OPTIMIZATION OF SWARM ROBOTIC SYSTEMS VIA MACROSCOPIC MODELS

Alcherio Martinoli
Collective Robotics Group, California Institute of Technology
1200 E. California Blvd., M/C 136-93, Pasadena, CA, 91125-9300
U.S.A.[*]
alcherio@caltech.edu

Kjerstin Easton
Collective Robotics Group, California Institute of Technology
1200 E. California Blvd., M/C 136-93, Pasadena, CA, 91125-9300
U.S.A.
easton@caltech.edu

Abstract In this paper, we propose a time-discrete, macroscopic model able to capture the dynamics of a robotic swarm system engaged in a collaborative manipulation task. The case study is concerned with pulling sticks out of the ground, an action that requires the collaboration of two robots to be successful. We will show that the model can deliver not only quantitatively correct predictions but also be a very useful tool for optimization. In particular, we will show how a mathematical analysis of a simplified model leads to counterintuitive results which can then be exploited in the full model or more detailed microscopic simulations to quantitatively assess the dynamic of the whole system. We conclude the paper with a discussion of strengths and limitations of the current model-based optimization method.

Keywords: Modeling, Optimization, Swarm Robotics, Swarm Intelligence

[*]This work is supported by the Caltech Center for Neuromorphic Systems Engineering as part of the NSF Engineering Research Center program under grant EEC-9402726.

1. Introduction

Swarm Intelligence (SI) is a computational and behavioral metaphor for solving distributed problems that takes its inspiration from biological examples mainly provided by social insects [1]. The abilities of such natural systems appear to transcend the abilities of the constituent individual agents. In most biological cases studied so far, the robust and capable high-level group behavior is mediated by nothing more than a small set of simple low-level interactions between individuals and between individuals and the environment.

One way to analyze and understand underlying common principles of swarm systems (both natural and artificial) is to capture their dynamics at more abstract levels. Modeling is a means for saving time, enabling generalization to different platforms, and estimating optimal system parameters, including control parameters and number of individuals. We present a detailed macroscopic model for a case study concerned with pulling sticks out of the ground with a team of simple, reactive autonomous robots. The macroscopic model is quantitatively correct, time-discrete, and characterized by zero-free parameters. It has been derived from its corresponding microscopic model presented in [3], and is consistent with the microscopic and macroscopic methodologies originally introduced in [6] and [7] respectively. Finally, we show how macroscopic models can be used as an optimization tool for swarm robotic systems. In particular, we discuss strengths and limitations of this model-based approach in comparison to one based on learning and microscopic models, a solution we recently investigated for the same case study in [5].

2. A Case Study in Distributed Manipulation

In the case study described in this paper, robots must pull sticks out of the ground, an action which, due to the length of the sticks, requires the collaboration of two robots to be successful. The metric measured to quantitatively investigate and model the effects of variations of system parameters is the collaboration rate among robots, i.e. the number of sticks successfully taken out of the ground over time.

2.1 Physical Set-Up and Embodied Simulations

The experiment is carried out in a circular arena of fixed radius with several 15 cm-long sticks, protruding 5 cm above holes in the arena floor. Groups of Khepera robots, equipped with gripper turrets, are used to pull sticks out of the ground.

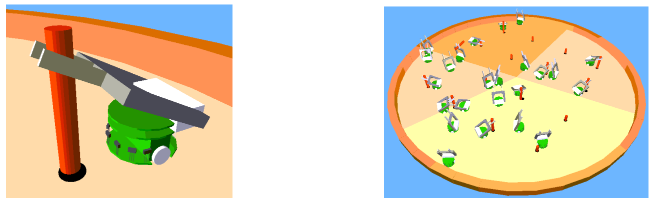

Figure 1. **Left:** Khepera robot with gripper in the embodied simulator. **Right:** Stick-pulling set-up in the embodied simulator (24 robots, 16 sticks).

