A Macroscopic Analytical Model of Collaboration in Distributed Robotic Systems

Kristina Lerman
Aram Galstyan
Information Sciences Institute
University of Southern California
Marina del Rey, CA 90292, USA
lerman@isi.edu

Alcherio Martinoli
Collective Robotics Group
California Institute of Technology
Pasadena, CA 91125, USA

Auke Ijspeert
Department of Computer Science
University of Southern California
Los Angeles, CA 90089, USA

Abstract In this article, we present a macroscopic analytical model of collaboration in a group of reactive robots. The model consists of a series of coupled differential equations that describe the dynamics of group behavior. After presenting the general model, we analyze in detail a case study of collaboration, the stick-pulling experiment, studied experimentally and in simulation by Ijspeert et al. [*Autonomous Robots, 11*, 149–171]. The robots' task is to pull sticks out of their holes, and it can be successfully achieved only through the collaboration of two robots. There is no explicit communication or coordination between the robots. Unlike microscopic simulations (sensor-based or using a probabilistic numerical model), in which computational time scales with the robot group size, the macroscopic model is computationally efficient, because its solutions are independent of robot group size. Analysis reproduces several qualitative conclusions of Ijspeert et al.: namely, the different dynamical regimes for different values of the ratio of robots to sticks, the existence of optimal control parameters that maximize system performance as a function of group size, and the transition from superlinear to sublinear performance as the number of robots is increased.

Keywords
robotics, mathematical modeling, swarm intelligence

1 Introduction

Swarm intelligence [3] is an innovative computational and behavioral metaphor for solving distributed problems that takes its inspiration from the biological examples provided by social insects [6]—ants, termites, bees, and wasps—and by swarming, flocking, herding, and shoaling phenomena in vertebrates [26]. The abilities of such 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 has been found to be mediated by nothing more than a small set of simple low-level interactions between individuals, and between individuals and the environment. The swarm intelligence approach emphasizes distributedness and exploitation of direct (robot-to-robot) or indirect (via the environment) local interactions among relatively simple agents.

The main advantages of the application of the swarm approach to the control of a group of robots are threefold: (a) scalability: the control architecture is kept exactly the same from a few units to thousands of units; (b) flexibility: units can be dynamically added or removed; they can be given the ability to reallocate and redistribute themselves in a self-organized way; (c) robustness: the resulting collective system is robust not only through unit redundancy but also through the unit minimalistic design [5, 20].

Although a formal and quantitative definition of minimalism has yet to be formulated for collective systems, minimalistic design in swarm intelligence implies an effort to keep the resources for computation, sensors, actuators, and communication as low as possible for each unit, while aiming at having as smart as possible group behavior.

In the last few years, the swarm intelligence control principles have been successfully applied to a series of case studies in collective robotics: aggregation [2, 20, 21] and segregation [13], beacon and odor localization [11, 12], collaborative mapping [4], collaborative transportation [15, 17], work division and task allocation [1, 16], and flocking and foraging [23]. All these works have been performed using groups of simple, autonomous robots or embodied simulated agents, exploiting local communication forms among teammates (implicit, through the environment, or explicit, wireless communication), and fully distributed control. Sometimes, due to technical difficulties in experimentation with real robots, local explicit communication [4, 12, 23] or specific environmental information (e.g., nest energy in [16]) has been obtained with the help of absolute positioning systems and/or global communication. While global communication capabilities, if used extensively, represent a bottleneck for the scalability of the collective system, global positioning systems (GPSs), depending on their specific implementation, can achieve performances independent of the team size (e.g., GPS or the system used in [4]) and, therefore, represent suitable technical aids for applying the swarm intelligence approach to artificial systems. Unfortunately, the lack of rigorous, scalable methodologies for designing and analyzing such fully distributed robotics systems has, for the moment, prevented a more extensive application of the swarm intelligence approach to real-world applications such as traffic regulation [31] or surveillance [7].

This article aims at contributing to research in swarm intelligence (a) by making a quantitative study of how collaboration in a group of simple reactive, autonomous robots can be obtained and controlled through the exploitation of local interactions, and (b) by proposing a novel methodology for mathematical analysis of group behavior based on a system of differential equations.

2 Collaboration in Robots

Collaboration can significantly improve the performance of a multi-agent system. In "strictly collaborative" systems [20], collaboration is an explicit requirement, because no single agent can successfully complete the task on its own. Such systems are common in insect as well as human societies, for example, transport of objects too heavy or awkward to be lifted by a single ant, flying the space shuttle, playing a soccer match, and so on. Collaboration in a group of robots has been studied by several groups [14, 17, 22, 24, 29, 30]. We will focus on a specific case study initiated by Martinoli and collaborators [22] and studied in detail by Ijspeert et al. [14], that takes a swarm intelligence approach to collaboration. In this system collaboration in a group of simple reactive agents was achieved entirely through local interactions, that is, without explicit global communication or coordination among robots. Because of a purely swarm approach, this system is a compelling and effective model of how collaboration may arise in natural systems, such as insect societies. In addition, the simplicity of the robots' interactions lends itself to mathematical analysis. In this article we will propose and study an analytical model of collaboration in a group of robots, presenting the general case first, then analyzing the system studied by Ijspeert et al. and comparing the results of analysis to experimental results and simulation.

As mentioned in the previous section, there has been relatively little prior work in mathematical analysis of multirobot systems in general and collaboration in particular, with the exception of Sugawara and coworkers' [28, 29] research. They carried out a

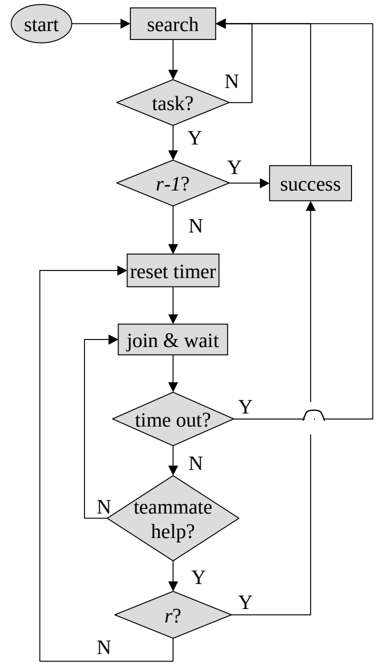

Figure 1. Schematic of a robot's controller for collaborative task completion.

quantitative study of cooperative foraging in a group of communicating robots. They have developed a simple state-based analytical model and analyzed it under different conditions. In their system when a robot finds a puck or a collection of pucks, it may broadcast a signal for a period of time to other robots, which move toward it. The robots pick up pucks and bring them home. Sugawara et al. did not take the interaction into account explicitly but in an approximate manner. In our model, we will include the duration of the interaction explicitly, resulting in a better description of the dynamics of the system. Another difference between their work and ours is that their system is not strictly collaborative—collaboration via signaling improves performance but is not a requirement for task completion.

2.1 A Model of Collaboration in Robots

Consider a homogeneous system composed of N robots and M spatially distributed tasks. The tasks are such that a single robot cannot execute one on its own—a collaboration of r ($r < N$) robots is required to complete each task successfully. The task could be long sticks that have to be pulled out of the ground or heavy objects that need to be transported by several robots. We consider a swarm intelligence approach that uses *simple locally interacting* robots to achieve collaboration in the absence of central or hierarchical control and explicit communication between robots. We consider a homogeneous system in which each robot has the same simple controller, schematically represented in Figure 1.

