are in the process of integrating them toward a fully unified framework for MRS synthesis and analysis. The analysis method will provide online analysis during synthesis, thus enabling more effective and efficient MRS design. Also, the current model does not consider robots capable of inter-robot communication. We are now extending the model to include such capabilities, thereby enabling the modeling and analysis of MRS composed of communicative robots that maintain internal state.

## 5 Case Study: Coordination in Multi-Robot Construction

We apply the macroscopic model to the analysis of coordinated behavior in a multi-robot construction task domain. This task requires the sequential placement of a series of cubic, colored bricks into a planar structure. For all examples, a brick’s color is denoted by the letters R, G, B, Y, C, and Bl which stand for Red, Green, Blue, Yellow, Cyan, and Black, respectively. The construction task starts with a seed structure, a small number of initially placed bricks forming the core of the structure.

Experimental data were collected through extensive physically-realistic simulation trials, using Player and Gazebo simulation environments, and through a limited number of real robot experiments. Player [4] is a server that connects robots, sensors, and control programs over a network; Gazebo [8] simulates a set of Player devices in a 3-D physically-realistic world with full dynamics. Together, the two represent a high-fidelity simulation tool for individual robots and teams that has been validated on a collection of real robot experiments using Player control programs transferred directly to physical Pioneer 2DXs. In all simulation experiments 8 robots were used, and in all real robot experiments 3 robots were used. The robots were either realistic models of or actual ActivMedia Pioneer 2DX mobile robots. 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 100-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 local bricks on the periphery of the structure (i.e., robots do not have a birds-eye view of the entire structure). Figure 2 shows snapshots of our simulation and real-world experimental setup.




Fig. 1. The sequence of brick placements defining construction task 1 (upper) and 2 (lower).





Fig. 2. Gazebo simulation (left) and real robot (right) MRS construction experiment snapshots.

Our robots do not have the ability to independently manipulate bricks during the construction process, either in simulation or in the physical world. To address this issue in simulation, when a robot wants to execute a brick placement action, it commands the simulator to place a brick of a given color at a specified location relative to the robot’s current pose; this process contains action uncertainty. In real robot experiments, the experimenter manually placed the appropriate brick in response to the robot’s audible command (e.g., “Place yellow brick in the corner formed by the red and blue bricks directly in front of my position”). Due to the manual brick placement, there is no associated action error in real robot experiments.

### 5.1 Formal Definitions for Construction Task

To ground the construction task in the formal framework presented in Section 3, we now define the world, task definitions, observations, and actions in the construction task domain. The *world* state is defined as a specific spatial configuration of bricks, including their color. A construction *task* is defined as a sequence of brick configurations (i.e., world states), providing a specific construction sequence. The two construction tasks discussed in this paper are shown in Figure 1. Observations in the construction domain are made up of the spatial configuration and color of bricks in the field-of-view of the robot’s laser rangefinder and color camera and within an appropriate range and bearing. Two categories of observations are considered. The



**Fig. 3.** Example observations and actions in the construction domain. Far left: Robot in position to make observation  $\langle \text{FLUSH R B} \rangle$ . Center left: Immediately after robot performs action  $\langle \text{G RIGHT FLUSH R B} \rangle$ . Center right: Robot in position to make observation  $\langle \text{CORNER R B} \rangle$ . Far right: Immediately after robot performs action  $\langle \text{G CORNER R B} \rangle$ .

first is two adjacent, aligned bricks. A situation in which such an observation would be made is shown in Figure 3 and is denoted as  $\langle \text{FLUSH R B} \rangle$ . The second is two adjacent bricks forming a corner. A situation in which such an observation would be made is shown in Figure 3 and is denoted as  $\langle \text{CORNER R B} \rangle$ . The observations  $\langle \text{FLUSH R B} \rangle$  and  $\langle \text{FLUSH B R} \rangle$  constitute two different observations in which the spatial relationships between the Red and Blue bricks are mirrored. A similar point holds for the observations  $\langle \text{CORNER R B} \rangle$  and  $\langle \text{CORNER B R} \rangle$ .

Actions 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 substructures consisting of multiple bricks may be connected together. Other actions performed by the robots, such as moving through the environment, do not affect the world state and therefore do not need to be explicitly considered. Three categories of actions are considered. The first is the placement of a brick on the right side (from the perspective of the acting robot) of a pair of adjacent, aligned bricks; the immediate result is demonstrated in Figure 3 and is denoted as  $\langle \text{G RIGHT FLUSH R B} \rangle$ . The second is identical to the first except that the brick is placed on the left side of a pair of adjacent, aligned bricks. This action is denoted as  $\langle \text{G LEFT FLUSH R B} \rangle$ . The third is the placement of a brick in the corner formed by two other bricks; the immediate result is demonstrated in Figure 3 and is denoted as  $\langle \text{G CORNER B R} \rangle$ .