Results of physical experiments (conducted with 4 sticks, team sizes from 2 to 6 robots, in an arena of radius 40 cm) are presented in [3, 7]. Embodied simulation, however, allows us to more systematically validate collaboration dynamics predicted by our models. Further experiments characterized by arenas with 4 to 16 sticks and team sizes between 2 and 24 robots were conducted in Webots, a 3D kinematic, sensor-based embodied simulator of Khepera robots (see Fig. 1). As shown in previous cases studies concerned with different tasks and different robotic platforms, this simulator is sufficiently faithful for the controllers to be transferred to real robots without changes [6, 3, 2].

2.2 The Robot's Controller

The behavior of a robot is determined by a simple hand-coded program which can be represented with a standard flow chart or a Finite State Machine (FSM), as depicted in Fig. 2 (left). The behavioral granularity shown in Fig. 2 is arbitrary and is chosen by the experimenter so that the FSM captures all the details of interest.

In addition to the default search behavior (moving in a straight line) and an obstacle avoidance behavior, the robot is endowed with a stick-gripping and -pulling procedure. The robot can determine from its arm elevation speed while pulling whether another robot is already gripping the same stick. While waiting for collaboration, another robot's attempt to lift the stick is similarly detected. If no other robot is holding the stick, we call such a grip a *grip1*. If another robot is already holding the stick, such a grip is called *grip2*. When a robot makes a grip1, it holds the stick raised half-way out of the ground and releases it when either the duration of the grip exceeds the *gripping time parameter* τ_g (a failed collaboration), or another robot comes to make a grip2 (a successful collaboration). Once the stick is released, the robot turns away, performs obstacle avoidance for a few seconds, then returns to the search procedure. When a robot makes a grip2, the robot making the grip1 will

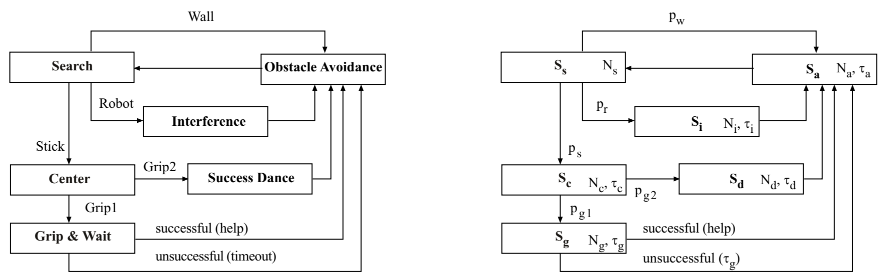

Figure 2. **Left:** FSM representing the robot controller. Transition between states are deterministically triggered by sensory measurements. **Right:** PFSM representing an agent in the microscopic model or the whole robotic team in the macroscopic model. The parameters characterizing probabilistic transitions and states are explained in the text.

release the stick, allowing its teammate to raise the stick completely. The robot making grip2 performs a short "success dance" (moving the arm up and down) to mark the successful collaboration, then releases the stick (which has to be replaced in the hole by the experimenter), performs obstacle avoidance for a few seconds, and resumes searching for sticks. Because of the way sticks are recognized (only by their thinness), a stick which is held by one robot can only be recognized when approached from the opposite side within a certain angle (approx. 125 degrees in the physical set-up). For the other angles of approach, both the stick and the robot are detected and the whole is taken for an obstacle. More details are reported in [3], but we note that the acceptable approach angle, expressed as a ratio R_g over the whole approaching perimeter, is an important parameter in the collaboration dynamics of the system and its role is explored in Section 4.

3. The Macroscopic Model

The central idea of both microscopic and macroscopic probabilistic modeling methodologies is to describe the experiment as a series of stochastic events with probabilities computed from the interactions' geometrical properties and systematic experiments with one or, due to the collaborative nature of the stick pulling experiment, two real robots or embodied agents [6, 3, 7]. Fig. 2 (right) shows a Probabilistic Finite State Machine (PFSM) whose state-to-state transitions depend on the interaction probabilities of a robot with another teammate and with the environment. While in microscopic models each robot is represented by its own PFSM, in macroscopic models, a single PFSM summarizes the whole robotic team, each of its states representing the average number of teammates in a particular state at a certain time step. In both types of models, the robots' PFSM(s) are coupled with the environment, a passive, shared resource whose modifications are generated by the parallel actions of the robots. In both cases, our metric, the collaboration rate, is computed from the modifications made to the environment.