Each robot explores the arena, looking for tasks and avoiding obstacles. If it finds itself at the location of the task, it prepares to execute it. If there are no robots present at this location, the robot stops and waits for some period of time τ. If no other robots come to its aid during this time interval (time out), the robot abandons the task and

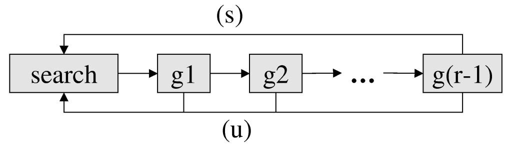

Figure 2. Macroscopic state diagram of collaboration in a multirobot system. The arrow marked "s" corresponds to the transition to the search state after a successful collaboration has occurred, while the arrow marked "u" corresponds to the transition after an unsuccessful collaboration, that is, when waiting time exceeded τ.

resumes the search. If another robot encounters it, the first robot resets its timer,[1] and both robots wait for the same time interval τ. Now there is a group of size two waiting to execute a task. If no other robot encounters the group during the waiting period, both robots abandon the task and return to the search mode, but if another robot does find it, the first two robots reset their timers and all three robots wait for a time τ. This is repeated until a group of size $r-1$ is waiting to perform a task. If a robot finds this group during the time interval τ, the task is completed successfully, and all r robots resume the search; otherwise, $r-1$ robots abandon the task and start searching again. Other designs can also lead to successful task execution (e.g., communication will help assemble a group of size r faster than random search); however, we will focus on this simple system and show how to construct a macroscopic mathematical model to describe the dynamics of collaboration.

Generally, two tools—experiment and simulation—have been available for the study of multirobot systems. Experiments with real robots allow researchers to observe swarms under real conditions; however, experiments are very costly and time consuming, and systematically varying individual robot parameters to study their effect on the group behavior is often impractical. Simulations, such as sensor-based simulations for robots, attempt to model the environment realistically, the robots' imperfect sensing of and interactions with it. Though simulations are much faster and are less costly than experiments, still they suffer from many of the same limitations, namely, they are tedious to perform, and it is often still impractical to explore the parameter space systematically. Mathematical analysis provides an alternative to experiment and simulation as a tool for the study of behavior of multirobot systems. Using mathematical analysis we can study dynamics of even large systems, predict their long-term behavior, and gain insight into system design: for example, verify the existence of optimal parameters and estimate their values. In Section 2.2, we present an analytical model of dynamics in the collaborative system described above. In Section 4, we will analyze a case study of collaboration in a multirobot system: the stick-pulling experiment.

2.2 The Dynamical Model

To construct a model of collaboration in a multirobot system, it is helpful to draw the macroscopic state diagram of the system (Figure 2). During a sufficiently short time interval, each robot can be thought to belong to the *search* state, or be part of a group of size one ($g1$), two ($g2$), and so on, up to a group of size $r-1$ [$g(r-1)$]. The search state consists of a set of behaviors associated with looking for tasks, such as wandering around the arena, detecting objects, and avoiding obstacles. We assume that successful completion of these actions takes place on a short enough time scale that it can be incorporated into the search state.

1 This operation would require communication between robots in the group. However, communication through the environment rather than explicit communication may be sufficient to accomplish this goal: for example, if a load gets lighter, the robots in the group know another robot has joined them.

In addition to states, we must specify transitions between states. When a searching robot locates a task in the arena and begins the wait for help, it makes a transition to state $g1$. If no help arrives (unsuccessful collaboration), it makes a transition back to the search state; otherwise, it makes a transition to state $g2$. Again, if no help arrives, it makes a transition to the search state; otherwise, it makes a transition to state $g3$. Therefore, except for search state, there is one transition *to* each state and two transitions *from* the state. The two transitions from the $g(r-1)$ state correspond to a successful task completion and unsuccessful collaboration.

Each of the boxes in the state diagram in Figure 2 becomes a dynamic variable in the mathematical model. Let $N_s(t)$ be the number of robots in the search state, $N_k(t)$, $1 \leq k \leq r-1$, be the number of groups of size 1 up to $r-1$ at time t. Also, let $M(t)$ be the number of uncompleted tasks at time t. This variable does not represent a macroscopic state; rather it tracks the state of the environment. A mathematical model describes how the dynamic variables change in time. We have a choice of two formalisms for the model: (a) a difference equation, $\Delta N = N(t+\Delta t) - N(t)$, that governs how N changes in time, or (b) a differential equation of the form $dN(t)/dt$. The first model deals with discrete variables, but its results depend on the choice of Δt. In the continuum limit, as $\Delta t \to 0$, the instantaneous change in N is given by the derivative $dN/dt = \lim_{\Delta t \to 0} \Delta N/\Delta t$. Here, N must be thought of as a continuous variable, an approximation of a discrete quantity. Though the approximation is more accurate for larger values of N, it is often used for moderately large and even smaller quantities. Additionally, the dynamic variables in our model are *average* quantities [18]; therefore, it is reasonable to treat them as continuous variables.

We assume that robots and tasks are distributed uniformly around the arena. A series of differential rate equations describes how the dynamic variables change in time:

$$\begin{aligned}
\frac{dN_s}{dt} &= -\alpha N_s(t)\left(M(t) - \sum_{k=1}^{r-1} N_k(t)\right) - N_s(t)\sum_{k=1}^{r-2}\tilde{\alpha}_k N_k(t) + (r-1)\tilde{\alpha}_{r-1}N_s(t)N_{r-1}(t) \\
&\quad + \alpha N_s(t-\tau)\left(M(t-\tau) - \sum_{k=1}^{r-1} N_k(t-\tau)\right)\Gamma_1(t;\tau) \\
&\quad + N_s(t-\tau)\sum_{k=1}^{r-2} k\tilde{\alpha}_k N_k(t-\tau)\Gamma_k(t;\tau) \qquad (1)
\end{aligned}$$

$$\begin{aligned}
\frac{dN_1}{dt} &= \alpha N_s(t)\left(M(t) - \sum_{k=1}^{r-1} N_k(t)\right) - \tilde{\alpha}_1 N_s(t)N_1(t) \\
&\quad - \alpha N_s(t-\tau)\left(M(t-\tau) - \sum_{k=1}^{r-1} N_k(t-\tau)\right)\Gamma_1(t;\tau) \qquad (2)
\end{aligned}$$

$$\begin{aligned}
\frac{dN_n}{dt} &= \tilde{\alpha}_{n-1}N_s(t)N_{n-1}(t) - \tilde{\alpha}_n N_s(t)N_n(t) \\
&\quad - \tilde{\alpha}_{n-1}N_s(t-\tau)(M(t-\tau) - N_{n-1}(t-\tau))\Gamma_n(t;\tau), \qquad n = 2, \ldots, r-1 \qquad (3)
\end{aligned}$$

$$\frac{dM}{dt} = -\tilde{\alpha}_{r-1}N_s(t)N_{r-1}(t) + \mu(t) \qquad (4)$$

where α, $\tilde{\alpha}_n$ are, respectively, the rates at which a searching robot encounters a task and a group of size n waiting to execute the task, and τ is the waiting period. $\Gamma_n(t;\tau)$ is the fraction of groups of size n to abandon their tasks, and it will be derived below. The rate at which new tasks are added is $\mu(t)$. The first two terms in Equation 1 describe a decrease in the number of searching robots because robots find isolated tasks or join a group waiting to execute a task. The last three terms describe an increase in the number of searching robots: the first due to successful task completion, and the last two due to unsuccessful collaboration, that is, when the group times out. The three terms in Equation 2 correspond to the three arrows entering and leaving state $g1$ in Figure 2. The first term accounts for the increase in the number of groups of size one because some robots find tasks that have not been found by other robots and begin the wait for help. Under the uniform distribution assumption, the rate at which robots encounter these tasks is proportional to the number of tasks in the arena, with the proportionality factor given by α. The second term describes the decrease in the number of groups of size one triggered by the arrival of searching robots during the waiting period τ, and the final term accounts for the failed collaborations (no help arrives during period τ), which also leads to a decrease in the number of groups of size one. The terms in Equations 3 and 4 have similar interpretations. Note that the total number of robots, $N_0 = N_s + \sum_{k=1}^{r-1} kN_k$, is conserved; therefore, one of the differential equations above, for example, for N_1, is superfluous, and the variable can be computed from the conservation of robots' condition.

