

Mathematical Analysis of Multi-Agent Systems*

Kristina Lerman Aram Galstyan Tad Hogg
USC Information Sciences Institute Hewlett-Packard Labs

Abstract

We review existing approaches to mathematical modeling and analysis of multi-agent systems in which complex collective behavior arises out of local interactions between many simple agents. Though the behavior of an individual agent can be considered to be stochastic and unpredictable, the collective behavior of such systems can have a simple probabilistic description. We show that a class of mathematical models that describe the dynamics of collective behavior of multi-agent systems can be written down from the details of the individual agent controller. The models are valid for Markov or memoryless agents, in which each agents future state depends only on its present state and not any of the past states. We illustrate the approach by analyzing in detail applications from the robotics domain: collaboration and foraging in groups of robots.

Keywords: multi-agent systems, robotics, mathematical analysis, stochastic systems

*The research reported here was supported in part by the Defense Advanced Research Projects Agency (DARPA) under contract number F30602-00-2-0573, in part by the National Science Foundation under Grant No. 0074790, and by the ISI/ISD Research Fund Award. 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.

1 Introduction

Distributed systems composed of large numbers of relatively simple autonomous agents are receiving increasing amount of attention in the Artificial Intelligence (AI), robotics and networking communities. Unlike complex deliberative agents, the subject of much of AI research of the past two decades, simple agents have no, or limited, capacity to reason about data, plan action or negotiate with other agents. Although individual agents are far less powerful than traditional deliberative agents, distributed multi-agent systems based on such simple agents offer several advantages over traditional approaches: specifically robustness, flexibility, and scalability. Simple agents are less likely to fail than more complex ones. If they do fail, they can be pulled out entirely or replaced without significantly impacting the overall performance of the system. They are, therefore, tolerant of agent error and failure. They are also highly scalable – increasing the number of agents or task size does not require changes in the agent control programs nor compromise the performance of the system. In systems using deliberative agents, on the other hand, the high communications and computational costs required to coordinate group behavior limit the size of the system to at most a few dozen agents. Larger versions of such systems require division into subgroups with limited and simplified interactions between the groups [63]. In many cases, these interacting subgroups can in turn be viewed abstractly as agents following relatively simple protocols, as, for example, in market-based approaches to multi-agent systems [14].

There is no central controller directing agents' behavior, rather, these multi-agent systems are *self-organizing*, meaning that constructive *collective* (macroscopic) behavior emerges from *individual* (microscopic) decisions agents make. In most cases, these decisions are based on purely local information (that comes from other agents as well as the environment). Self-organization is ubiquitous in nature — bacteria colonies, amoebas and social insects such as ants, bees, wasps, termites, among others — display interesting manifestations of this phenomenon. Indeed, in many of these systems while the individual and its behavior appear simple to an outside observer, the collective behavior of the colony can often be quite complex. The apparent success of these organisms has inspired computer scientists and engineers to design algorithms and distributed problem-solving systems modeled after them (*e.g.*, swarm intelligence [7, 12, 61] and biologically-inspired systems [52, 24, 35, 10]). Moreover, current developments in micromechanical systems (MEMS) [8] and proposals for coordinated behavior among microscopic robots [1, 21] will require agent control programs capable of scaling to extremely large numbers of agents. Such agents will encounter various microscopic environments, some with counterintuitive properties [59], making it difficult to design appropriate deliberative control programs. Furthermore, at least in their initial development, such machines are likely to face severe power, computational and communication limitations and hence require a focus on collective behavior from computationally simple agents.

The main difficulty in designing multi-agent systems (MAS) with desirable self-organized behavior is understanding the effect individual agent characteris-

tics have on the collective behavior of the system. In the past, few analysis tools have been available to researchers, and it is precisely the lack of such tools that has been a chief impediment to the wider deployment of biologically-inspired MAS. Another impediment has been the difficulty of building hardware required for large-scale experiments. Researchers had a choice of experiments with relatively few agents or simulation for studying behavior of a MAS. Experiments with real agents, *e.g.*, robots, allow them to observe MAS under real conditions; however, experiments are very costly and time consuming, and systematically varying individual agent parameters to study their effect on the group behavior is often impractical. Simulations, such as sensor-based simulations of robots [23, 53], attempt to realistically model the environment, the robots' imperfect sensing of and interactions with it. Though simulations are much faster and less costly than experiments, they suffer from many of the same limitations, namely, they are tedious and the results are not generalizable. Exhaustive scan of the entire 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.