## 5.2 Analysis Using the Macroscopic Model

We now apply the model to the analysis of task performance in the construction of the two tasks shown in Figure 1. The controllers executed by each robot for both tasks are shown in Table 1. The formal method by which they were synthesized, complete controller description and associated simulation and real robot experimental results are found in [7].

Task 1 Action Function
$A(\langle \text{CORNER R G} \rangle, m_0, \langle \text{B CORNER R G} \rangle) = 1$
$A(\langle \text{FLUSH R B} \rangle, m_0, \langle \text{G RIGHT FLUSH R B} \rangle) = 1$
$A(\langle \text{FLUSH B G} \rangle, m_1, \langle \text{Y LEFT FLUSH B G} \rangle) = 1$
$A(\langle \text{CORNER B Y} \rangle, m_1, \langle \text{R CORNER B Y} \rangle) = 1$
$A(\langle \text{FLUSH R Y} \rangle, m_1, \langle \text{B LEFT FLUSH R Y} \rangle) = 1$
$A(\langle \text{FLUSH B R} \rangle, m_2, \langle \text{Y RIGHT FLUSH B R} \rangle) = 1$
$A(\langle \text{FLUSH G Y} \rangle, m_2, \langle \text{R RIGHT FLUSH G Y} \rangle) = 1$


Task 1 Internal State Transition Function
$L(m_0, \langle \text{CORNER G B} \rangle, m_1) = 1$
$L(m_1, \langle \text{CORNER R B} \rangle, m_2) = 1$

Task 2 Action Function
$A(\langle \text{FLUSH R B} \rangle, m_0, \langle \text{G RIGHT FLUSH R B} \rangle) = 1$
$A(\langle \text{CORNER G B} \rangle, m_0, \langle \text{Y CORNER G B} \rangle) = 1$
$A(\langle \text{FLUSH B R} \rangle, m_1, \langle \text{R LEFT FLUSH B R} \rangle) = 1$
$A(\langle \text{CORNER B R} \rangle, m_1, \langle \text{G CORNER B R} \rangle) = 1$
$A(\langle \text{FLUSH R G} \rangle, m_2, \langle \text{C LEFT FLUSH R G} \rangle) = 1$
$A(\langle \text{FLUSH B G} \rangle, m_3, \langle \text{Y LEFT FLUSH B G} \rangle) = 1$
$A(\langle \text{CORNER R C} \rangle, m_4, \langle \text{G CORNER R C} \rangle) = 1$
$A(\langle \text{FLUSH Y B} \rangle, m_5, \langle \text{B LEFT FLUSH Y B} \rangle) = 1$
$A(\langle \text{CORNER R G} \rangle, m_6, \langle \text{B1 CORNER R G} \rangle) = 1$
$A(\langle \text{CORNER Y B} \rangle, m_7, \langle \text{C CORNER Y B} \rangle) = 1$

Task 2 Internal State Transition Function
$L(m_0, \langle \text{FLUSH G Y} \rangle, m_1) = 1$
$L(m_1, \langle \text{FLUSH G R} \rangle, m_2) = 1$
$L(m_2, \langle \text{CORNER R C} \rangle, m_3) = 1$
$L(m_3, \langle \text{CORNER B Y} \rangle, m_4) = 1$
$L(m_4, \langle \text{FLUSH G C} \rangle, m_5) = 1$
$L(m_5, \langle \text{CORNER Y B} \rangle, m_6) = 1$
$L(m_6, \langle \text{FLUSH B1 G} \rangle, m_7) = 1$

**Table 1.** Controller action and internal state transition functions for construction task 1 (top) and 2 (bottom) as shown in Figure 1.  $m_0, m_1, \dots, m_7 \in M$ . The robots' initial internal state value is  $m_0$ . All probabilities not shown are 0.

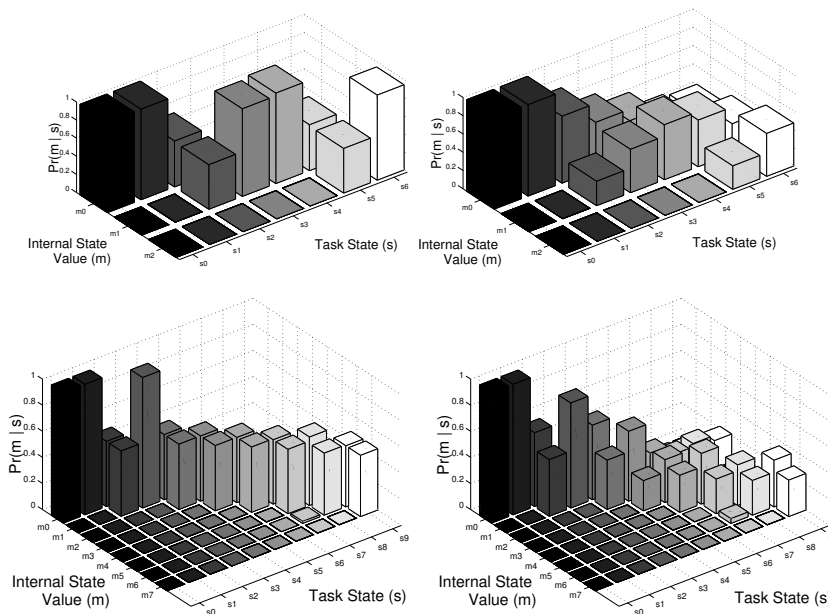
In order to build the model needed in the analysis, we ran extensive simulation trials on a baseline construction task different from the tasks used in this section. Using experimental data from these trials, we constructed the required model, including the world state transition function and the observation function, which was applied unchanged to the analysis of both construction tasks discussed in this section.