The fraction of groups of size n that abandoned their tasks at time t, $\Gamma_n(t;\tau)$, is equivalent to the probability that no robot came "to help" the group during the time interval $[t-\tau, t]$. To calculate $\Gamma_n(t;\tau)$ let us divide the time interval $[t-\tau, t]$ into K small intervals of length $\delta t = \tau/K$. The probability that no robot comes to help during the time interval $[t-\tau, t-\tau+\delta t]$ is simply $1-\tilde{\alpha}N_s(t-\tau)\delta t$. Hence, the probability for a failed collaboration is

$$\Gamma_n(t;\tau) = \prod_{i=1}^{K}[1-\tilde{\alpha}_n\delta t N_s(t-\tau+i\delta t)]\Theta(t-\tau) \quad (5)$$

$$\equiv \exp\left[\sum_{i=1}^{K}\ln[1-\tilde{\alpha}_n\delta t N_s(t-\tau+i\delta t)]\right]\Theta(t-\tau)$$

The step function $\Theta(t-\tau)$ ensures that $\Gamma_n(t;\tau)$ is zero for $t<\tau$. Finally, expanding the logarithm in Equation 6 and taking the limit $\delta t \to 0$ we obtain

$$\Gamma_n(t;\tau) = \exp\left[-\tilde{\alpha}_n\int_{t-\tau}^{t} dt' N_s(t')\right]\Theta(t-\tau) \quad (6)$$

The collaboration rate is defined as the rate at which tasks are completed: $R(t) = \tilde{\alpha}_{r-1}N_s(t)N_{r-1}(t)$. Once we know the solutions $N_s(t)$ and $N_{r-1}(t)$ at some time, we can compute the value of the collaboration rate at that time. Note that if no new tasks are added, $R(t) = -dM/dt$.

To solve Equations 1–4, we need to specify initial conditions. One possible set of initial conditions may be that at $t=0$ all the robots are searching and there are no groups. We will not solve the general case; rather, in Section 4 we will describe and analyze a case study of collaboration in robots given by the stick-pulling experiments.

3 Case Study: Physical Implementation of the Stick-Pulling Experiment

The stick-pulling experiments were carried out by Ijspeert et al. [14] to study the dynamics of collaboration among locally interacting simple reactive robots. Figure 3 is a snapshot of the physical setup of the experiments. The robots' task is to locate sticks

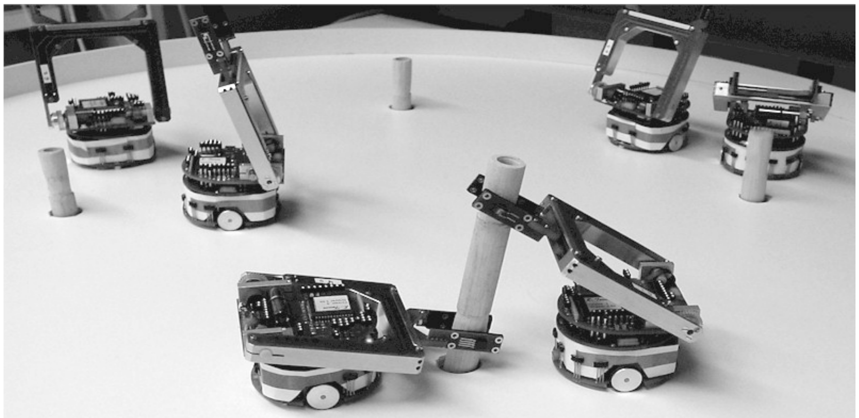

Figure 3. Physical setup of the stick-pulling experiment showing six Khepera robots.

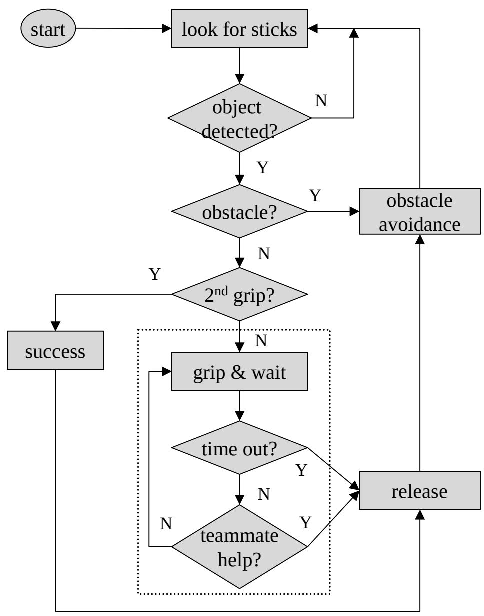

Figure 4. Flowchart of the robots' controller reported from [14] with overlapped state blocks.

scattered around the arena and pull them out of their holes. A single robot cannot pull the stick out by itself—a collaboration between two robots is required for the task to be successfully completed. Collaboration occurs in the following way: One robot finds a stick, lifts it partly out of the ground, and waits for a second robot to find it and complete the task by pulling the stick out of its hole completely.

The actions of each robot are governed by the same simple controller, outlined in Figure 4. The robot's default behavior is to wander around the arena looking for sticks and avoiding obstacles, which could be other robots or walls. When a robot finds a stick that is not being held by another robot, it grips it, lifts it halfway out of the ground and waits for a period of time specified by the *gripping time parameter*. If no other robot comes to its aid during the waiting period (time out), the robot releases the stick and resumes the search for other sticks. If another robot encounters a robot holding a stick, a successful collaboration will take place during which the second robot will grip the stick, pulling it out of the ground completely, while the first robot releases the stick and resumes the search. After the task is completed, the second robot also releases the stick and returns to the search mode, and the experimenter replaces the stick in its hole.

3.1 Real Robots, Embodied Simulations, and Microscopic Modeling

Ijspeert et al. [14] studied the dynamics of collaboration in the stick-pulling experiment at three different levels: by conducting experiments with physical robots; using a sensor-based simulator of robots; and using a microscopic probabilistic model. The physical experiments were carried out in groups of two to six Khepera robots in an arena containing four sticks. Because experiments with physical robots are very time consuming, Webots, the sensor-based simulator of Khepera robots, was used to explore systematically parameters affecting the dynamics of collaboration. The Webots simulator [25] attempts to model the environment faithfully and replicate the experiment by reproducing the robots' (noisy) sensory input and the (noisy) response of the on-board actuators to compute the trajectory and interactions of all the robots in the arena. The probabilistic microscopic model, on the other hand, does not attempt to compute trajectories of individual robots. Rather, it is a numerical model in which the robot's actions—encountering a stick, a wall, another robot, a robot gripping a stick, or wandering around the arena—are represented as a series of stochastic events, with probabilities based on simple geometric considerations and systematic tests with one or two real robots. For example, the probability of a robot encountering a stick is equal to the product of the number of ungripped sticks, and the detection area of the stick normalized by the arena area. Probabilities of other interactions can be similarly calculated. The microscopic simulation consists of running several processes in parallel, one for each robot, while keeping track of the global state of the environment, such as the number of gripped and ungripped sticks. According to Ijspeert et al. [14] the acceleration factor for Webots and real robots can vary between one and two orders of magnitude for the experiments presented here. Because the probabilistic model does not require calculations of the details of the robots' trajectories, it is about 300 times faster than Webots for this experiment.