3.1 Modeling Assumptions and Properties

The microscopic and macroscopic probabilistic modeling methodologies share several assumptions and properties [7]. The methodologies currently rely on the assumption that the coverage of the arena by the groups of robots is spatially uniform. Because we model the system as a PFSM, we must also assume that the robot's future state depends only on its present state and on how much time it has spent in that state. The details of computing the transition probabilities for the stick-pulling experiment can be found in [3]. Furthermore, consistent with previous publications [6, 3, 7], each iteration of our models corresponds to a time step T in real time equivalent to the time needed for a robot moving with a certain speed V_r and having a certain detection width W_r to cover the area of the smallest object in the arena (in our case, a stick).

3.2 Mathematical Description

The PFSM depicted in Fig. 2 (right) can also be translated into a set of difference equations (DE), one for each state, which mathematically represent the dynamics of the whole system at the macroscopic level. Conservation laws based on invariant numbers of sticks and robots during the experiment can be also exploited by replacing one of the DEs with a conservation equation.

All the different average durations measured with real robots or embodied agents can be then translated into numbers of iterations. It follows that the time needed for centering on and gripping a stick τ_c corresponds to T_c iterations, the duration of obstacle avoidance τ_a to T_a iterations, the duration of the interference between two robots τ_i to T_i iterations, the duration of the success dance τ_d to T_d iterations, and the gripping time parameter τ_g to T_g iterations. In this paper, the numerical values used in the microscopic and macroscopic models are exactly the same as those reported in [3], with the exception of the robot's detection width, W_r[1]. The following set of difference equations represent the macroscopic model of the stick pulling experiment:

$$N_s(k+1) = \tag{1}$$
$$N_s(k) - (p_w + p_R)N_s(k) - p_{g2}N_g(k)N_s(k) - p_{g1}[M_0 - N_g(k)]N_s(k)$$
$$+ p_w N_s(k - T_a) + p_R N_s(k - T_{ia}) + p_{g2}N_g(k - T_{ca})N_s(k - T_{ca})$$
$$+ p_{g2}N_g(k - T_{cda})N_s(k - T_{cda})$$
$$+ p_{g1}[M_0 - N_g(k - T_{cga})]N_s(k - T_{cga})\Gamma(k; T_{ga})$$

$$N_a(k+1) = \tag{2}$$
$$N_a(k) + p_w N_s(k) + p_R N_s(k - T_i) + p_{g2}N_g(k - T_c)N_s(k - T_c)$$
$$+ p_{g2}N_g(k - T_{cd})N_s(k - T_{cd})$$
$$+ p_{g1}[M_0 - N_g(k - T_{cg})]N_s(k - T_{cg})\Gamma(k; T_g) - p_w N_s(k - T_a)$$
$$- p_R N_s(k - T_{ia}) - p_{g2}N_g(k - T_{ca})N_s(k - T_{ca})$$
$$- p_{g2}N_g(k - T_{cda})N_s(k - T_{cda})$$
$$- p_{g1}[M_0 - N_g(k - T_{cga})]N_s(k - T_{cga})\Gamma(k; T_{ga})$$

$$N_i(k+1) = \tag{3}$$
$$N_i(k) + p_R N_s(k) - p_R N_s(k - T_i)$$

$$N_c(k+1) = \tag{4}$$
$$N_c(k) + p_{g2}N_g(k)N_s(k) + p_{g1}[M_0 - N_g(k)]N_s(k)$$
$$- p_{g2}N_g(k - T_c)N_s(k - T_c) - p_{g1}[M_0 - N_g(k - T_c)]N_s(k - T_c)$$

