

A Review of Probabilistic Macroscopic Models for Swarm Robotic Systems

Kristina Lerman¹, Alcherio Martinoli², and Aram Galstyan¹

¹ USC Information Sciences Institute,
Marina del Rey CA 90292, USA,
lermand@isi.edu,

WWW home page: <http://www.isi.edu/~lerman/>

² Swarm-Intelligent Systems Group, Nonlinear Systems Laboratory, EPFL
CH-1015 Lausanne, Switzerland

Abstract. In this paper, we review methods used for macroscopic modeling and analyzing collective behavior of swarm robotic systems. Although the behavior of an individual robot in a swarm is often characterized by an important stochastic component, the collective behavior of swarms is statistically predictable and has often a simple probabilistic description. Indeed, we show that a class of mathematical models that describe the dynamics of collective behavior can be generated using the individual robot controller as modeling blueprint. We illustrate the macroscopic modelling methods with the help of a few sample results gathered in distributed manipulation experiments (collaborative stick pulling, foraging, aggregation). We compare the models' predictions to results of probabilistic numeric and sensor-based simulations as well as experiments with real robots. Depending on the assumptions, the metric used, and the complexity of the models, we show that it is possible to achieve quantitatively correct predictions.

1 Vision

Swarm Robotics is an emerging area in collective robotics which uses a fully distributed control paradigm and relatively simple robots to achieve coordinated behavior at the group level. Swarm robotic systems are *self-organizing*, meaning that constructive *collective* (or macroscopic) behavior emerges from *individual* (or microscopic) decisions robots make. These decisions are based on purely local information that comes from other robots as well as the environment. Swarm Robotics takes its inspiration from examples of collective behavior exhibited by biological systems, such as social insects [3], and the swarming, flocking, herding, and shoaling phenomena in vertebrates. In all these systems, the abilities of the collective appear to transcend the abilities of the constituent individuals.

The main advantages of the application of the swarm approach to the control of a group of robots are: (i) scalability: the control architecture can be kept exactly the same from a few units to thousands of units; (ii) flexibility: units can be dynamically added or removed, they can be given the ability to reallocate

and redistribute themselves in a self-organized way; (iii) robustness: the resulting collective system is robust not only through unit redundancy but also through unit simplicity and an appropriate balance between exploitative and exploratory behavior.

The main difficulty in designing swarm robotic systems with desirable self-organized behavior is understanding the effect individual robot characteristics have on the collective behavior. In the past, few analysis tools have been available to researchers. Experiments with physical robots are very costly and time consuming, and systematically studying group behavior is often impractical. Simulations, such as with embodied simulators [5, 13], attempt to realistically model the environment, the robots' imperfect sensing of and interactions with it. Though simulations are much faster and much more reliable than experiments, their results are not easily generalizable. Exhaustive scan of the design parameter space is often required to reach any conclusion. Moreover, simulations do not scale well with the system size — unless computation is performed in parallel, the greater the number of agents, the longer it takes to obtain results.

Macroscopic modeling and mathematical analysis offer an alternative to experiments and simulations. Using mathematical analysis we can quickly and efficiently study swarm robotic systems, predict their long term behavior, gain insight into system design: *e.g.*, how individual robot characteristics affect group behavior. Additionally, mathematical analysis may be used to select parameters that optimize group performance, prevent instabilities, *etc.* Finally, results of analysis can be used as feedback to guide performance-enhancing modifications of the robot controller.

In this paper we survey existing work on modeling collective behavior of robot swarms with macroscopic models. The robots themselves in these systems are simple, usually using reactive control: robots decide about future actions based solely on input from sensors (including communication with other robots) and the action they are currently executing. They do not rely on abstract representation, planning, or higher order reasoning functions. Such robots can be represented as stochastic Markov processes. An equation, known as the Rate Equation, describes the dynamics of their collective behavior. The Rate Equation formalism can be derived from theory of stochastic processes [8], although in practice, the equations are usually phenomenological and can be easily written down by considering details of the individual robot controller. The Rate Equation approach has been applied to study several distributed robot systems [14, 10, 7, 11, 1]. Below we review the elements of the mathematical formalism and illustrate with a few sample results from the robotics domain.