3.2 Results Obtained at the Three Lower-Level Implementations

Ijspeert et al. [14] systematically studied the collaboration rate, that is, the number of sticks successfully pulled out of the ground in a given time interval, and its dependence on the group size and the gripping time parameter. Though in that work they also investigated the effects of robot heterogeneity and explicit communication, we will focus on a homogeneous system of noncommunicating robots. Ijspeert et al. [14] report very good qualitative and quantitative agreement between the three different levels of experiments, as shown in Figure 5. The main result is that, depending on the ratio of robots to sticks (or workers to the amount of work), there appear to be two different regimes in the collaboration dynamics. When there are fewer robots than sticks, the collaboration rate decreases to zero as the value of the gripping time parameter grows. In the extreme case, when the robot grabs a stick and waits indefinitely for another robot to come and help it, the collaboration rate is zero, because after some period of time each robot ends up holding a stick, and no robots are available to help. When there are more robots than sticks, the collaboration rate remains finite even in the limit as the gripping time parameter becomes infinite, because there will always be robots available to help pull the sticks out. Another finding of Ijspeert et al. [14] was that when there are fewer robots than sticks, there is an optimal value of the gripping time

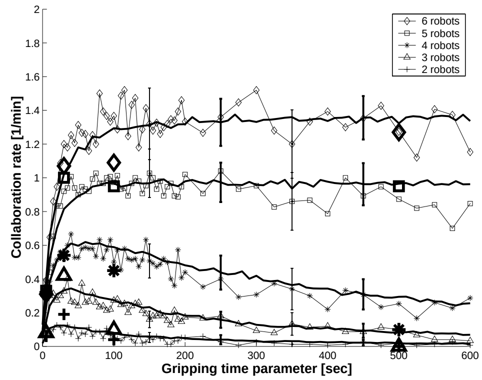

Figure 5. Collaboration rate as a function of the gripping time in homogeneous groups of two to six robots and four sticks. The large single markers correspond to the results with the real robots, the linked small markers to those with the Webots simulator, and the underlying continuous lines to those with the probabilistic simulation (from Ijspeert et al. [14]).

parameter that maximizes the collaboration rate. In the other regime, the collaboration rate appears to be independent of the gripping time parameter above a specific value, so the optimal strategy is for the robot to grip a stick and hold it indefinitely. They also found that the system is one of few collaborative systems known to the authors that demonstrates superlinearity, that is, for some range of robot group sizes and a given number of sticks, adding a robot not only increases the global performance of the system but also the relative performance of the other robots. However, as the robot group size increases, the overcrowding and interference effects cause the relative collaboration rate to saturate and become sublinear.

4 The Macroscopic Analytical Model of the Stick-Pulling Experiment

In the following sections we present a macroscopic analytical model of the stick-pulling experiments in a homogeneous multirobot system. Such a model is useful for the following reasons. First, the complexity of a macroscopic model is independent of the system size, that is, the number of robots: therefore, the time required to obtain solutions for a system of 5,000 robots is as long as that to obtain solutions for a system of 5 robots, whereas for a microscopic description the time required for computer simulation scales at least linearly with the number of robots. Second, our approach allows us to derive analytic expressions for certain important parameters (e.g., those for which the performance is optimal). It also enables us to study the stability properties of the system and see whether solutions are robust under external perturbation or noise. These capabilities are important for the design and control of large multi-agent systems.

To construct a model of the stick-pulling experiments, it is helpful to write the macroscopic state diagram of the system. During a sufficiently short time interval, each robot can be thought to be in one of two states: *searching* or *gripping*. The state labels several related robot behaviors and it is a useful shorthand for thinking about the system. Using a flowchart of the robots' controller, shown in Figure 4, as a reference, we can consider the search state to be the set of behaviors associated with looking for sticks, such as wandering around the arena ("look for sticks" action), detecting objects, and avoiding obstacles; the gripping state is composed of the decisions and actions inside the dotted box of Figure 4. We assume that actions "success" (pull the stick out

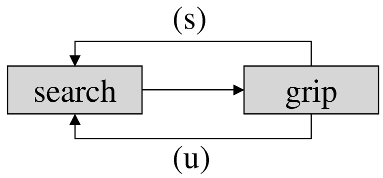

Figure 6. Macroscopic state diagram of the multirobot system. The arrow marked "s" corresponds to the transition from the gripping to the searching state after a successful collaboration, while the arrow marked "u" corresponds to the transition after an unsuccessful collaboration, that is, when the robots time out.

completely) and "release" (release the stick) take place on a short enough time scale that they can be incorporated into the search state. While the robot is in the obstacle avoidance mode, it cannot detect and try to grip objects; therefore, avoidance serves to decrease the number of robots that are searching and capable of gripping sticks. We incorporate avoidance into the model explicitly in Section 4.3. For now, we are interested in the *minimal* model required to explain the main experimental results.

In addition to states, we must also specify all possible transitions between states. When it finds a stick, the robot makes a transition from the search state to the gripping state. After both a successful collaboration and when it times out (unsuccessful collaboration) the robot releases the stick and makes a transition into the searching state, as shown in Figure 6. These arrows correspond to the arrow entering and the two arrows leaving the dotted box in Figure 4. We will use the macroscopic state diagram as the basis for writing down the differential rate equations that describe the dynamics of the stick-pulling experiments.

4.1 The Dynamical Model

The dynamic variables of the model are $N_s(t)$ and $N_g(t)$, the number of robots in the searching and gripping states, respectively. Also, let $M(t)$ be the number of unextracted sticks at time t. The latter variable does not represent a macroscopic state; rather, it tracks the state of the environment. We assume that robots and sticks are distributed uniformly around the arena.

A series of differential rate equations govern the dynamics of the stick-pulling system:

$$\frac{dN_s}{dt} = -\alpha N_s(t)(M(t) - N_g(t)) + \tilde{\alpha} N_s(t) N_g(t) \\ + \alpha N_s(t-\tau)(M(t-\tau) - N_g(t-\tau))\Gamma(t;\tau) \tag{7}$$

$$N_g = N_0 - N_s \tag{8}$$

$$\frac{dM}{dt} = -\tilde{\alpha} N_s(t) N_g(t) + \mu(t) \tag{9}$$

where α, $\tilde{\alpha}$ are the rates at which a searching robot encounters a stick and a gripping robot, respectively, τ is the gripping time parameter, and $\mu(t)$ is the rate at which new tasks are added. The parameters α, $\tilde{\alpha}$, and τ connect the model to the experiment. Parameters α and $\tilde{\alpha}$ are related to the size of the object, the robot's detection radius, or footprint, and the speed at which it explores the arena. The three terms in Equation 7 correspond to the three arrows in Figure 6. The first term accounts for the decrease in the number of searching robots because some robots find and grip sticks. Under the uniform distribution assumption, the rate at which robots encounter ungripped sticks is proportional to the number of ungripped sticks in the arena, with the proportionality factor given by α. The second term describes the successful collaborations between two robots, and the third term accounts for the failed collaborations, both of which lead to an increase in the number of searching robots. The fraction of failed collaborations, $\Gamma(t;\tau)$, is given by Equation 6, with $\tilde{\alpha}_n = \tilde{\alpha}$.