Mathematical modeling and analysis offer an alternative to the time-consuming and costly experiments and simulations. Using mathematical analysis we can study dynamics of multi-agent systems, predict long term behavior of even very large systems, gain insight into system design: *e.g.*, what parameters determine group behavior and how individual agent characteristics affect the MAS. Additionally, mathematical analysis may be used to select parameters that optimize group performance, prevent instabilities, *etc.* Conversely, such analytical tools can also provide design guidelines for agent programs. Specifically, these tools rely on various simplifying approximations to the agent behaviors. By deliberately designing agents to closely match these approximations, the resulting collective behavior will correspond to the analytic predictions. In this case, the tools will give a good indication of how to optimize the design to achieve desired behaviors. As one example, in market-based systems [14], designing agents to satisfy the assumptions of purely competitive markets allows a simple analysis of resulting behaviors in terms of market equilibria. Of course, restricting the agent design choices to achieve a close correspondence with analytic tools may limit the performance of the system, but this may be a useful tradeoff to achieve a simpler understanding of overall system behavior. Thus an important question for applying mathematical analysis for multi-agent systems is identifying situations in which simple analytic tools give a useful approximation to the collective behavior.

Mathematical modeling and analysis of large-scale collective behaviors is being increasingly used outside of the physical sciences where it has had much success. It has been applied to ecology [25], epidemiology [18], social dynamics [27], artificial intelligence [30], and behavior of markets [29], to name just a few disciplines.

In this paper we survey existing work on mathematical modeling and analysis of artificial multi-agent systems. We also describe a methodology for creating

analytic models of collective behavior of a multi-agent system. This type of analysis is valid for systems composed of agents that obey the Markov property: where each agent’s future state depends only on its present state. Many of the currently implemented multi-agent systems, specifically, reactive and behavior-based robotics, satisfy this property. We illustrate the approaches on robotics problems, comparing theoretical predictions with experimental and simulations results whenever possible.

2 Mathematical Models

A mathematical model is an idealized representation of a process. Constructing a mathematical model proceeds incrementally. To be useful, the model must explicitly include the salient details of the process it describes so its predictions reasonably match the actual behaviors of interest. On the other hand, the model should also be as simple as possible, ideally to allow analytic treatment and identification of qualitatively important relationships between individual and system behaviors. The precise choice of model involves a tradeoff between accuracy in describing reality and ease of use in providing explanations of the behavior. In our analysis we will strive to construct the simplest mathematical model that captures all of the most important details of the multi-agent system we are trying to describe.

Mathematical models can generally be broken into two classes: *microscopic* and *macroscopic*. Microscopic descriptions treat the agent as the fundamental unit of the model. There are several variations of the microscopic approach, as described in the following section. Macroscopic models, on the other hand, directly describe the collective behavior of a group of agents. Such models are the focus of the present paper.

2.1 Microscopic Models

Microscopic models treat the individual agent as the fundamental unit of the model. These models describe the agent’s interactions with other agents and the environment. Solving or simulating a system composed of many such agents gives researchers an understanding of the global behavior of the system.

2.1.1 Equations of Motion Approach

A common method used by physicists to study a system consisting of multiple entities consists of writing down and solving equations of motion for each entity. This approach has been adapted by some to describe agent-based pattern forming systems [11], including behaviors exhibited by colonies of biological organisms, such as slime mold [60] and social insects [17]. Of particular relevance to the agents community is the work of Schweitzer and coworkers [62, 28] on the active walker model of trail formation by ants and people. Active walkers are randomly walking agents that can influence the environment (*e.g.*, by depositing pheromone), in addition to being influenced by it. Schweitzer *et al.*

proposed a microscopic model of the interaction of the ground potential (created by pheromone or pedestrians' footprints) with the equations of motion of active walkers.

For large systems, solving equations with many degrees of freedom is often impractical. In some cases, it may be possible to derive a macroscopic model with fewer degrees of freedom from the microscopic model. Helbing, Schweitzer and coworkers did this in a later work [28], where they derived a model that describes the behavior of subpopulations of active walkers. Although these models of trail formation may be faulted for not being biologically realistic, they reproduce trail-forming behavior of real ants, such as the ability to discover and link distributed food sources without *a priori* knowledge of their location. Such models may be especially useful to study pheromone-based trail formation and navigation in robots [70, 69].

The main disadvantages of microscopic models are their poor scaling properties and that it is not always easy or obvious how to write down the equations of motion of each agent. Even if a model can be written down, in most cases it will not be analytically tractable, and the solution will have to be simulated on a computer. By resorting to simulation, one loses much of the power of mathematical analysis.