$$N_d(k+1) = \tag{5}$$
$$N_d(k) + p_{g2}N_g(k - T_c)N_s(k - T_c) - p_{g2}N_g(k - T_{cd})N_s(k - T_{cd})$$

$$N_g(k+1) = \tag{6}$$
$$N_0 - N_s(k+1) - N_a(k+1) - N_i(k+1) - N_c(k+1) - N_d(k+1)$$

$$\Gamma(k; T_g) = \prod_{j=k-T_g}^{k} [1 - p_{g2}N_s(j)] \tag{7}$$

$$\Gamma(k; T_{ga}) = \prod_{j=k-T_{ga}}^{k-T_a} [1 - p_{g2}N_s(j)] \tag{8}$$

The current iteration is represented by k, $k = 0, 1, 2, ..$ and for $k < 0$ the arena is empty (i.e. $N_s = N_a = N_i = N_c = N_d = N_g = 0$). M_0 is the number of sticks in the arena, N_0 is the total number of robots, and $T_{xyz} = T_x + T_y + T_z$. N_s represents the mean number of robots in the searching state, N_a those in the obstacle avoidance state, N_i those in the interference state, N_c those in the stick centering state, N_d those in the success dance state, and N_g those in the gripping state. Furthermore, p_w is the probability to find a wall, $p_s = p_{g1}$ that of finding a stick for grip1, $p_{g2} = R_g p_{g1}$ that of finding a stick for grip2, p_r that of finding a robot, and $p_R = p_r(N_0 - 1)$ that of finding any other robot in the arena. Equations (7) and (8) represent the fraction of robots that abandoned the grip1 state after the time spent in this state exceeded their gripping time parameter τ_g. This is equivalent to calculating the probability that no other robot came "to help" during a time interval of duration τ_g.

Finally, our team metric, the collaboration rate $\bar{C}_t$, can be computed from the number of successful collaborations C per time unit[2] over the maximal number of iterations T_e:

$$C(k) = p_{g2} N_s(k - T_{cd}) N_g(k - T_{cd}) \tag{9}$$

$$\bar{C}_t = \frac{\sum_{k=0}^{T_e} C(k)}{T_e} \tag{10}$$

3.3 Validation of the Model

Using an embodied simulator to validate our models' predictions, we present the results of several experiments. In particular, in this section we investigate the influence of the gripping time parameter and team size. Experiments using the embodied simulator have been repeated 10 times, and those using the microscopic model 100 times. All error bars in the plots correspond to the standard deviation among runs.

Macroscopic models base their collective performance forecast on one single run whose computation time is independent of the number of teammates but with small group sizes their predictions is only qualitatively correct (see Fig. 3, left) because they base their prediction on the validity of the law of large numbers. Fig. 3 (right) shows that, without changing the density of sticks in the arena (area and number of sticks are also multiplied), it suffices to multiply the number of robots by four in order to obtain quantitative agreement between microscopic and macroscopic models without changing any implementation details.

Observing Fig. 3, our intuition suggests than there is a different relation between the collaboration rate and the gripping time parameter depending on the ratio between number of robots and number of sticks. When there are more robots than sticks, the collaboration rate increases monotonically with the gripping time parameter and eventually saturates in a plateau corresponding to the optimal collaboration rate. In other words, under these conditions, it is a good strategy for a robot gripping a stick to wait a very long time for another robot to give a hand, because there will always be at least one "free" robot available. By contrast, when there are fewer robots than sticks, waiting a very long time becomes a bad strategy, as the few robots lose time holding different