We do not need a differential equation for N_g, the number of gripping robots, because this quantity may be computed using the conservation of robots' condition, Equation 8. The last equation, Equation 9, says that the number of unextracted sticks $M(t)$ decreases in time at the rate of successful collaborations. The equations are subject to the initial conditions that at $t = 0$ the number of searching robots is N_0 and the number of unextracted sticks is M_0.

To proceed further let us introduce $n(t) = N_s(t)/N_0$, $m(t) = M(t)/M_0$, $\beta = N_0/M_0$, $R_G = \tilde{\alpha}/\alpha$, $\tilde{\beta} = R_G\beta$ and a dimensionless time $t \to \alpha M_0 t$, $\tau \to \alpha M_0 \tau$. The dimensionless rate at which new tasks (sticks) are added is μ'. The fraction of robots in the search state is $n(t)$ and $m(t)$ is the fraction of unextracted sticks at time t. Due to the conservation of the number of robots, the fraction of robots in the gripping state is simply $1 - n(t)$. Equations 7–9 can be rewritten in dimensionless form as

$$\frac{dn}{dt} = -n(t)[m(t) + \beta n(t) - \beta] + \tilde{\beta} n(t)[1 - n(t)] + n(t-\tau)[m(t-\tau) + \beta n(t-\tau) - \beta] \times \gamma(t;\tau) \tag{10}$$

$$\frac{dm}{dt} = -\beta\tilde{\beta} n(t)[1 - n(t)] + \mu' \tag{11}$$

$$\gamma(t;\tau) = \exp\left[-\tilde{\beta}\int_{t-\tau}^{t} dt'\, n(t')\right] \tag{12}$$

Equations 10–12 together with initial conditions $n(0) = 1$, $m(0) = 1$ determine the dynamical evolution of the system. Note that only two parameters, β and τ, appear in the equations and, thus, determine the behavior of solutions. The third parameter $\tilde{\beta} = R_G\beta$ is fixed experimentally and is not independent. Note that we do not need to specify α and $\tilde{\alpha}$—they enter the model only through R_G (throughout this article we will use $R_G = 0.35$, the value reported in [14]).[2] Below we provide a detailed analysis of these equations.

4.2 Analysis

Let us assume that new sticks are added to the system at the same rate that the robots pull them out. This situation was realized experimentally by replacing the sticks in their holes after they were pulled out by robots. Therefore, the number of sticks does not change with time ($m(t) = m(0) = 1$). A steady-state solution, if it exists, describes the long-term time-independent behavior of the system. To find it, we set the left-hand side of Equation 10 to zero. Equation 10 has a nontrivial steady-state solution that satisfies the following transcendental equation:

$$-1 + (\beta + \tilde{\beta})(1 - n) + (1 - \beta(1 - n))e^{-\tilde{\beta}\tau n} = 0 \tag{13}$$

2 The parameter α can be easily calculated from experimental values quoted in [14]. As a robot travels through the arena, it sweeps out some area during time dt and will detect objects that fall in that area. This detection area is $V_R W_R dt$, where $V_R = 8.0$ cm/s is the robot's speed, and $W_R = 14.0$ cm is the robot's detection width. If the arena radius is $R = 40.0$ cm, a robot will detect sticks at the rate $\alpha = V_R W_R/\pi R^2 = 0.02\ s^{-1}$. According to [14], a robot's probability to grab a stick already being held by another robot is 35% of the probability of grabbing a free stick. Therefore, $R_G = \tilde{\alpha}/\alpha = 0.35$. R_G is an experimental value obtained with systematic experiments with two real robots, one holding the stick and the other one approaching the stick from different angles.

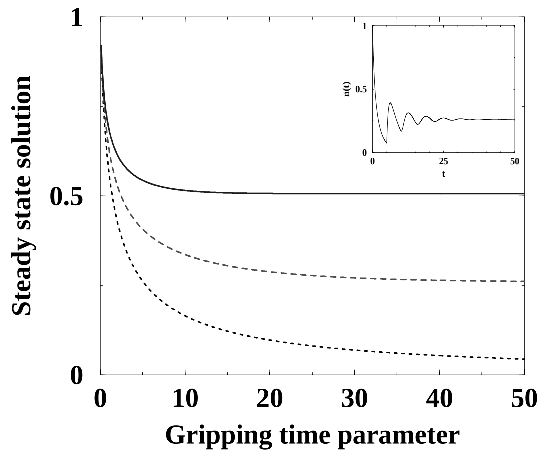

Figure 7. Steady-state solution versus (dimensionless) gripping time parameter τ: for $\beta = 0.5$ (short dash), 1 (long dash), 1.5 (solid line). Inset shows a typical relaxation to the steady state for $\tau = 5$, $\beta = 0.5$.

Figure 7 shows the dependence of the fraction of searching robots in the steady state on the gripping time τ for different values of the parameter β. Note that for small enough β's $n(\tau) \to 0$ as $\tau \to \infty$. The intuitive reason for this is the same one given in Section 3.2: When there are fewer robots than sticks, and each robot holds the stick indefinitely, after a while every robot is holding a stick, and no robots are searching. For $\beta > 1/(1 + R_G)$, however, $n(\tau) \to \text{const} \neq 0$ as $\tau \to \infty$. The inset in Figure 7 shows how a typical solution, $n(t)$, relaxes to its steady-state value. The oscillations are characteristic of time-delay differential equations, and their period is determined by τ.

The collaboration rate is the rate at which robots successfully pull sticks out of their holes. The steady-state collaboration rate $R(\tau; \beta)$ is given by the following equation:

$$R(\tau, \beta) = \beta\tilde{\beta}\, n(\tau, \beta)[1 - n(\tau, \beta)], \tag{14}$$

where $n(\tau, \beta)$ is the number of searching robots in the steady state for a particular value of τ and β, and $(1 - n(\tau, \beta))$ is the number of gripping robots in the steady state. Figure 8 depicts the collaboration rate as a function of τ. There exists a critical value β_c of β such that for $\beta > \beta_c$ the collaboration rate increases monotonically with τ. However, for $\beta < \beta_c$ there is an optimal gripping time, $\tau = \tau_{opt}$, that maximizes the collaboration rate. To understand this behavior note that the maximum collaboration rate for a given β is achieved for $n(\tau, \beta) = 1/2$. For $\beta > \beta_c$, however, the solution of Equation 13 is always greater than 1/2, so an optimal solution does not exist. For $\beta < \beta_c$ a simple analysis gives

$$\tau_{opt} = \frac{2}{\tilde{\beta}} \ln \frac{1 - \beta/2}{1 - 1/2(\beta + \tilde{\beta})}, \qquad \beta < \beta_c = \frac{2}{1 + R_G} \tag{15}$$

The three curves in Figure 8 are qualitatively similar to those in Figure 5 for two robots ($\beta = 0.5$), four robots ($\beta = 1.0$), and six robots ($\beta = 1.5$). Mathematical analysis reproduces the following conclusions of Ijspeert et al. [14]: the different dynamical regimes depending on the value of the ratio of robots to sticks (β) and the optimal gripping time parameter for $\beta < \beta_c$.

