Modeling Swarm Robotic Systems

Alcherio Martinoli and Kjerstin Easton

California Institute of Technology, M/C 136-93, 1200 E. California Blvd.
Pasadena, CA 91125, U.S.A.
alcherio,easton@caltech.edu, http://www.coro.caltech.edu

Abstract. In this paper, we discuss strengths and limitations of different abstraction levels for distributed robotics experiments. We support the discussion with a concrete case study which has been investigated at four different levels: real robots, embodied simulations, microscopic modeling, and macroscopic modeling. Both modeling methodologies presented represent the collective dynamics of the experiment as a set of stochastic events based on simple geometrical considerations and systematic tests with one or two real robots instead of computing trajectories and sensory information like an embodied simulator would do. The case study we describe is concerned with pulling sticks out of the ground - an action which requires the collaboration of two robots to be successful. Experiments were carried out with teams consisting of two to 24 individuals endowed with simple reactive controllers. In addition to showing that models can deliver both qualitatively and quantitatively correct predictions in time lapses that are three or four orders of magnitude smaller than those required by embodied simulations, we discuss differences, assumptions, and subtle numerical effects of the current modeling methodologies.

1 Introduction

Swarm Intelligence (SI) is a new computational and behavioral metaphor for solving distributed problems [1]; it is based on the principles underlying the behavior of natural systems consisting of many agents, such as ant colonies and bird flocks. The abilities of such systems appear to transcend the abilities of the constituent individual agents; in all the biological cases studied so far, the emergence of high-level control has been found to be mediated by nothing more than a small set of simple low-level interactions among individuals, and between individuals and the environment. The application of SI principles to multiple-robot systems leads to three main advantages: first, scalability of the control architecture, from a few to thousands of units; second, self-organization, as units can be dynamically added, removed, or reallocated to different tasks without explicit reorganization; and third, increased system robustness, not only through unit redundancy but also through the design of simple units.

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. There has been relatively little prior work in modeling of multi-robot systems with the exception of our research [2–5] and that of Sugawara and coworkers [6], which explored macroscopic models for a foraging task. This paper combines the strengths of our previous *quantitative* microscopic and *qualitative* macroscopic models in order to move towards a macroscopic *quantitative* modeling methodology. In addition, we discuss how the new macroscopic methodology relates to the microscopic one, in particular by describing how to build up a macroscopic model starting from a state diagram of the robot controller used at the microscopic level. Finally, we discuss some of the problems and limitations that affect either one or both types of modeling, in particular in comparison to other popular simulation tools such as sensor-based, embodied simulators. We support our discussion with a concrete case study concerned with pulling sticks out of the ground with a team of simple, reactive, autonomous robots.

2   A Case Study in Distributed Manipulation

The case study described in this paper, first introduced by Martinoli and Mondada [7], has been systematically investigated using embodied simulations and a microscopic model [3], and, more recently, using a qualitative macroscopic model [5]. In this case study, the task involves pulling sticks out of the ground, an action which requires the collaboration of two robots to be successful. All the experiments have been carried out to quantitatively investigate and model the effects of variations of robot controller implementations and number of robots on group behavior, in particular, on 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 (40 cm of radius) delimited by a white wall. Four holes situated at the corners of a square with 30 cm edges, hold white sticks (15 cm long, diameter of 1.6 cm) which, in their lowest position, protrude 5 cm above the ground (see Fig. 1, left). Groups of two to six Khepera robots, equipped with gripper turrets, are used to pull the sticks out of the ground. Because of their thinness, the sticks can be distinguished from the wall and from other robots using the Khepera's six frontal IR proximity sensors. Because the sticks are too long to be pulled from the ground by a single robot's lifting motion, collaboration between two robots is required. After a successful collaboration, the stick taken out of the ground is released by the robot, and replaced in its hole by the experimenter.

In order to more systematically investigate the collaboration dynamics, we also implemented the experiment in Webots, a 3D kinematic, sensor-based simulator of Khepera robots (see Fig. 1, right). Teams of two to 24 robots were
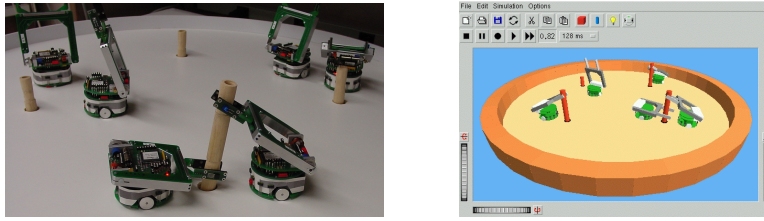
Fig. 1. Left: Overview of the physical set-up for the stick-pulling experiments. **Right:** Corresponding set-up in the embodied simulator

simulated using Webots. The simulator computes trajectories and sensory input of the robots in an arena corresponding to the physical set-up. The simulation is sufficiently faithful for the controllers to be transferred to real robots without changes and for the robot behaviors to be very similar to those of the real robots, as shown in several previous papers [2,3]. The mean speed ratio for this experiment with five robots between Webots and real time is about 18 on a Pentium III, 900 MHz machine.