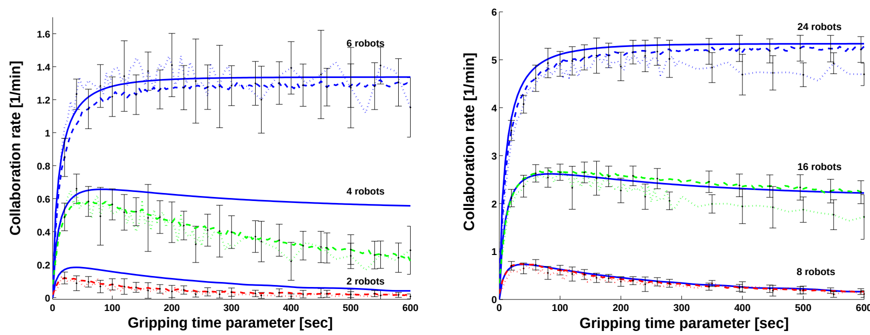

Figure 3. **Left:** Collaboration rate as a function of the gripping time parameter for several group sizes. Results gathered for using embodied simulations (dotted lines), the microscopic model (dashed lines), and the macroscopic model (solid lines) in the 4-stick, 40 cm-radius arena and **Right:** in the 16-stick, 80 cm-radius arena.

sticks, while no other robots are available to collaborate. For instance (an extreme case), an infinite gripping time parameter would lead to a null collaboration rate with all robots eventually holding a different stick permanently. In Section 4, we will demonstrate via the steady-state macroscopic analysis that this intuition is not generally correct since it is only true for a particular range of values of the R_g parameter. The specific value corresponding to the physical set-up investigated happened to be precisely in this range.

4. Model Analysis and System Optimization

In this section, we will show that macroscopic models not only allow a researcher to predict the behavior of the whole system but represent also a useful tool for optimization. For the sake of simplicity and availability of space in this paper, we will focus our analysis on a simplified model which is, however, good enough to qualitatively capture the collaboration dynamics of the system. In Subsection 4.3, with the help of Eqs. (1–10), we use the insight gained with the simple model to investigate the behavior of the full-system model quantitatively.

4.1 Steady-State Analysis of a Simplified Model

As described in Subsection 3.2, several states of the macroscopic models (N_a, N_i, N_c, N_d) can be thought of simple delay states: in reality, robots assume these modes for a short period of time, for instance the time needed to perform a success dance or avoiding an obstacle. These operations usually last a few seconds and are therefore much shorter than most of the values of the gripping time parameter considered in this paper. As a consequence, for sake of simplicity, we can neglect these delays and reduce the model to two states: searching and gripping. Mathematically, this can be obtained by setting $T_a = T_i = T_c = T_d = 0$ in Eqs. (1–10) and reformulating the model as follows:

$$\begin{aligned} N_s(k+1) &= N_s(k) - p_{g1}[M_0 - N_g(k)]N_s(k) + p_{g2}N_g(k)N_s(k) \\ &\quad + p_{g1}[M_0 - N_g(k-T_g)]N_s(k-T_g)\Gamma(k;T_g) \quad (11) \\ N_g(k+1) &= N_0 - N_s(k+1) \quad (12) \end{aligned}$$

$\Gamma(k;T_g)$ is described by Eq. (7), the collaboration rate by Eq. (10), while Eq. (9) becomes:

$$C(k) = p_{g2}N_s(k)N_g(k) \quad (13)$$

4.2 The Role of the Collaboration Parameter R_g

Much as Lerman et al. have proposed in [4] for a time-continuous version of the simplified model, we now perform an analysis of the system in its steady state.

By setting $N_s(i) = N_s^*$ and $N_g(i) = N_g^*$ for all i between $k - T_g$ and $k+1$ in Eq. (7), Eq. (11), and Eq. (12), we obtain:

$$\begin{aligned} 0 &= -p_{g1}(M_0 - N_g^*)N_s^* + p_{g2}N_g^*N_s^* + \quad (14) \\ &\quad p_{g1}(M_0 - N_g^*)N_s^*(1 - p_{g2}N_s^*)^{T_g} \\ N_g^* &= N_0 - N_s^* \quad (15) \\ C^* &= p_{g2}N_s^*N_g^* \quad (16) \end{aligned}$$

First of all, we would like to know when the number of collaborations is maximized as a function of the robots in the searching state (or in the gripping state, respectively). In order to answer this question, we insert Eq. (15) in Eq. (16), perform a partial derivative over N_s^*, and set the result equal to zero. C^* is maximal when $N_s^* = N_0/2$.