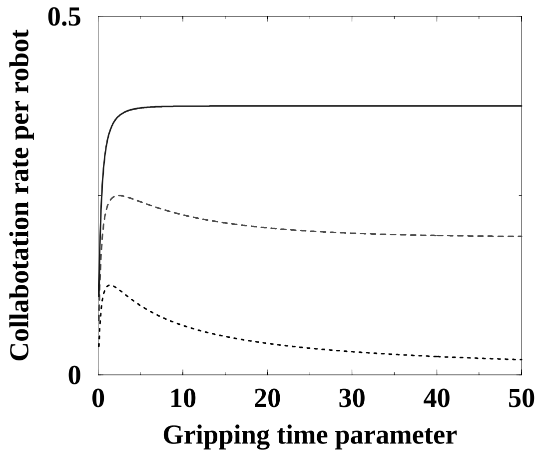

Figure 8. Collaboration rate per robot versus (dimensionless) gripping time parameter τ for $\beta = 0.5$ (short dash), $\beta = 1$ (long dash), $\beta = 1.5$ (solid line). These values of β correspond, respectively, to two, four, and six robots in the experiments with four sticks (cf. Figure 5).

4.3 Interference Effects

In the previous section we neglected the effects of interference between robots. Interference is the result of competition for space between spatially extended robots. When a robot finds itself within sensing distance of an obstacle (another robot or a wall), it will execute obstacle avoiding behavior to reduce the risk of a potentially damaging collision. Obstacle avoidance takes time; therefore, interference may impact the performance of the system. Ijspeert et al. [14] showed that adding more robots can lead to a drastic deterioration in the system's performance. We now address this question in the framework of the approach developed in the previous sections.

To model the avoiding behavior we assume that each time a robot encounters an obstacle it "halts" for a certain amount of time τ_{av} and resumes the search afterward. Although this is a very simplified version of the real situation, we found that this approach reproduces the main effects of the experiment. The macroscopic state diagram, Figure 6, will be modified by an inclusion of a new *avoiding* state, with arrows to and from from the *searching* state. A searching robot will make a transition to the avoiding state when it encounters another robot, which can be in the searching, gripping, or avoiding states. After a period of time τ_{av}, the robot will finish executing the avoiding behavior and resume the search. We neglect avoidance of walls. This effect contributes a constant term for each robot and becomes less important as the arena area is increased.

Let $N_{\mathrm{av}}(t)$ be the number of robots in the *avoiding* state at time t. Again, we will consider a static environment only, where the number of sticks remains constant. Taking the avoiding behavior into account modifies the model (cf. Equation 7) as follows:

$$\frac{dN_{\mathrm{s}}}{dt} = -\alpha N_{\mathrm{s}}(t)(M_0 - N_{\mathrm{g}}(t)) + \tilde{\alpha} N_{\mathrm{s}}(t) N_{\mathrm{g}}(t) + \alpha N_{\mathrm{s}}(t-\tau)(M_0 - N_{\mathrm{g}}(t-\tau))\Gamma(t;\tau) - 2\alpha_1 N_{\mathrm{s}}(t)(N_{\mathrm{s}}(t) - 1) - \alpha_2 N_{\mathrm{s}}(t) N_{\mathrm{av}}(t) - \alpha_3 N_{\mathrm{s}}(t) N_{\mathrm{g}}(t) + \frac{1}{\tau_{\mathrm{av}}} N_{\mathrm{av}} \tag{16}$$

$$\frac{dN_{\mathrm{av}}}{dt} = -\frac{1}{\tau_{\mathrm{av}}} N_{\mathrm{av}} + 2\alpha_1 N_{\mathrm{s}}(t)(N_{\mathrm{s}}(t) - 1) + \alpha_2 N_{\mathrm{s}}(t) N_{\mathrm{av}}(t) + \alpha_3 N_{\mathrm{s}}(t) N_{\mathrm{g}}(t) \quad (17)$$

$$N_{\mathrm{g}} = N_0 - N_{\mathrm{s}}(t) - N_{\mathrm{av}}(t). \quad (18)$$

The first three terms in Equation 16 have the same meaning as for Equation 7. The next three terms describe the loss in the number of searching robots due to avoidance. The rates at which a searching robot encounters, and has to avoid, another searching, avoiding, or gripping robot are given by parameters α_1, α_2, and α_3, respectively. The rate at which avoiding robots finish the avoiding behavior and resume searching is given, on average, by $N_{\mathrm{av}}/\tau_{\mathrm{av}}$. The conservation of the total number of robots is given by Equation 18.

We are interested in the steady-state properties of system (Equations 16–17), that is, $dN_{\mathrm{s}}/dt = dN_{\mathrm{av}}/dt = 0$. Steady-state solutions describe the long-term, time-independent behavior of the system. Let $n_{\mathrm{s}} = N_{\mathrm{s}}/N_0$, $n_{\mathrm{av}} = N_{\mathrm{av}}/N_0$ be the fraction of robots in searching and avoiding states, respectively. Then, the steady-state solutions n_{s} and n_{av} satisfy the following equations:

$$(1-\gamma)[1-\beta(1-n_{\mathrm{s}}-n_{\mathrm{av}})] + \tilde{\beta}[1-n_{\mathrm{s}}-n_{\mathrm{av}}] = 0 \quad (19)$$

$$-\frac{1}{\tau_{\mathrm{av}}} n_{\mathrm{av}} + 2\beta_1 n_{\mathrm{s}}(n_{\mathrm{s}}-\varepsilon) + \beta_2 n_{\mathrm{s}} n_{\mathrm{av}} + \beta_3 n_{\mathrm{s}}(1-n_{\mathrm{s}}-n_{\mathrm{av}}) = 0 \quad (20)$$

where $\varepsilon = 1/N_0$, $\tilde{\beta} = \beta\tilde{\alpha}/\alpha, \beta_i = \beta\alpha_i/\alpha, (i = 1, 2, 3)$. The parameter ε describes the finite size effect. For relatively small systems, such as the ones studied in the experiments, ε is finite, but it approaches zero as the number of robots in the system becomes large.

We solved Equations 19 and 20 numerically to obtain steady-state values of n_{s} and n_{av} that we can use to calculate the collaboration rate. The collaboration rate, the rate at which the robots pull sticks out, is given by the following dimensionless expression:

$$R(t) = \beta\tilde{\beta} n_{\mathrm{s}}(t)[1 - n_{\mathrm{s}}(t) - n_{\mathrm{av}}(t)]. \quad (21)$$

Including the effects of interference does not qualitatively change the behavior of the collaboration rate as a function of τ and β. However, we have found that it does affect the performance of the system as the group size, N_0, increases.

Figure 9 shows the optimal (maximal) collaboration rate per robot as a function of the robot group size and for three different interference strengths. For small N_0, the performance of each robot increases with group size for all interference strengths, which suggests that the system as a whole performs superlinearly. However, interference and overcrowding, as measured by τ_{av} and the total number of robots, degrade the performance of the system. As the number of robots grows, the superlinear regime is followed by an almost linear and then a sublinear regime for nonzero interference strengths. The saturation and decrease of the relative performance already occurs for moderately large groups and this agrees qualitatively with results of Ijspeert et al. (see [14], Figure 12).