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 on the left hand of Fig. 2. 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.
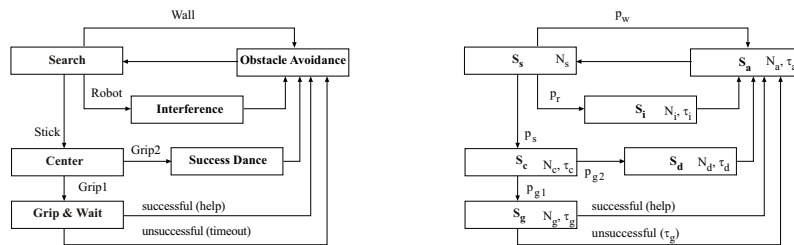

Fig. 2. Left: FSM representing the controller of a real robot. Transition between states are deterministically triggered by sensory measurements. **Right:** PFSM representing an agent in the microscopic model or the whole robotic team in the macroscopic model. The parameters characterizing probabilistic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

In addition to the default search behavior (wandering in a straight line) and an obstacle avoidance behavior, the robot is endowed with an original stick gripping and pulling procedure. During pulling, the robot can determine whether another robot is already gripping the same stick by measuring the speed of elevation of the gripper arm. If the elevation is fast, the robot assumes that the stick is free (no other robot holding it) and we call such a grip a *grip1*. If the elevation is slow, the robot assumes that another robot is already holding that stick and therefore "braking" the elevation. Such a grip is called *grip2*. When a robot makes a grip1, it holds the stick half out of the ground and releases it when either the duration of the grip exceeds a *gripping time parameter* (which is then considered as a failed collaboration) or another robot comes to make a grip2. The robot can detect when another robot is making a grip2 because the force exerted by that robot on the stick leads to a slight elevation of its arm's position compared to the arm's programmed position. Once the stick is released, the robot turns away, performs obstacle avoidance for a few seconds, and returns to the search procedure. When a robot makes a grip2, the pulling of the stick is temporarily braked until the robot making the grip1 releases the stick. The arm eventually reaches the programmed position, which corresponds to a complete lift of the stick out of the ground. To mark the successful collaboration, a small "success dance" (moving the arm up and down) is performed. Like at the end of grip1, the robot then releases the stick (which has to be replaced in the hole by the experimenter) and resumes looking for sticks.

Note that, because of the way sticks are recognized (i.e. 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). For the other angles of approach, both the stick and the robot are detected and the whole is therefore taken for an obstacle [3].

3   Microscopic and Macroscopic Models

The central idea of both 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. Fig. 2 (right) shows a Probabilistic Finite State Machine (PFSM) or a Markov chain whose state-to-state transitions depends on the interaction probabilities of a robot with another teammate and with the environment. As first introduced in [5], 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 then coupled with the environment. The environment can be considered as a passive, shared resource whose

modifications are generated by the parallel actions of the robots. In order to compute the arbitrary metric we are interested in (collaboration rate) we keep track of the modifications of the environment. The mean speed ratio for this experiment with five robots between microscopic model (implemented in C) and Webots simulations is about 25'000 on a Pentium III, 900 MHz machine. That of the macroscopic model (currently implemented in Matlab) on the same machine is instead of $4 \times N \times R$, N being the total number of robots in the team and R the total number of runs for obtaining the mean performance.

3.1   Common Modeling Assumptions and Properties

Both modeling methodologies share common assumptions and properties, each of them briefly described in turn.

1. **Spatial Uniformity** - Both methodologies currently rely on the assumption that the coverage of the arena by the groups of robots is uniform. Robots' trajectories therefore are not considered in the models. We also assume that the absolute position of a given object to manipulate in the arena does not play a role: the object will have the same probability to be manipulated if it is placed in the center or in the periphery of the arena.

2. **Markov Properties** - We assume that the robot's future state depends only on its present state and on how much time it has spent in that state. This assumption is correct for a reactive robot controller extended with a time-out or following a predetermined sequence of actions (e.g. gripping a stick, dancing) that lasts a certain amount of time. The robots (and the environment) in the stick-pulling case study clearly obey this Markov property if we assume that trajectories (i.e. position and heading history) can be neglected for computing the desired metric.