2.1.2 Microscopic Simulations

Microscopic simulations, such as molecular dynamics [19], cellular automata [1, 72] and particle hopping models [13], are a popular tool for studying dynamics of large multi-agent systems. In these simulations, agents change state stochastically or depending on the state of their neighbors. The popular Game of Life is an example of cellular automata simulation. Another example of the microscopic approach is the probabilistic model developed by Martinoli and coworkers [45, 46, 35] to study collective behavior of a group of robots. Rather than compute the exact trajectories and sensory information of individual robots, Martinoli *et al.* model each robot's interactions with other robots and the environment as a series of stochastic events, with probabilities determined by simple geometric considerations. Running several series of stochastic events in parallel, one for each robot, allowed them to study the group behavior of the multi-robot system.

2.2 Macroscopic Models

A macroscopic description offers several advantages over the microscopic approach. It is more computationally efficient, because it uses many fewer variables. A macroscopic model can often be solved analytically, yielding important insights into the behavior of quantities of interest. The macroscopic descriptions also tend to be more universal, meaning the same mathematical description can be applied to other systems governed by the same abstract principles. At the heart of this argument is the concept of separation of scales, which holds that the details of microscopic interactions (among agents) are only relevant for com-

puting the values of the parameters of the macroscopic model. This idea has been used by physicists to construct a single model that describes the behavior of seemingly disparate systems, *e.g.*, pattern formation in convecting fluids and chemical reaction-diffusion systems [71]. This principle of systems consisting of nearly decomposable parts applies broadly not only to physical systems but also to naturally evolved systems, as found in biology and economics, and designed technological artifacts [15, 64, 63]. From the perspective of large-scale agent systems, this decomposition often arises from processing, sensory and communication limitations of the individual agents. In effect, these limits mean agents can only pay attention to a relatively small number of variables in the full system [30], and will generally communicate concise summaries of their states to other agents. Of course, the two description levels are related, and it may be possible in some cases to derive the parameters of the macroscopic model from microscopic theory.

Macroscopic models are very popular and have been successfully applied to a wide variety of problems in physics, chemistry, biology and the social sciences. In most of these applications, the microscopic behavior of individual entity (a Brownian particle in a volume of gas or an individual residing in US) is quite complex, often stochastic and unpredictable, 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 (volume of gas or population of US). Such macroscopic descriptions often have a very simple form and are analytically tractable. They are sometimes called phenomenological models, because they are not derived from microscopic theories. 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. Such a probabilistic approach is the basis for statistical physics.

Usually, the relative size of fluctuations in statistical systems decreases with the number of components. In these cases, the system is almost always found near its average behavior and so the average is a good description for most individual experiments. It is this observation that allows the convenient study of average properties to describe behaviors actually seen in most experiments. In some cases, fluctuations can become large, *e.g.*, near phase transitions in physical systems. Such behaviors are also seen in computational systems, particularly those involved with combinatorial search [31], where long-tailed distributions have typical behavior far from that of the average. Thus when using macroscopic models to determine behavior of averages, it is important to keep in mind the possibility that actual system behaviors could be far from average. Fortunately, in the context of multi-agent systems, such large fluctuations will require an unexpectedly large statistical correlation in agent behaviors, which is unlikely in situations in which the agents are fairly independent and each act on only a few aspects of the overall system state (*e.g.*, based on local sensory information).

2.2.1 Finite Difference Equations

A macroscopic model can be frequently written down as a finite difference equation describing the change in the value of a dynamic variable over some time interval Δt . For example, in a model of population dynamics of US,

$$\begin{aligned} N(t + \Delta t) &= N(t) + \Delta t R(t) N(t) \\ R(t) &= \frac{N(t + \Delta t) - N(t)}{\Delta t N(t)}, \end{aligned}$$

where $N(t)$ is the (time-dependent) US population, Δt is a decade used by the Census Bureau, and $R(t)$ is the rate of change of population due to births, deaths, immigration and emigration. In general, $R(t)$ will also depend on the choice of Δt . The modeler finds an appropriate R to describe population growth of US, and solves the equations to project population growth into the future.

This description provides a finite-difference equation for the behavior of N at integer multiples of Δt . This has been used to model a robotic system [48, 49], and is particularly appropriate for *synchronous* systems, *i.e.*, where all agents make decisions at the same time (such as parallel update cellular automata).