5 Discussion and Future Work

This article, together with [14], presents *three* levels of abstraction for modeling a robotic experiment: (a) sensor-based simulations, (b) microscopic numerical model, (c) macroscopic analytical model. Each level has its advantages and drawbacks. The

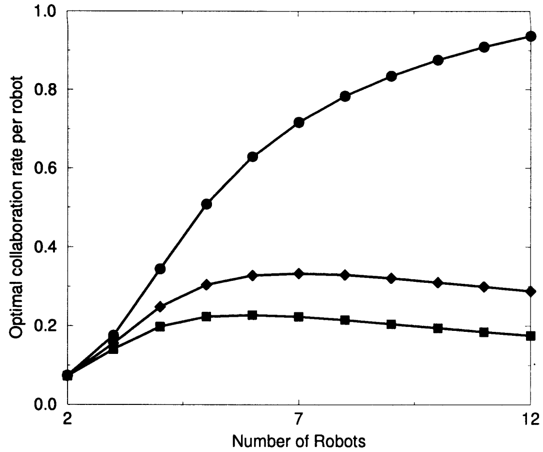

Figure 9. Relative collaboration rate versus the number of robots for different values of interference strength: $\tau_{av} = 0$ (circles), $\tau_{av} = 0.5$ (diamonds), and $\tau_{av} = 1.0$ (squares).

sensor-based simulation is perhaps the most flexible: it allows one to include different types of controllers (both reactive and nonreactive, homogeneous and heterogeneous systems) and easily incorporate specific environmental constraints (e.g., nonuniform stick distribution, special arena shapes, etc). Because the sensor-based simulation attempts to reproduce faithfully the environment and the robots' imperfect interactions with it, its results are most easily linked to the physical system. However, using this type of simulator requires a substantial investment in time—from implementing the logic of the controllers, to running the simulations. Also, the bigger the group size, the bigger the computational resources required to produce results. It is, therefore, impractical to use this model to study very large systems.

The microscopic numerical model does not require the computation of the details of the robots' trajectories; therefore, they run much faster than the sensor-based simulations and require fewer computational resources. The microscopic model can be adapted to different experiments with relative ease and it can deliver quantitatively accurate data. The macroscopic mathematical model is slightly more difficult to implement but very fast—unlike the microscopic model or sensor-based simulations, the time required to obtain results is independent of the robot group size. Using the macroscopic model, one can often study the system analytically, obtain expressions for many parameters of interest, and estimate the desired values using these expressions. However, the predictions of this type of model are sometimes only qualitatively correct, at least for small groups of robots such as those presented in this article. In addition, heterogeneous robot systems are easily studied using microscopic models since individuals are not summarized in a single caste. Each caste would require a different set of equations in a macroscopic model. Particular spatial or temporal probability distributions are more easily introduced in microscopic models. The macroscopic model is a deterministic model: given the same initial conditions, the same solution will always be reached. If it is required to know what the variation or the noise envelope in the performance of the system is, the probabilistic microscopic model is a better candidate. All three levels of abstraction are complementary to one another and can be used alone or together to gain insight into the behavior of multirobot systems.

In addition to the case of a homogeneous system of noncommunicating robots, Ijspeert et al. [14] studied, in simulation and using the probabilistic model, the cases of communicating as well as heterogeneous noncommunicating robot systems. In the future, we would like to expand the analytical model to include these cases. Introducing communication is perhaps the easier of the tasks. Ijspeert et al. [14] describe a simple signaling scheme in which a gripping robot emits a continuous signal ("call for help"), and searching robots within the hearing distance move toward the source of the signal. This simple scheme can be treated by the mathematical model by introducing two dynamic variables: signaling robots and signal-following robots. Just such a model of interacting foraging robots was studied by Sugawara and Sano [28]. Constructing a mathematical model of a heterogeneous system of robots is more challenging. It is an important task, however, because it is difficult to imagine that in practice, multirobot systems will be composed of identical robots. One approach is to treat a heterogeneous system as a collection of several homogeneous populations, or castes, of robots. Each population would be described by a set of equations like the ones presented in this article, but each with a different set of parameters (in the experimental article [14], the robots were differentiated by their gripping time parameter), and possibly new terms to describe interactions between populations. Though this approach may appear simplistic, it has been used successfully in population dynamics, for example, to describe predator–prey systems [10]. If both tasks and robots are heterogeneous, more complex coordination strategies will be required. For instance, market-based approaches offer simple analyzable distributed coordination strategies that may be used with robots [9].

We would like to test our approach by applying it to analyze other multirobot systems, including larger systems for which we could do a rigorous quantitative comparison between theoretical predictions and experimental results. Because of the practical difficulties involved in implementing a large multirobot system, a detailed comparison with embodied simulations may be more feasible. This is the approach taken by the studies of threshold-based algorithms for labor division [1, 16].

Another important challenge is to expand the model to allow learning. The type of models presented in this article apply to the simplest Markov-based systems. The next step is to generalize the model so that the robot's future state depends not only on the latest past state, but on the latest n past states. By introducing memory, we would allow robots to learn from past states and adapt to changing environmental conditions.

6 Conclusion

We have presented a macroscopic analytical model of collaboration in a homogeneous group of noncommunicating reactive robots. We first introduced a general model for the prediction of the collaboration dynamics for a task that requires r robots to be solved. We then validated the model in a specific case study: the stick-pulling experiment. The robots' task was to pull sticks out of their holes, and it could be successfully achieved only through a collaboration between two robots. Mathematical analysis reproduces the main qualitative conclusions of Ijspeert et al. [14], namely: the different dynamical regimes for different values of the ratio of robots to sticks (β), the optimal gripping time parameter for β less than the critical value, superlinearity of the group performance for small group sizes, as well as saturation and decrease in the relative collaboration rate as the size of the group grows. More significantly, these results were obtained without time-consuming simulations. In fact, some conclusions, such as the importance of the parameter β, fall directly out of simple analysis of the model, while others, such as an analytical expression for the optimal gripping time parameter, cannot be obtained without mathematical analysis. Another advantage of the macroscopic model is the ease of application. Once the macroscopic state diagram is drawn (from the details of the microscopic robot controller), the rate equations can be written down directly from it and numerically solved using available packages, such as Mathematica or Matlab, or by implementing algorithms from [27]. For example, starting from the model without interference, it took one of us (AG) one day to implement the model with interference and obtain numeric and analytic results.

In the simple state-based model we studied, the robot's future state depends only on its present state (and on how much time it has spent in that state). While the reactive robots in the stick-pulling study clearly obey this Markov property, other systems composed of robots with memory, learning, or deliberative capabilities do not, and therefore, they cannot be described by the simple models presented here. As it is common to do, we made some simplifying assumptions to make the mathematical model tractable. The most important assumption was that of spatial uniformity—that is, we assume that robots and sticks were uniformly distributed in space. The spatial uniformity assumption is used to calculate how many sticks and gripping robots a searching robot will encounter. For instance, the rate at which a searching robot encounters sticks (and makes a transition to the gripping state) is proportional to the number of ungripped sticks in the arena, with the proportionality factor given by α. This is a reasonable assumption for robots, because the searching behavior will tend to smooth out any inhomogeneities in the robots' distribution; however, it is not a good description of systems in which the sticks are strongly localized in some area of the arena.

The mathematical approach presented here is very general and can be applied to other multi-agent systems. We have used it to study coalition formation in an electronic marketplace [19], platoon formation in traffic flow [8], and in a work in progress on foraging in a group of robots.