2 Methods for Modeling Swarm Robotic Systems

Models can generally be broken into two classes: *microscopic* and *macroscopic*. Microscopic descriptions treat the robot as the fundamental unit of the model. These models describe the robot's interactions with other robots and the environment. Solving or simulating a system composed of many such agents gives

researchers an understanding of the global behavior of the system. Examples of such microscopic models are reported in [12, 6]; they have been used to study collective behavior of a swarm of robots engaged in object aggregation and collaborative pulling. Rather than compute the exact trajectories and sensory information of individual robots, the robot’s interactions with other robots and the environment are modeled as a series of stochastic events, with probabilities determined by simple geometric considerations and systematic experiments with one or two real robots. Running several series of stochastic events in parallel, one for each robot, allows researchers to study the collective behavior of the swarm.

A macroscopic model, on the other hand, directly describes the collective behavior of the robotic swarm. It is computationally efficient because it uses fewer variables. Macroscopic models have been successfully applied to a wide variety of problems in physics, chemistry, biology and the social sciences. In these applications, the microscopic behavior of an individual (e.g., a Brownian particle in a volume of gas or an individual residing in US) is quite complex, often stochastic and only partially predictable, and certainly analytically intractable. Rather than account for the inherent variability of individuals, scientists model the behavior of some *average* quantity that represents the system they are studying (e.g., volume of gas or population of US). Such macroscopic descriptions often have a very simple form and are analytically tractable. It is important to remember that such models do not reproduce the results of a single experiment — rather, the behavior of some observable averaged over many experiments or observations. The two description levels are, of course, related: we can start from the Stochastic Master Equation that describes the evolution of a robot’s probability density and get the Rate Equation, a macroscopic model, by averaging it [8]. In most cases, however, Rate Equations are phenomenological in nature, *i.e.*, not derived from first principles. Below we show how to formulate the Rate Equations describing dynamics of a homogeneous robot swarm by examining the details of individual robot controller.

The Rate Equation is not the only approach to modeling collective behavior. Anderson [2], for example, shows how geometric analysis can be used to predict distribution of individuals playing spatial participative games from the microscopic rules each individual is following.

2.1 Stochastic Approach to Modeling Robotic Swarms

The behavior of individual robots in a swarm has many complex influences, even in a controlled laboratory setting. Robots are influenced by external forces, many of which may not be anticipated, such as friction, battery power, sound or light signals, *etc.* Even if all the forces are known in advance, the robots are still subject to random events: fluctuations in the environment, as well as noise in the robot’s sensors and actuators. A robot will interact with other robots whose exact trajectories are equally complex, making it impossible to know which robots will come in contact with one another. Finally, the designer can take advantage of the unpredictability and incorporate it directly into the robot’s behavior: *e.g.*, the simplest effective policy for obstacle avoidance is for the robot to turn a

random angle and move forward. In summary, the behavior of robots in a swarm is so complex, it is best described probabilistically, as a stochastic process.




Fig. 1. Diagram of a robot controller for the simplified foraging scenario

Consider Figure 1, it depicts a controller for a simplified foraging scenario. Each box represents a robot’s state — the action it is executing. In the course of accomplishing the task, the robot will transition from searching to puck pick-up to homing. Transitions between states are triggered by external stimuli, such as encountering a puck. This robot can be described as a stochastic Markov process³, and the diagram in Figure 1 is, therefore, the Finite State Automaton (FSA) of the controller.