In the continuous limit ($\Delta t \rightarrow 0$ or large N), the difference equation becomes a differential equation, known as the rate equation. For the example above, the rate equation is $\frac{dN(t)}{dt} = R(t)N(t)$, which is also applicable to *asynchronous* systems. In many cases, the behavior of this differential equation matches that of the difference equation [6]. However, this is not always the case: synchronous and asynchronous systems can have very different collective behaviors [33]. Large scale agent systems interacting with an environment, *e.g.*, robots, will often need to respond to environmental signals that arrive at unpredictable times. Such systems are likely to be better viewed as asynchronous, for which the differential equation approach is most suited.

2.2.2 Rate Equations

An alternate way to derive the rate equation is to start with the master equation for a stochastic system and macroscopically average it to get the rate equation for the dynamics of average quantities. Section 3.2 presents a tutorial on this approach. However, in order to create a model of a multi-agent system, one does not need to start with the master equation — one can easily write down the rate equations by examining the details of individual agent controller.

The rate equations are deterministic. In stochastic systems, however, rate equations 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. Pacala *et al.*[57] showed that in models of task allocation in ants, the exact stochastic and the average deterministic models *quantitatively* agree in systems containing as few as ten ants. The agreement increases as the size of the system grows. Martinoli and

Easton [48] have shown quantitative agreement with simulations in a system of 16–24 robots.

The rate equation has been used to model dynamic processes in a wide variety of systems. The following is a short list of applications: in chemistry, it has been used to study chemical reactions [22]; in physics, the growth of semiconductor surfaces [4] among others; in ecology to study dynamics of populations [16], including predator-prey systems [26]; in biology to model the behavior of social insects [57, 68]. The rate equation has also found application in the social sciences [27] and in AI. Huberman, Hogg and coworkers [34, 37] mathematically studied collective behavior of a system of agents (they called a computational ecology), where each agent chooses between two alternative strategies. They start with underlying probability distributions and derive rate equations for the average numbers of agents using each strategy. In fact, the same equations can be written down by examining the macroscopic state diagram of the agents. Yet another application of the approach presented here is coalition formation in electronic marketplaces [42].

In the robotics domain, Sugawara and coworkers [65, 66] developed simple analytical models of cooperative foraging in groups of communicating and non-communicating robots. Kazadi *et al.* [36] study the general properties of multi-robot aggregation using phenomenological macroscopic models. Agassounon and Martinoli [2] present a model of aggregation in which the number of robots taking part in the clustering task is based on the division of labor mechanism in ants. Lerman *et al.* have analyzed collaborative [41] and foraging [40] behavior in groups of robots. Results from these works will be used to illustrate the modeling methodology described in this paper. The focus of the present paper is to show that there is a principled way to construct a macroscopic analytic model of collective dynamics of a MAS, and, more importantly, a practical “recipe” for creating such a model from the details of the microscopic agent controller.

3 Macroscopic Analytic Models

3.1 MAS as Stochastic Systems

The behavior of individual agents in a multi-agent system has many complex influences, even in a controlled laboratory setting. Agents are influenced by external forces, many of which may not be anticipated. For robots, external forces include friction, which may vary with the type of surface the robot is moving on, battery power, sound or light signals, *etc.* Even if all the forces are known in advance, the agents are still subject to random events: fluctuations in the environment, as well as noise in the robot’s sensors and actuators. Each agent will interact with other agents that are influenced by these and other events. In most cases it is difficult to predict the agents’ exact trajectories and thus know which agents will come in contact with one another. Finally, the agent designer can take advantage of the unpredictability and incorporate it

directly into the agent’s behavior. For example, the simplest effective policy for obstacle avoidance in a robot is for it to turn a random angle and move forward. In summary, the behavior of simple agents in a complicated environment is so complex, the MAS is best described probabilistically, as a stochastic system.

Before we present a methodology for mathematical analysis of stochastic systems, we need to define some terms. *State* labels a set of related agent behaviors required to accomplish a task. For example, when a robot is engaged in a foraging task, its goal is to collect objects, such as pucks, scattered around the arena and bring them to a home base. The foraging task can be thought of as consisting of the following high-level behavioral requirements [3] or states

Homing — return the puck to a home base after it is picked up (includes collision avoidance)

Pickup — if a puck is detected, close gripper

Searching — wander around the arena in search of pucks (includes collision avoidance)

Each of these high level states may consist of a single action or behavior, or a set of behaviors. For example, when a robot is in the *Searching* state, it is wandering around the arena, detecting objects and avoiding obstacles. In the course of accomplishing the task, the robot will transition from the *Searching* to *Pickup* and finally to *Homing* states. We define these states to be mutually exclusive, *i.e.*, so that each agent in a multi-agent system is in exactly one of a finite number of states during a sufficiently short time interval. Note that there can be one-to-one correspondence between agent actions/behaviors and states. However, in order to keep the mathematical model compact and tractable, it is useful to *coarse-grain* the system by choosing a smaller number of states, each incorporating a set of agent actions or behaviors. Such coarse-graining is particularly relevant when we are only interested in behaviors described at this coarser level of abstraction.