Acknowledgments

This research was supported by the National Science Foundation under grant no. 0074790, in part by the Defense Advanced Research Projects Agency (DARPA) under contract number F30602-00-2-0573, and in part by the Center for Neuromorphic Systems Engineering at the California Institute of Technology as part of the National Science Foundation Engineering Research Center Program under grant EEC-9402726. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of any of the above organizations or any person connected with them.

References

1. Agassounon, W., Martinoli, A., & Goodman, R. M. (2001). A scalable, distributed algorithm for allocating workers in embedded systems. In *Proceedings of the IEEE Conference on System, Man and Cybernetics SMC-01* (pp. 3367–3373). New York: IEEE Press.

2. Beckers, R., Holland, O. E., & Deneubourg, J. L. (1994). From local actions to global tasks: Stigmergy and collective robotics. In R. A. Brooks & P. Maes (Eds.), *Proceedings of the 4th International Workshop on the Synthesis and Simulation of Living Systems Artificial Life IV* (pp. 181–189). Cambridge, MA: MIT Press.

3. Beni, G., & Wang, J. (1989). Swarm intelligence. In *Proceedings of the Seventh Annual Meeting of the Robotics Society of Japan* (pp. 425–428). Tokyo: RSJ Press.

4. Billard, A., Ijspeert, A. J., & Martinoli, A. (1999). A multi-robot system for adaptive exploration of a fast changing environment: Probabilistic modelling and experimental study. *Connection Science, 11*, 359–379.

5. Boehringer, K., Brown, R., Donald, B., Jennings, J., & Rus, D. (1995). Distributed robotic manipulation: Experiments in minimalism. In O. Khatib & J. K. Salisbury (Eds.), *Proceedings of the Fourth International Symposium on Experimental Robotics* (pp. 11–25). Lecture Notes in Control and Information Sciences. New York: Springer.

6. Bonabeau, E., Dorigo, M., & Theraulaz, G. (1999). *Swarm intelligence: From natural to artificial systems*. SFI Studies in the Science of Complexity. Oxford: Oxford University Press.

7. Everett, H. R., Gilbreath, G. A., Heath-Pastore, T. A., & Laird, R. T. (1993). Coordinated control of multiple security robots. In *Mobile Robots VII* (Vol. 2058, pp. 292–305). Boston: SPIE.

8. Galstyan, A., & Lerman, K. (2001). A stochastic model of platoon formation in traffic flow. In J. Hendler (Ed.), *Proceedings of the TASK Workshop*. Published on CD-ROM by DARPA.

9. Gerkey, B. P., & Matarić, M. J. (2001). *Sold!: Market methods for multi-robot control*. (Tech. Rep. IRIS-01-399). Los Angeles: University of Southern California. See also http://iris.usc.edu/irislib/.

10. Haberman, R. (1998). *Mathematical models: Mechanical vibrations, population dynamics, and traffic flow*. Philadelphia: Society of Industrial and Applied Mathematics (SIAM).

11. Hayes, A. T., Martinoli, A., & Goodman, R. M. (2000). Comparing distributed exploration strategies with simulated and real autonomous robots. In L. E. Parker, G. Bekey, & J. Bahren (Eds.), *Proceedings of the Fifth International Symposium on Distributed Autonomous Robotic Systems DARS-00* (pp. 261–270). New York: Springer.

12. Hayes, A. T., Martinoli, A., & Goodman, R. M. (in press). Swarm robotic odor localization. In *Proceedings of the IEEE Conference on Intelligent Robots and Systems IROS-01*. New York: IEEE Press.

13. Holland, O., & Melhuish, C. (1999). Stigmergy, self-organization, and sorting in collective robotics. *Artificial Life, 5*, 173–202.

14. 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*, 149–171.

15. Johnson, P. J., & Bay, J. S. (1995). Distributed control of simulated autonomous mobile robot collectives in payload transportation. *Autonomous Robots, 2*, 43–63.

16. Krieger, M. J. B., & Billeter, J.-B. (2000). The call of duty: Self-organised task allocation in a population of up to twelve mobile robots. *Robotics and Autonomous Systems, 30*, 65–84.

17. Kube, C. R., & Bonabeau, E. (2000). Cooperative transport by ants and robots. *Robotics and Autonomous Systems, 30*, 85–101.

18. Lerman, K., & Galstyan, A. (2001). *A general methodology for mathematical analysis of multi-agent systems*. (Tech. Rep. ISI-TR-529). Los Angeles: University of Southern California, Information Sciences Institute. See also http://www.isi.edu/isi-technical-reports.html.

19. Lerman, K., & Shehory, O. (2000). Coalition formation for large-scale electronic markets. In E. Durfee (Ed.), *Proceedings of the International Conference on Multi-Agent Systems* (*ICMAS '2000*) (pp. 167–174). New York: IEEE Press.

20. Martinoli, A. (1999). *Swarm intelligence in autonomous collective robotics: From tools to the analysis and synthesis of distributed control strategies*. Unpublished doctoral dissertation. (Tech. Rep. No. 2069). Lausanne, École Polytechnique Fédérale de Lausanne.

21. Martinoli, A., Ijspeert, A. J., & Mondada, F. (1999). Understanding collective aggregation mechanisms: From probabilistic modelling to experiments with real robots. *Robotics and Autonomous Systems, 29*, 51–63.

22. Martinoli, A., & Mondada, F. (1995). Collective and cooperative group behaviors: Biologically inspired experiments in robotics. In O. Khatib & J. K. Salisbury (Eds.), *Proceedings of the Fourth International Symposium on Experimental Robotics ISER-95* (pp. 3–10). New York: Springer.

23. Matarić, M. (1994). *Interaction and intelligent behavior*. Unpublished doctoral dissertation, Department of Electrical Engineering and Computer Science, MIT, Cambridge, MA.

24. Matarić, M. J., Nilsson, M., & Simsarian, K. (1995). Cooperative multi-robot box pushing. In *Proceedings of the 1995 IEEE/RSJ International Conference on Intelligent Robots* (pp. 556–561). New York: IEEE Press.

25. Michel, O. (1998). Webots: Symbiosis between virtual and real mobile robots. In J.-C. Heudin (Ed.), *Proceedings of the First International Conference on Virtual Worlds* (pp. 254–263). New York: Springer. See also http://www.cyberbotics.com/webots/.

26. Parrish, J. K., & Hamner, W. M. (1997). *Animal groups in three dimensions*. Cambridge: Cambridge University Press.

27. Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. (1993). *Numerical recipes in C: The art of scientific computing* (2nd ed.). Cambridge: Cambridge University Press.

28. Sugawara, K., & Sano, M. (1997). Cooperative acceleration of task performance: Foraging behavior of interacting multi-robots system. *Physica D, 100*, 343–354.

29. Sugawara, K., Sano, M., Yoshihara, I., & Abe, K. (1998). Cooperative behavior of interacting robots. *Artificial Life and Robotics, 2*, 62–67.

30. Vaughan, R. T., Støy, K., Sukhatme, G. S., & Matarić, M. J. (2000). Blazing a trail: Insect-inspired resource transportation by a robot team. In *Proceedings of the 5th International Symposium on Distributed Autonomous Robotic Systems* (*DARS*) (pp. 111–120).

31. Wang, J., & Premvuti, S. (1995). Distributed traffic regulation and control for multiple autonomous mobile robots operating in discrete space. In *Proceedings of the IEEE International Conference on Robotics and Automation* (pp. 1619–1624). New York: IEEE Computer Science.