The stochastic process approach allows us to mathematically study the behavior of robot swarms. Let $p(n, t)$ be the probability robot is in state n at time t . The Markov property allows us to write change in probability density as [8]

$$\begin{aligned} \Delta p(n, t) &= p(n, t + \Delta t) - p(n, t) \\ &= \sum_{n'} p(n, t + \Delta t | n', t) p(n', t) - \sum_{n'} p(n', t + \Delta t | n, t) p(n, t). \end{aligned} \quad (1)$$

The conditional probabilities define the transition rates for a Markov process

$$W(n|n'; t) = \lim_{\Delta t \rightarrow 0} \frac{p(n, t + \Delta t | n', t)}{\Delta t}. \quad (2)$$

The quantity $p(n, t)$ also describes a macroscopic variable — the fraction of robots in state n , with Equation 1 describing how this variable changes in time. Averaging both sides of the equation over the number of robots (and assuming only individual transitions between states are allowed), we obtain in the continuous limit ($\lim_{\Delta t \rightarrow 0}$)

$$\frac{dN_n(t)}{dt} = \sum_{n'} W(n|n', t) N_{n'}(t) - \sum_{n'} W(n'|n, t) N_n(t), \quad (3)$$

where $N_n(t)$ is the *average* number of robots in state n at time t . This is the so-called Rate Equation. It is sometimes also written in a discrete form, as a finite

³ A Markov process’s future state depends only on its present state and none of the past states.

difference equation that describes the behavior of $N(kT)$, k being an integer and T the discretization interval: $(N(t + T) - N(t))/T$. Equation 3 has the following interpretation: the number of robots in state n will increase in time due to transitions *to* state n from other states, and it will decrease in time due to the transitions *from* state n to other states.

Rate Equations are deterministic. In stochastic systems, however, they describe the dynamics of average quantities. How closely the average quantities track the behavior of the actual dynamic variables depends on the magnitude of fluctuations. Usually the larger the system, the smaller are the (relative) fluctuations. In a small system, the experiment may be repeated many times to average out the effect of fluctuations. The agreement increases as the size of the system grows.

2.2 A Recipe for Model Construction

The Rate Equation is a useful tool for mathematical analysis of collective dynamics of robot swarms. To facilitate the analysis, we begin by drawing the macroscopic state diagram of the system. *The collective behavior of the swarm is captured by an FSA that is functionally identical to the individual robot FSA, except that each state of the automaton now represents the number of robots executing that action* [10, 7, 11]. Not every microscopic robot behavior need to become a macroscopic state. In order to keep the model tractable, it is often useful to coarse-grain it by considering several related actions or behaviors as a single state. For example, we may take the searching state of robots to consist of the actions *wander in the arena*, *detect objects* and *avoid obstacles*. When necessary, the searching state can be split into three states, one for each behavior; however, we are often interested in the *minimal* model that captures the important behavior of the system. Coarse-graining presents a way to construct such a minimal model.

The macroscopic automaton can be directly translated into the Rate Equations. Each state in the automaton becomes a dynamic variable $N_n(t)$, with its own Rate Equation. Every transition will be accounted for by a term in the equation: a positive term for the incident ($W(n|n')N_{n'}$) arrows and negative term for the outgoing ($W(n'|n)N_n$) arrows.

Finding an appropriate mathematical form for the transition rates is the main challenge in studying real systems. The transition is triggered by some stimulus — be it another robot in a particular state, an object to be picked up, *etc.* In order to compute the transition rates, we assume, for simplicity, that robots and stimuli are uniformly distributed. The transition rates then have the following form: $W(n|n') \approx M$, where M is the environmental stimulus encountered (*e.g.*, number of sticks in the arena). The proportionality factor connects the model to experiments, and it depends on the rate at which a robot detects sticks. It can be roughly estimated from first principles (“scattering cross section” approach [10]), measured from simulations or experiments with one or two robots, or left as a model parameter. There will be cases where the uniformity assumption fails: *e.g.*, in overcrowded scenarios where robots, depending on their obstacle avoidance

controller, tend to clump, forming “robotic clouds” [11]. If the transition rates cannot be calculated from first principles, it may be expedient to leave them as parameters of the model and obtain them by fitting the model to data.

3 Application to Swarm Robotic Experiments

The Rate Equation has been used to study a variety of distributed robot systems. Below we illustrate the approach with a few sample results from swarm robotic experiments, for which a body of experimental and simulations data exists.