In general, the full description of an agent in its environment could involve an arbitrary amount of detail, *e.g.*, its exact location in the environment. While such continuous states could be included in the formalism we present here, for simplicity we take the possible states of interest to be a finite set. Even when continuous values, such as location, are relevant to particular applications, it may be sufficient to treat these values with just a few coarse-grained regions.

We associate a unit vector \hat{q}_k with each state $k = 1, 2, \dots, L$. The configuration of the system is defined by the occupation vector

$$\vec{n} = \sum_{k=1}^L n_k \hat{q}_k \quad (1)$$

where n_k is the number of agents in state k . The probability distribution $P(\vec{n}, t)$ is the probability the system is in configuration \vec{n} at time t .

3.2 The Stochastic Master Equation: A Tutorial

For systems that obey the Markov property, the future is determined only by the present and not by the past. Clearly, agents that plan or use memory of past actions to make decisions, will not meet this criterion directly. While it is always possible to include the contents of an agent's memory as part of its state to maintain the Markov property, this can result in a vast expansion in the number of states to consider.

Fortunately, many MAS studied by various researchers, specifically those based on reactive and behavior-based robots and many types of software agents and sensors, do satisfy the Markov property with a fairly modest number of states. We restate the Markov property in the following way: the configuration of a system at time $t + \Delta t$ depends only on the configuration of the system at time t . In terms of coarse-grained states, we require this property to apply to the system described at this level of abstraction, at least to sufficient approximation.

The Markov property allows us to rewrite the marginal probability density $P(\vec{n}, t + \Delta t)$ in terms of conditional probabilities for transition from \vec{n}' to \vec{n} :

$$P(\vec{n}, t + \Delta t) = \sum_{\vec{n}'} P(\vec{n}, t + \Delta t | \vec{n}', t) P(\vec{n}', t).$$

Using the fact that

$$\sum_{\vec{n}'} P(\vec{n}', t + \Delta t | \vec{n}, t) = 1,$$

allows us to write the change in probability density as

$$\begin{aligned} P(\vec{n}, t + \Delta t) - P(\vec{n}, t) &= \sum_{\vec{n}'} P(\vec{n}, t + \Delta t | \vec{n}', t) P(\vec{n}', t) \\ &\quad - \sum_{\vec{n}'} P(\vec{n}', t + \Delta t | \vec{n}, t) P(\vec{n}, t). \end{aligned} \quad (2)$$

In the continuum limit, as $\Delta t \rightarrow 0$, Eq. 2 becomes

$$\frac{\partial P(\vec{n}, t)}{\partial t} = \sum_{\vec{n}'} W(\vec{n} | \vec{n}'; t) P(\vec{n}', t) - \sum_{\vec{n}'} W(\vec{n}' | \vec{n}; t) P(\vec{n}, t), \quad (3)$$

with transition rates defined as

$$W(\vec{n} | \vec{n}'; t) = \lim_{\Delta t \rightarrow 0} \frac{P(\vec{n}, t + \Delta t | \vec{n}', t)}{\Delta t}, \quad (4)$$

provided this limit exists, *e.g.*, changes are not synchronized to a global clock which would make $P(\vec{n}, t + \Delta t | \vec{n}', t)$ equal to zero for all $\vec{n} \neq \vec{n}'$ when Δt is sufficiently small.

Equation 3 says that the configuration of the system is changed by transitions to and from states. It is known as the Master Equation and is used widely to study dynamics of stochastic systems in physics and chemistry [22], traffic

flow [44, 43] and sociodynamics [27], among others. The Master Equation also applies to semi-Markov processes in which the future configuration depends not only on the present configuration, but also on the time the system has spent in this configuration. As we will see in a later section, transition rates in these systems are time dependent, while in pure Markov systems, they are time-independent.