By inserting this result, $p_{g2} = R_g p_{g1}$, and Eq. (15) in Eq. (14), we obtain following transcendental equation:

$$0 = -(M_0 - \frac{N_0}{2}) + R_g\frac{N_0}{2} + (M_0 - \frac{N_0}{2})(1 - p_{g1}R_g\frac{N_0}{2})^{T_g^{opt}} \quad (17)$$

Introducing $\beta = N_0/M_0$ and solving the equation over T_g^{opt}, we obtain:

$$T_g^{opt} = \frac{1}{\ln(1 - p_{g1}R_g\frac{N_0}{2})}\ln\frac{1 - \frac{\beta}{2}(1+R_g)}{1 - \frac{\beta}{2}} \quad (18)$$

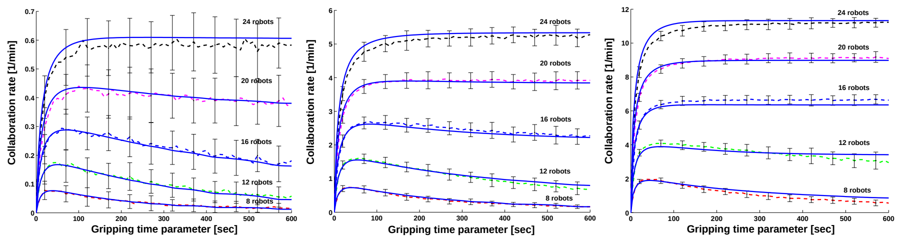

Figure 4. Collaboration rate as a function of R_g (from left to right $R_g = 0.035; 0.35; 1.0$). Results gathered using the microscopic model (dashed lines) and the macroscopic model (solid lines) in the 16-stick, 80 cm-radius arena.

Equation (18) tells us that an optimal T_g exists if all the arguments of the logarithms are greater than zero. While this condition is always verified for the first logarithm ($N_0 < 2/p_{g1} \approx 312$ with $R_g = 1$, our worst case), the argument of the second logarithm depends on β and R_g. It can be demonstarted that an optimum exists if and only if[3]:

$$\beta < \beta_c = \frac{2}{1 + R_g} \tag{19}$$

Equation (19) tells us that the bifurcation of the system (optimal T_g vs. no optimum at all) is a function of the collaboration parameter R_g. For instance, this means that if the collaboration is very difficult (R_g is very small), there could be situations where β_c is quite greater than one. This in turn means that, even in the case we have an enclosed arena and more robots than sticks, an optimal gripping time parameter will still exist. In other words, when it is difficult to collaborate, in order to enhance the number of collaborations, it is worth abandoning the sticks after a while and increasing the critical mass working in another side of the arena. Although the precise team size at which the bifurcation happens in the real system cannot correctly be computed with Eq. (19), this equation allows us to better situate intuitive considerations such as those presented in Subsection 3.3.

4.3 Validation using Full-System Models

We assessed the validity of the assertion presented in the previous section using three different values of R_g: 0.035, 0.35 (corresponding to the physical set-up), and 1.0. Figure 4 shows the results obtained using the full-system microscopic and macroscopic models. We notice the correctness of Eq. (19), for instance, by observing the collaboration rate of a team of 20 robots in the left plot of Fig. 4: although there are more robots than sticks (16) a τ_g of about 100 s represents the optimal gripping time parameter for these experimental conditions.