3.1 Collaborative Stick Pulling

The stick-pulling experiments were carried out to study dynamics of collaboration in robots [6]. The robots’ task was to locate sticks scattered around the arena and pull them out of their holes. A single robot cannot complete the task on its own: rather, when a robot finds a stick, it lifts it partially out of the hole and waits for a period specified by its *gripping time parameter* for a second robot to find it. If a second robot finds the first during this time interval, it will pull the stick out; otherwise, the first robot releases the stick and returns to the searching state.

Lerman *et al.* [10] studied a minimal continuous time model of the system. A minimal model includes only the salient details of the process it describes. They found that this model reproduced key experimental observations and qualitatively agreed with results of experiments and simulations (see Figure 2(a)). Martinoli & Easton [11] formulated a more detailed model based on finite difference equations that accounts for every state in the robot control diagram.

Figure 2 depicts the collaboration rate, the rate at which robots pull sticks out, as a function of the individual robot gripping time parameter for the minimal (a) and the detailed (b) models. Figure 2(b) also shows results of embodied and probabilistic numeric simulations for the same set of parameters. One can see quantitative agreement already with swarms as small as 8 robots. The minimal model shows the same qualitative behavior as the more detailed model.

3.2 Collective Object Collection

Mathematical models have been applied to study collective collection experiments (aggregation and foraging). In the aggregation experiments, the task was to gather small objects in a single cluster starting from a situation where they were all randomly scattered in an arena [12, 1]. Swarms of robots of different group size, or differing in the sensing and actuation capabilities, were used to aggregate different types of objects. These publications considered both microscopic and macroscopic models as well as a few metrics for measuring the evolution of aggregation (average cluster size, number of clusters, size of the biggest cluster). Figure 3(a) shows the results of macroscopic model’s predictions compared to realistic embodied simulation for swarm sizes of one and five robots (see




Fig. 2. Collaboration rate per robot vs gripping time parameter for different robot group sizes and 16 sticks. (a) Results of the minimal model for 8 (short dash), 16 (long dash) and 24 (solid line) robots. (b) Results for detailed model (solid lines), embodied simulations (dotted lines), the microscopic model (dashed lines).

[1] for details). It is worth nothing that, although certain swarm sizes considered were extremely small, quantitative agreement between model and realistic simulation was achieved. The authors also report experiments using variable swarm sizes, by enabling robots to decide whether to continue aggregating the objects or rest. Also in this scenario, theoretical predictions were extremely faithful not only in predicting dynamics of aggregation but also the number of active workers over time.

In foraging experiments, Lerman and Galstyan studied the influence of physical interference on the swarm performance [7]. Interference is a critical issue in swarm robotics, in particular in foraging experiments where there is a spatial bottleneck at the predefined “home” region where the collected objects must be delivered. When two robots find themselves within sensing distance of one another, they will execute obstacle avoidance maneuvers. Because this behavior takes time, interference decreases robots’ efficiency. Clearly, a single robot working alone is relatively more efficient, because it does not experience interference from other robots (the larger the swarm, the greater the degree of interference). However, parallel work helps speed up the foraging process and increases the system robustness in case of individual robot failures.

Figure 3(b) shows the total time required to complete the task for two different interference strengths, as measured by the avoiding time τ . For both cases task completion time is minimized for some swarm size and increases for larger swarms. The greater the effect of interference (larger τ), the smaller the optimal swarm size. Results show good quantitative agreement with embodied simulations with swarms of one to 20 robots.