3.3 The Rate Equation

The Master equation (Eq. 3) fully determines the evolution of a stochastic system. Once the probability distribution $P(\vec{n}, t)$ is found, one can calculate the characteristics of the system, such as the average and the variance of the occupation numbers. The Master equation is almost always too complex to be analytically tractable. Fortunately we can vastly simplify the problem by working with the average occupation number, $\langle \vec{n} \rangle$ (the Rate Equation). To derive an equation for $\langle \vec{n} \rangle$, we multiply Eq. 3 by \vec{n} and take the sum over all configurations:

$$\begin{aligned}
\frac{\partial}{\partial t} \langle \vec{n} \rangle &\equiv \frac{\partial}{\partial t} \sum_{\vec{n}} \vec{n} P(\vec{n}, t) \\
&= \sum_{\vec{n}} \sum_{\vec{n}'} \vec{n} W(\vec{n} | \vec{n}'; t) P(\vec{n}', t) - \sum_{\vec{n}} \sum_{\vec{n}'} \vec{n} W(\vec{n}' | \vec{n}; t) P(\vec{n}, t) \\
&= \sum_{\vec{n}} \sum_{\vec{n}'} (\vec{n}' - \vec{n}) W(\vec{n}' | \vec{n}; t) P(\vec{n}, t) \\
&= \left\langle \sum_{\vec{n}'} (\vec{n}' - \vec{n}) W(\vec{n}' | \vec{n}; t) \right\rangle
\end{aligned} \tag{5}$$

where $\langle \dots \rangle$ stands for averaging over the distribution function $P(\vec{n}, t)$. The time-evolution of a particular occupation number is obtained from the vector equation Eq. 5 as

$$\frac{\partial}{\partial t} \langle n_k \rangle = \left\langle \sum_{\vec{n}'} (n'_k - n_k) W(\vec{n}' | \vec{n}; t) \right\rangle \tag{6}$$

Let us assume for simplicity that only individual transitions between states are allowed, *i.e.*, $W(\vec{n}' | \vec{n}; t) \neq 0$ only if $\vec{n}' - \vec{n} = \hat{q}_i - \hat{q}_j$, $i \neq j$, and let w_{ij} be the transition rate from state j to state i . This is appropriate for systems with asynchronous updates for the agents since in that case it is very unlikely that two agents would make a change at the same time. Note that in general w_{ij} may be a function of the occupation vector \vec{n} , $w_{ij} = w_{ij}(\vec{n})$. Define a matrix \mathbf{D} with off-diagonal elements w_{ij} and with diagonal elements $D_{ii} = -\sum_k w_{ki}$. Then we can rewrite Eq. 5 in a matrix form as

$$\frac{\partial}{\partial t} \langle \vec{n} \rangle = \langle \mathbf{D}(\vec{n}) \cdot \vec{n} \rangle \approx \mathbf{D}(\langle \vec{n} \rangle) \cdot \langle \vec{n} \rangle \tag{7}$$

where we have used so the called mean-field approximation $\langle F(\vec{n}) \rangle \approx F(\langle \vec{n} \rangle)$. The mean-field approximation is often used in statistical physics and is well

justified for unimodal and sharp distribution functions [55]. The average occupation numbers obey the following system of coupled linear equations

$$\frac{\partial}{\partial t} \langle n_k \rangle = \sum_j w_{jk}(\langle \vec{n} \rangle) \langle n_j \rangle - \langle n_k \rangle \sum_j w_{kj}(\langle \vec{n} \rangle) \quad (8)$$

The above equation is known as the Rate Equation. It has the following interpretation: occupation number n_k will increase in time (first term in Eq. 8) due to transitions from other states to state k , and it will decrease in time due to the transitions from the state k to other states (second term).

3.3.1 Transition rates

Finding an appropriate mathematical form for the transition rates is the main challenge in applying the rate equations to real systems. Usually, the transition is triggered when an agent encounters some stimulus — be it another agent in a particular state, an object, its location, *etc.* For simplicity, we will assume that agents and triggers are uniformly distributed in space (though we will consider systems where agents interact in space, it does not necessarily have to be physical space, but a network, the Web, *etc.*). The assumption of spatial uniformity may be reasonable for agents that randomly explore space (*e.g.*, searching behavior in robots tends to smooth out any inhomogeneities in the robots' initial distribution); however, it fails for systems that are strongly localized, for instance, where all the objects to be collected by robots are located in the center of the arena. In these anomalous cases, the transition rates will have a more complicated form and in some cases it may not be possible to express them analytically altogether. If the transition rates cannot be calculated from first principles, it may be expedient to leave them as parameters in the model and estimate them by fitting the model to data.