3. **Transition Probabilities** - Consistent with previous publications [2–4], we compute the transition probabilities from a state to another based on simple geometrical considerations about the interaction (e.g. detection areas, approaching perimeters). The numerical values used for these geometrical parameters are measured in systematic tests with one or two real robots, as mentioned above. At each iteration or time step, the probability that a robot in the search mode will encounter a wall, a stick, or another robot is determined by their corresponding detection area divided by the whole arena area A_a. For instance, the probability of finding a stick can be computed as $p_s = A_s/A_a$, A_s being the detection area of a stick. Similarly, $p_w = A_w/A_a$ and $p_r = A_r/A_a$ represent the probability of encountering a wall and another robot respectively. Additionally, since robots can perform a grip1 from any angle of approach, $p_{g1} = p_s$. On the other hand, since a stick available for grip2 can only be approached from a certain angle, an additional factor R_g based on the approaching-whole perimeter ratio has to be introduced. The probability of a grip2 event is therefore $p_{g2} = R_g p_{g1}$. See [3] for more details.

4. **Time Discretization** - Consistent with previous publications [2–4], each iteration of our models correspond to a time step of a finite duration in real time. The duration of a time step is 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). Equation (1) shows how to compute the duration T of one time step in both modeling methodologies:

$$T = \frac{A_s}{V_r W_r} \tag{1}$$

All the different average durations measured with one or two real robots, such as the time for centering on and gripping a stick τ_c, the duration of obstacle avoidance τ_a, the duration of the interference between two robots τ_i, the duration of the success dance τ_d, and the gripping time parameter τ_g can also be discretized using the time step T. In this paper, the numerical values used in the microscopic and macroscopic models are exactly the same as those reported in [3].

4   Mathematical Description of the Macroscopic Model

The Markov chain depicted in Fig. 2 (right) can also be translated into a set of difference equations (DE), one for each state, which mathematically represent the dynamic of the whole system at the macroscopic level. An equation stating the conservation of the number of robots during the experiment can be also exploited to replace one of the DEs. For instance, the mean number of robots in search state N_s can be expressed as follows:

$$
\begin{aligned}
N_s(kT + 1) = {}& N_s(kT) - p_{g1}[M_0 - N_g(kT)]N_s(kT) \\
& - p_{g2}N_g(kT)N_s(kT) - (p_w + p_R)N_s(kT) \\
& + p_{g2}\Theta(kT - \tau_{cda})N_s(kT - \tau_{cda})N_g(kT - \tau_{cda}) \\
& + p_{g2}\Theta(kT - \tau_{ca})N_s(kT - \tau_{ca})N_g(kT - \tau_{ca}) \\
& + p_{g1}\Theta(kT - \tau_{cga})[M_0 - N_g(kT - \tau_{cga})]N_s(kT - \tau_{cga})\Gamma \\
& + p_w\Theta(kT - \tau_a)N_s(kT - \tau_a) + p_R\Theta(kT - \tau_{ia})N_s(kT - \tau_{ia})
\end{aligned} \tag{2}
$$

$$\Gamma = \prod_{j=k-\tau_{ga}/T}^{k-\tau_a/T}[1 - p_{g2}N_s(jT)]\Theta(kT - \tau_{ga}/T) \tag{3}$$

$$\Theta(kT - \tau) = \begin{cases} 0 \text{ if } kT < \tau \\ 1 \text{ if } kT \geq \tau \end{cases} \tag{4}$$

where $k = 0..n$, n being the maximal number of iterations, M_0 the number of sticks in the arena, N_0 the total number of robots, N_g the mean number of

robots in gripping state, $p_R = (N_0-1)p_r$, $\tau_{cda} = \tau_c+\tau_d+\tau_a$, $\tau_{cga} = \tau_c+\tau_g+\tau_a$, $\tau_{ca} = \tau_c + \tau_a$, $\tau_{ia} = \tau_i + \tau_a$, and $\tau_{cg} = \tau_c + \tau_g$. Similar expressions can be derived for the other states (N_a, N_c, N_d, N_g, N_i) of the PFSM depicted in Fig. 2 (right).