**Fig. 4.** Analysis for construction task 1 (left) and task 2 (right). The probability of correct task execution given for each task state is the probability the task was correctly executed up to and including the action performed in that task state.

For each of the two tasks shown in Figure 1, we ran 500 simulation experiments. Figure 4 compares the results to those provided by the model for tasks 1 and 2. As can be seen, the model provides quantitatively correct predictions of task performance. The inaccuracies seen in the model's predictions are primarily due to two factors. First, the action and sensing uncertainties observed in the experimental trials can deviate from the values used by the model. Second, the method used to calculate the distribution of internal state values over task states is in need of improvement. As is evident in Figures 5, the probability distribution of internal





**Fig. 5.** Predicted (left) and actual (right) probability distribution of internal state values over task states for construction task 1 (top) and 2 (bottom).

state values predicted by the model is generally in qualitative agreement with the experiments, but quantitatively leaves much to be desired. Errors in this part of the model are most likely to manifest themselves due to perceptual aliasing in temporally proximate task states.

## 6 Conclusions

We have presented a macroscopic modeling approach for the predictive analysis of homogeneous MRS task performance. The model was shown to provide quantitatively correct results in predicting the probability that a MRS would correctly execute a given sequential Markovian task in a multi-robot construction task domain. Extensive experiments in physically-realistic simulations were conducted to validate the model. Real robot experiments are currently underway in the construction task domain to provide more extensive validation of the model.

## 7 Acknowledgments

This work was supported by Defense Advanced Research Projects Agency (DARPA) Grant F30602-00-2-0573.

## References

1. B. Donald, J. Jennings, and D. Rus. Information invariants for distributed manipulation. In K. Goldberg, D. Halperin, J.-C. Latombe, and R. Wilson, editors, *International Workshop on the Algorithmic Foundations of Robotics*, pages 431–459, Boston, MA, 1995.
2. B. R. Donald. Information invariants in robotics. *Artificial Intelligence*, 72(1–2):217–304, 1995.
3. B. Gerkey and M. Mataric. Multi-robot task allocation: Analyzing the complexity and optimality of key architectures. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 3862–3867, Taipei, Taiwan, Sep 2003.
4. B. Gerkey, R. Vaughan, K. Stoey, A. Howard, G. Sukhatme, and M. Mataric. Most valuable player: A robot device server for distributed control. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1226–1231, Maui, Hawaii, Oct 2001.
5. C. Jones and M. Mataric. Adaptive division of labor in large-scale minimalist multi-robot systems. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1969–1974, Las Vegas, Nevada, Oct 2003.
6. C. Jones and M. Mataric. Automatic synthesis of communication-based coordinated multi-robot systems. Technical report, University of Southern California Center for Robotics and Embedded Systems, CRES-04-007, 2004.
7. C. Jones and M. Mataric. Toward a multi-robot coordination formalism. Technical report, University of Southern California Center for Robotics and Embedded Systems, CRES-04-008, 2004.
8. N. Koenig and A. Howard. Design and use paradigms for gazebo, an open-source multi-robot simulator. Technical report, USC Center for Robotics and Embedded Systems, CRES-04-002, 2004.
9. K. Lerman and A. Galstyan. Macroscopic analysis of adaptive task allocation in robots. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1951–1956, Las Vegas, Nevada, Oct 2003.
10. K. Lerman, A. Galstyan, A. Martinoli, and A.J. Ijspeert. A macroscopic analytical model of collaboration in distributed robotic systems. *Artificial Life*, 7(4):375–393, 2001.
11. A. Martinoli and K. Easton. Modeling swarm robotic systems. In B. Siciliano and P. Dario, editors, *Experimental Robotics VIII*, pages 297–306. Springer-Verlag, 2002.
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. L. Parker. Toward the automated synthesis of cooperative mobile robot teams. In *Proceedings of SPIE Mobile Robots XIII*, volume 3525, pages 82–93, Boston, MA, Nov 1998.
14. L. Sugawara and M. Sano. Cooperative acceleration of task performance: Foraging behavior of interacting multi-robot system. In *Physica D*, volume 100, pages 343–354, 1997.