3.3.2 Rate Equation and Multi-Agent Systems

The rate equation is a useful tool for mathematical analysis of macroscopic, or collective, dynamics of many agent-based systems. To facilitate the analysis, we begin by drawing the macroscopic state diagram of the system. The state diagram can be constructed from the details of the individual agent's behavior, as will be illustrated in the applications. Not every microscopic, or individual agent, behavior need become a macroscopic state. In order to keep the model tractable, it is often useful to coarse-grain the system by considering several related behaviors as a single state. As an example considered in detail below, 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. In addition to states, we must also specify transitions between states. These will be represented as arrows leading from one state to another.

Each state in the macroscopic state diagram corresponds to a dynamic variable in the mathematical model — the average number of agents in that state — and it is coupled to other variables via transitions between states. The mathematical model will consist of a series of coupled rate equations, one for each state, which describes how the number of agents in that state changes in time. This quantity may increase due to *incoming* transitions from other states, and it may decrease due to *outgoing* transitions to other states. Every transition will be accounted for by a term in each equation, with transition rates specified by the details of the interactions between agents.

In the next sections we will illustrate the details of the approach by applying it to study problems in the robotics domain. Our examples include foraging (Section 4) and collaboration (Section 5) in groups of robots.

4 Foraging in a Group of Robots

Robot collection and foraging are two of the oldest and most studied problems in robotics. In these tasks a single robot or a group of robots has to collect objects scattered around the arena and to assemble them either in some random location (collection task [5, 46]) or a pre-specified “home” location (foraging task [50, 24, 54]). These tasks have been studied under a wide variety of conditions and architectures, both experimentally and in simulation: in homogeneous and heterogeneous [24] systems, using behavior-based [50, 24] and hybrid control [54], no communication [24], direct communication [54, 65], as well as indirect communication through the environment [32]. The broad appeal of this problem is explained both by ubiquity of collection in general and foraging in particular in nature — as seen in the food gathering behavior of many insects — as well as its relevance to many military and industrial applications, such as de-mining, mapping and toxic waste clean-up. Foraging has been a testbed for the design of physical robots and their controllers, as well as a framework for exploring many issues in the design and implementation of multi-robot teams.

In this section, we focus on analysis of foraging in a homogeneous non-communicating multi-robot systems using behavior-based control, the type of systems studied by Mataric and collaborators [50, 24]. In Sec. 4.2 we will present Sugawara *et al.*'s model of foraging in communicating robots. Figure 1 is a snapshot of a typical experiment with four robots. The robots' task is to collect small pucks randomly scattered around the arena. The arena itself is divided into a search region and a small “home”, or goal, region where the collected pucks are deposited. The “boundary” and “buffer” regions are part of the home region and are made necessary by limitations in the robots' sensing capabilities, as described below. Each robot has an identical set of behaviors governed by the same controller. The behaviors that arise in the collection task are [24]:

Avoiding obstacles, including other robots and boundaries. This behavior is critical to the safety of the robot.

Searching for pucks: robot moves forward and at random intervals turns left




Figure 1: Diagram of the foraging arena (courtesy of D. Goldberg).

or right through a random arc. If the robot enters the Boundary region, it returns to the search region. This prevents the robot from collecting pucks that have already been delivered.

Detecting a puck.

Grabbing a puck.

Homing if carrying a puck, move towards the home location.

Creeping activated by entering Buffer region. The robot will start using the close-range detectors at this point to avoid the boundaries.

Home robot drops the puck. This activates the exiting behavior.

Exiting robot exits the home region and resumes search.

Interference, caused by competition for space between spatially extended robots, has long been recognized as a critical issue in multi-robot systems [20, 65]. When two robots find themselves within sensing distance of one another, they will execute obstacle avoiding maneuvers in order to reduce the risk of a potentially damaging collision. The robot stops, makes a random angle turn and moves forward. This behavior takes time to execute; therefore, avoidance increases the time it takes the robot to find pucks and deliver them home. Clearly, a single robot working alone will not experience interference from other robots. However, if a single robot fails, as is likely in a dynamic, hostile environment, the collection task will not be completed. A group of robots, on the other hand, is robust to an individual's failure. Indeed, many robots may fail but the performance of the group may be only moderately affected. Many

robots working in parallel may also speed up the collection task. Of course, the larger the group, the greater the degree of interference — in the extreme case of a crowded arena, robots will spend all their time avoiding other robots and will not bring any pucks home.