In other words, (2) tells us that the mean number of robots in search state at any time is decreased by the robots which transition to a gripping state (grip1 and grip2) and those which start avoiding a wall or a teammate; N_s is increased by the robots which come back from a successful collaboration either as first or second robot, those which come back from an unsuccessful collaboration, and those which finish their wall or robot avoidance maneuver. Equation (3) represents the fraction of robots that abandoned the grip1 state after the time spent in this state exceeded their gripping time parameter τ_g. As explained more extensively in [5], this is equivalent to calculating the probability that no other robot came "to help" during the time interval $[kT - \tau_{cg}, kT]$. Step functions Θ (4) are instead introduced in order to ensure that delayed factors do not play a role in the DE set when their corresponding time step would assume negative values. Finally, our team metric, the collaboration rate $\bar{C}_t$, can be computed from the number of successful collaborations C per time unit over the observation time T_e:

$$C(kT) = p_{g2}\Theta(kT - \tau_{cda})N_s(kT - \tau_{cda})N_g(kT - \tau_{cda}) \tag{5}$$

$$\bar{C}_t = \frac{\sum_{k=0}^{T_e/T} C(kT)}{T_e} \tag{6}$$

5   Results

We present the results of several experiments implemented at four different levels: real robots, embodied simulations, microscopic, and macroscopic modeling. In particular, we investigate the influence of the gripping time parameter and team size. We first investigate the standard experimental conditions (two to six robots, four sticks distributed symmetrically in an arena of 40 cm in radius) and we then discuss the influence of several experimental parameters such as the stick distribution, the robot density, and the absolute team size on the quality of the model predictions. All the real robot experiments lasted about 20 minutes (duration of the on-board batteries) while those in simulation 30 minutes (simulated time). Experiments using real robots have been repeated three times, those using the embodied simulator ten times, and those using the microscopic model 100 times. At the macroscopic level, of course, one run suffices, since only the mean team performance can be predicted. Although the simulation of microscopic models scales at least linearly with the team size and requires multiple runs in order to deliver statistically accurate predictions, with small team sizes, their results are quantitatively accurate (see Fig. 3, left). 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 are only qualitatively correct (see Fig. 3, right).
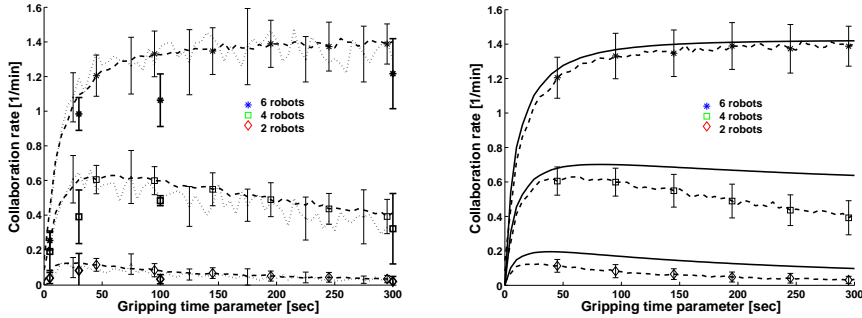

Fig. 3. Left: Collaboration rate as a function of the gripping time parameter for group sizes of two, four, and six robots. Results gathered using real robots ($\tau_g = [5, 30, 100, 300]$ seconds), embodied simulations ($\tau_g = [0 : 5 : 300]$ seconds), and the microscopic model ($\tau_g = [0 : 5 : 300]$ seconds) are overlapped. **Right:** Comparison of the prediction of the microscopic and the macroscopic model for the same experimental conditions

5.1   Numerical Effects due to the Stick Distribution

Both modeling methodologies assume a uniform distribution of objects on the arena. As long there are no overlapping detection areas between the objects (in this case walls, sticks), this assumption is correct. In order to verify this, we ran several experiments characterized by different stick distributions (see Fig. 4) using the embodied simulator. In each of these cases, the predictions of both models were as good as those shown in Fig. 3. The current modeling methodologies could be also easily adapted in order to take into account overlapped detection areas. This should be particularly easy for non-mobile objects. However, subtle effects due to robot clustering and mutual influence in search and manipulation activities could arise in dense scenarios and, as we will show in the next section, these effects are more difficult to incorporate in the models.