5. Conclusion

In this paper, we have presented a time-discrete, macroscopic model able to deliver quantitatively correct predictions about the collaboration dynamics of a specific distributed manipulation experiment concerned with pulling sticks out of the ground. Furthermore, we have shown how this type of model can be very useful as tool to estimate optimal parameters of the robotic system as a function of the environmental or task conditions. While a model-based approach shows several advantages, the major one being that if the model is mathematically tractable, it allows researchers to draw general conclusions very quickly, also shows several limitations. First, although modeling collective robotic system has recently received more attention than in the past, a lot of work must still be done to develop methodologies which describe the system at an abstract level (e.g. behavioral level) but are, at the same time, soundly anchored to the real physical system. The methodology we have developed in recent years is a good attempt in this direction but alternative approaches should be explored. Second, from an optimization point of view, a homogeneous solution may not necessarily achieve the best performance and we would like to have tools to explore heterogeneous solutions as well. The use of macroscopic models forces us to introduce a new set of DEs for each new type of agent involved in the system with the risk that, if the optimal solution requires high differentiation (in the worst case each individual may differ from the other), the quantitative correctness of the prediction is no longer insured, as we have shown in Subsection 3.3. To explore heterogeneous solutions, in particular those with non pre-established heterogeneity, microscopic models combined with machine-learning algorithms (see for instance [5]) appear to be a more efficient solution.

In conclusion, we strongly believe that the combination of model-based and machine-learning-based approaches will be a winning strategy for understanding, designing, and optimizing future swarm robotic systems able to solve real world tasks.

Acknowledgments

The authors wish to thank William Agassounon and Ling Li for helpful comments on this work.

Notes

1. Systematic experiments in Webots indicate a more accurate value of W_r to be 12.8 cm, i.e., twice the center-to-center distance at which a robot can detect a seed. This is consistent with the original microscopic methodology presented in [6].

2. A collaboration is considered successfully terminated when the second robot is at the end of its success dance.

3. It can be easily demonstrated that by introducing $\beta = N_0/M_0$ in Eq. (17), solving over β and comparing with Eq. (19), N_s in Eq. (17) will be always greater than $N_0/2$, i.e. not an optimal value.

References

[1] Bonabeau, E., Dorigo, M., and Theraulaz, G. (1999) Swarm Intelligence: From Natural to Artificial Systems. SFI Studies in the Science of Complexity, Oxford University Press, New York, NY.

[2] Hayes, A. T., Martinoli, A., and Goodman, R. M. (2002) Distributed Odor Source Localization. Special Issue on Artificial Olfaction, Nagle, H. T., Gardner, J. W., and Persaud, K. (Eds.), IEEE Sensors **2(3)**, 260–271.

[3] Ijspeert, A. J., Martinoli, A., Billard, A., and Gambardella, L. M. (2001) Collaboration through the Exploitation of Local Interactions in Autonomous Collective Robotics: The Stick Pulling Experiment. Autonomous Robots **11(2)**, 149–171.

[4] Lerman, K., Galstyan, A., Martinoli, A., Ijspeert, A. J. (2001) A Macroscopic Analytical Model of Collaboration in Distributed Robotic Systems. Artificial Life **7(4)**, 375–393.

[5] Li, L., Martinoli, A., and Abu-Mostafa, Y. (2002) Emergent Specialization in Swarm Systems. In Yin, H. et al.(Eds.), Proc. of the Third Int. Conf. on Intelligent Data Engineering and Automated Learning, August 2002, Manchester, UK. Lecture Notes in Computer Science (2002), Vol. 2412, 261–266.

[6] Martinoli, A., Ijspeert, A. J., and Gambardella, L. M. (1999) A Probabilistic Model for Understanding and Comparing Collective Aggregation Mechanisms. In: Floreano, D., Mondada, F., and Nicoud, J.-D. (Eds.), Proc. of the Fifth European Conf. on Artificial Life, September, Lausanne, Switzerland. Lectures Notes in Computer Science (1999), 575–584.

[7] Martinoli, A., and Easton, K. (2002) Modeling Swarm Robotic Systems. In Siciliano, B. and Dario, P. (Eds.), Proc. of the Eight Int. Symp. on Experimental Robotics, July, 2002, Sant'Angelo d'Ischia, Italy. Springer Tracts in Advanced Robotics (2003), 297–306.