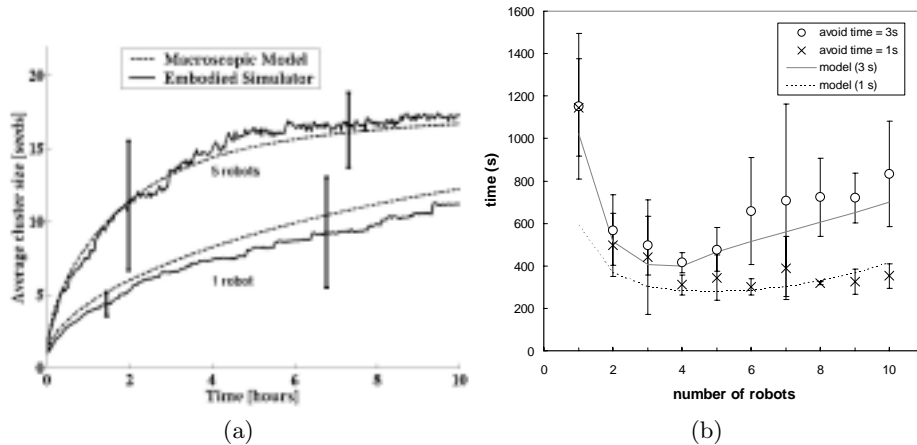


Fig. 3. (a) Evolution of the mean cluster size in an arena with 20 objects to gather and swarms of one and five robots. Macroscopic model (dashed lines) and embodied simulation (continuous lines) are compared. (b) Time it takes the swarm of robots to collect objects in the arena for two different interference strengths. Symbols are results of embodied simulations, while lines give the model’s predictions.

4 Discussion

The macroscopic methods used to analyze collective behavior of robot swarms are based on viewing individual robots as stochastic Markov processes. In order to construct a description of the behavior of a swarm, we do not need to know the exact trajectories of every robot; instead, we derive a model that governs the dynamics of the aggregate, or average, swarm behavior.

A number of simplifying assumptions and specific conditions were used in the methods presented in this paper. While these are not strictly necessary for the validity of the overall approach, they are important for producing mathematically tractable macroscopic models. First, we assume that robot’s actions are largely independent of one another (dilute limit), and the transition rates can be represented by aggregate quantities that are spatially uniform and independent of individual robots or their trajectories. Second, up to date we considered exclusively nonspatial metrics for evaluating collective swarm performance. As long as detection areas do not overlap between the objects placed in the arena and the metric does not specifically address spatiality, these assumptions are correct. Third, we assumed that modeled robots have perfectly centered, uniform, and precise range of detection for each object they may encounter in the arena, in contrast to the individual, heterogeneously distributed, noisy sensors available to the real robots and in the embodied simulation. Fourth, modeled robots are characterized by a set of parameters, each of them representing the mean value of some real robot feature: mean speed, mean duration for performing a certain maneuver, and so on. We do not consider parameter distributions in our models. Fifth, further difficulties may arise due to behavioral granularity captured in the

models. For instance, certain controllers can be approximated by a FSA, though certain routines (e.g., obstacle avoidance and interference) might rely on different control architectures. For instance, proximal control architectures such as neural networks are often used for such routines. They tightly couple actuators with sensors without passing through a distal representation as, for instance, is the case for behavior-based implementations. Parameters used to describe the states corresponding to such routines can still be measured in systematic tests with one or two real robots for achieving well calibrated models without using free parameters but this usually implies some inaccuracies. While for predicting high level metrics such as those considered in the distributed manipulation experiments mentioned above, this approximation is quite sufficient, for other metrics closer to the the robot movements (e.g., average number of robots in search or in obstacle avoidance) such inaccuracies might have a more relevant effect. As a consequence, it might be much more difficult to achieve quantitative agreement between models' predictions and experimental results without fitting the data. Finally, depending on the type of experiment performed and the metrics used, nonlinear mapping between microscopic and macroscopic representations might generate prediction discrepancies between the two type of models simply because average quantities in closed form at the macroscopic level cannot be calculated from the linear combination of the individual Markov chains constituting the microscopic model. As a general rule, discrepancies between the two modeling categories are usually more important with smaller swarm sizes. In such cases microscopic models are often more faithful than macroscopic ones.

5 Conclusion and Outlook

In this paper we have reviewed methods for macroscopically modeling and analyzing the behavior of robot swarms. Our analysis is based on the theory of distributed stochastic processes, which is applicable to robot swarms because the behavior of each robot is inherently probabilistic in nature and often not completely predictable, and its future state depends only on its present state. Despite the inherent unpredictability, the probabilistic description of the collective behavior is surprisingly simple. We showed that Rate Equations describe how the average collective system properties change in time. These equations can be easily written down from the details of the individual robot controller. We illustrated the formalism by reporting a few sample results from swarm robotics experiments presented in the past. Analysis yields important insights into the system, such as what are the important parameters that determine the behavior, how to optimize swarm performance, etc.