5.2   Numerical Effects due to the Team Size and Robot Density

Macroscopic models base their prediction on the validity of the law of large numbers. Fig. 5 (left) shows that, without changing the density of objects 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 a perfect quantitative
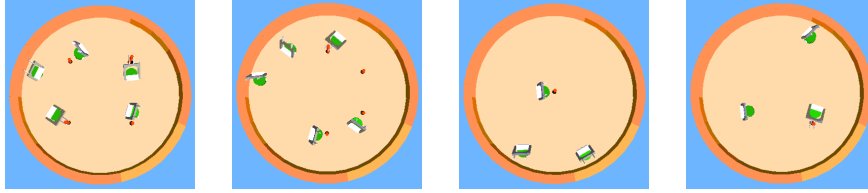
Fig. 4. Four examples of implemented stick distributions (four sticks, centered and eccentric distribution; one stick, centered and eccentric distribution)
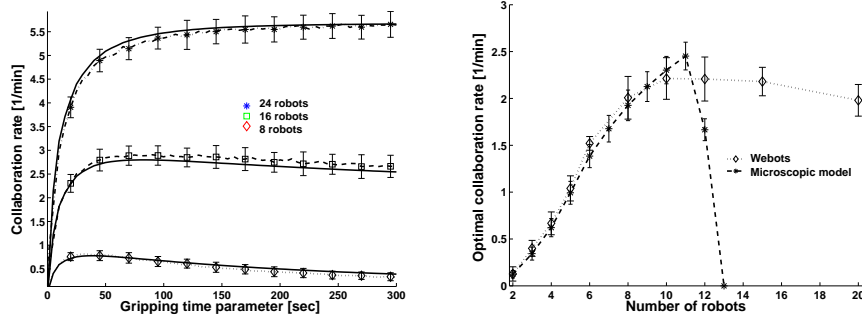

Fig. 5. Left: Results of microscopic and macroscopic models for 8, 16, and 24 robots, 16 sticks, and an arena of 80 cm in radius. **Right:** Comparison of the predictions obtained using embodied simulations and the microscopic model on an overcrowded arena (up to 20 robots and arena of 40 cm in radius). For each group size we plotted the collaboration rate achieved after optimization of τ_g (systematic search)

agreement between microscopic and macroscopic models without changing any implementation details.

Finally, both modeling methodologies achieve quantitatively correct predictions based on the assumption that robots are homogeneously distributed in the arena. As soon as this assumption is no longer valid, such as in an overcrowded scenario, the current methodologies reach their limitations and predictions are no longer correct. Fig. 5 (right) illustrates this effect for the microscopic model. The plot shows a clear discrepancy between the model's and embodied simulation's results for group sizes greater than 12 robots. In fact, the model predicts that with 13 robots in an arena of 40 cm radius, the detection area covered by the robots is equivalent to the arena area and therefore each robot will continuously avoid its teammates without any chance of collaborating in stick pulling. The reality is instead that robots form clusters (overlapped detection areas) and free up space around sticks which can, in turn, be exploited by other robots for successful collaborations.

6   Conclusion

In this paper, we have presented microscopic and macroscopic methodologies based on Markov models for predicting the dynamics of a swarm robotic system. We have supported the discussion with a specific case study in distributed manipulation: the stick pulling experiment. The methodologies are characterized by zero-free parameters, are three or four order of magnitude faster than embodied simulations, and achieve fairly accurate predictions if the density of robots is sufficiently low ("gaseous phase") and boundary conditions (e.g. sticks very close to the wall) do not play a major role in the metric considered. Future work will include a rigorous investigation of the role of time discretization in the current methodologies as well as possible extensions to further case studies characterized by different tasks, environmental features, and robots' controllers.

7   Acknowledgments

References

1. Bonabeau, E., Dorigo, M., Theraulaz, G.(1999) Swarm Intelligence: From Natural to Artificial Systems. SFI Studies in the Science of Complexity, Oxford University Press, New York, NY
2. Martinoli, A., Ijspeert, A. J., Gambardella, L. M. (1999) A Probabilistic Model for Understanding and Comparing Collective Aggregation Mechanisms. In: Floreano, D., Mondada, F., Nicoud, J.-D. (Eds.) Proc. of the Fifth European Conf. on Artificial Life, September, Lausanne, Switzerland, Lectures Notes in Computer Science, 575–584
3. Ijspeert, A. J., Martinoli, A., Billard, A., 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. Agassounon, W., Martinoli, A., Goodman, R. M. (2001) A Scalable, Distributed Algorithm for Allocating Workers in Embedded Systems. In: Proc. of the IEEE Conf. on System, Man and Cybernetics SMC-01, October 2001, Tucson, AR, USA, 3367–3373
5. 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
6. Sugawara, K., Sano, M. (1997) Cooperative Acceleration of Task Performance: Foraging Behavior of Interacting Multi-Robots System. Physica D **100**, 343–354
7. Martinoli, A., Mondada, F. (1995) Collective and Cooperative Group Behaviours: Biologically Inspired Experiments in Robotics. In: Khatib, O., Salisbury, J. K. (Eds.) Proc. of the Fourth Int. Symp. on Experimental Robotics, Stanford, June, Lecture Notes in Control and Information Sciences, 3–10