Much work remains to be done in extending stochastic mathematical models to new domains and overcoming limitations of the current models. For example, Lerman & Galstyan [9, 4] have moved beyond simple Markov processes to study distributed systems composed of adaptive robots that can change their behavior based on their estimates of the global state of the system. Another unexplored area is in modeling systems in which position has to be taken into account.

Such systems include any that are based on diffusing pheromone fields. Another research direction is to move beyond the mean-field approximation and develop exact statistical formulations of problems. Such formulations will enable us to study directly stochastic effects, including the strength of fluctuations.

6 Acknowledgments

Martinoli is currently sponsored by a Swiss NSF Professorship contract Nr. PP002-68647/1. Lerman and Galstyan are supported in part by the Defense Advanced Research Projects Agency (DARPA) under contract F30602-00-2-0573.

References

1. Agassounon, W., Martinoli, A. and Easton, K. 2004 Macroscopic Modeling of Aggregation Experiments using Embodied Agents in Teams of Constant and Time-Varying Sizes. Special issue on Swarm Robotics, Dorigo, M. and Sahin, E. editors, *Autonomous Robots*, **17**(2-3):163–191.
2. Anderson, C. 2003. Linking Micro- to Macro-level Behavior in the Aggressor-Defender-Stalker Game, in *Workshop on the Mathematics and Algorithms of Social Insects (MASI-2003)*, December, 2003, Atlanta, GA.
3. Bonabeau, E., Dorigo M. and Theraulaz, G. 1999. *Swarm Intelligence: From Natural to Artificial Systems*. Oxford University Press, New York.
4. Galstyan, A. and Lerman, K. 2004. Analysis of a Stochastic Model of Adaptive Task Allocation in Robots, to appear in *Workshop on Engineering Self-Organizing Systems at AAMAS-2004*.
5. Gerkey, B. P., Vaughan, R. T., Stoy, K., Howard, A., Sukhatme G. S., Matarić, M. J. 2001. Most Valuable Player: A Robot Device Server for Distributed Control, in *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2001)*, Wailea, Hawaii, October 29 - November 3, 2001.
6. 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.
7. Lerman, K. and Galstyan, A. 2002a. Mathematical model of foraging in a group of robots: Effect of interference. *Autonomous Robots*, **13**(2):127–141.
8. Lerman, K. and Galstyan, A. 2002b. Two paradigms for the design of artificial collectives. In *Proc. of the First Annual workshop on Collectives and Design of Complex Systems*, NASA-Ames, CA.
9. Lerman, K. and Galstyan, A. 2003. Macroscopic Analysis of Adaptive Task Allocation in Robots. pp. 1951–1956. In *Proc. of the Int. Conf. on Intelligent Robots and Systems (IROS-2003)*, Las Vegas, NV.
10. Lerman, K., Galstyan, A., Martinoli, A. and Ijspeert, A. 2001. A macroscopic analytical model of collaboration in distributed robotic systems. *Artificial Life Journal*, **7**(4):375–393.
11. Martinoli, A., Easton, K. and Agassounon, W. 2004. Modeling of Swarm Robotic Systems: A Case Study in Collaborative Distributed Manipulation. Special Issue on Experimental Robotics, Siciliano, B., editor, *Int. Journal of Robotics Research*, **23**(4):415–436.

12. Martinoli, A., Ijspeert, A. J., and Gambardella, L. M. 1999. A probabilistic model for understanding and comparing collective aggregation mechanisms. pp. 575–584. In D. Floreano, J.-D. Nicoud, and F. Mondada, editors, *LNAI:1674*, Springer, New York, NY.
13. Michel, O. 2004. Webots: Professional Mobile Robot Simulation. *Int. J. of Advanced Robotic Systems*, **1**(1):39–42.
14. Sugawara, K. and Sano, M. 1997. Cooperative acceleration of task performance: Foraging behavior of interacting multi-robots system. *Physica* **D100**:343–354.