Several approaches to minimize interference have been explored, including communication [58] and cooperative strategies such as trail formation [69] and bucket brigade [20, 56]. In some cases, the effectiveness of the strategy to minimize interference will also depend on the group size [56]. Therefore, it is important to quantitatively understand interference between robots and how it relates to the group and task sizes before choosing alternatives to the default strategy. For some tasks and a given controller, there may exist an optimal group size that maximizes the performance of the system [54, 20, 56]. Beyond this size the adverse effects of interference become more important than the benefits of increased robustness and parallelism, and it may become beneficial to choose an alternate foraging strategy. We will study interference mathematically and attempt to answer these questions.

Nitz *et al.*[54] briefly addressed the question of what is an appropriate number of robots for a foraging task in a given environment. By simulating foraging in groups of up to five communicating robots, they observed an increase in performance when adding one to three robots as compared to a single worker. However, the performance seemed to level out and even degrade with further additions. Performance of non-communicating robots seemed to improve as the group size grew, at least up to the group size of five. No simulations for larger group sizes were carried out.

4.1 Rate Equation Model of Foraging

As mentioned above, interference is the result of competition between two or more robots for the same resource, be it physical space, the puck both are trying to pick up, energy, communications channel, *etc.* In the foraging task, competition for physical space, and the resulting avoidance of collisions with other robots, is the most common source of interference. In this section we examine the foraging scenario where robots are required to collect pucks and bring them to a specified “home” location.

At a macroscopic level, during some short time interval, every robot can be considered to be in the searching, homing or avoiding states, as shown in Fig. 2. We assume that actions like detecting and grabbing a puck take place on a sufficiently short time scale that they can be incorporated into the search state. Likewise, creeping, *etc.* can be incorporated into the homing state.¹ Initially the robots are in the search state. When a robot finds a puck, it picks it up and moves to the “home” region. Execution of the homing behavior requires a period of time τ_h . At the end of this period, the robot deposits the puck at

¹If we find that the given descriptive level does not adequately capture the behavior of a real or simulated system, we can consider more states in the model. For now, we are interested in the minimal model that reproduces salient features of the foraging system.

home and resumes search for more pucks. While the robot is homing, it will encounter and try to avoid other robots.




Figure 2: State diagram of a multi-robot foraging system with homing.

Each state in the diagram corresponds to a dynamic variable. Let $N_s(t)$, $N_h(t)$, $N_{av}^s(t)$, $N_{av}^h(t)$ be the number of searching, homing, avoiding while searching and avoiding while homing robots at time t , with the total number of robots, $N_0 = N_s(t) + N_h(t) + N_{av}^s(t) + N_{av}^h(t)$, a constant. We model the environment by letting $M(t)$ be the number of undelivered pucks at time t . Also, let α_r be the rate of detecting another robot and α_p the rate of detecting a puck. These parameters connect the model to the experiment, and they are related to the size of the robot and the puck, robot's detection radius and the speed of the robot. It was shown experimentally [24] that interference is most pronounced near the home region, because the density of robots is, on average, greater there. Therefore, we expect the rate of encountering other robots to be greater near the home region and introduce α'_r , the rate of detecting another robot while homing. The following equations describe the time evolution of the dynamic variables²:

$$\begin{aligned} \frac{dN_s(t)}{dt} &= -\alpha_p N_s(t)[M(t) - N_h(t) - N_{av}^h(t)] \\ &\quad -\alpha_r N_s(t)[N_s(t) + N_0] \\ &\quad + \frac{1}{\tau_h} N_h(t) + \frac{1}{\tau} N_{av}^s(t), \end{aligned} \quad (9)$$

$$\begin{aligned} \frac{dN_h(t)}{dt} &= \alpha_p N_s(t)[M(t) - N_h(t) - N_{av}^h(t)] \\ &\quad -\alpha'_r N_h(t)[N_h(t) + N_0] \\ &\quad + \frac{1}{\tau} N_{av}^h(t) - \frac{1}{\tau_h} N_h(t), \end{aligned} \quad (10)$$

$$\frac{dN_{av}^h(t)}{dt} = \alpha'_r N_h(t)[N_h(t) + N_0] - \frac{1}{\tau} N_{av}^h(t), \quad (11)$$

$$\frac{dM(t)}{dt} = -\frac{1}{\tau_h} N_h(t). \quad (12)$$

The first term in Eq. 9 accounts for a decrease in the number of searching robots when robots find pucks and start homing. The second term says that

²For simplicity, we do not include wall avoidance in the equations, but do take it into account when fitting model to the data.

the number of searching robots decreases when two searching robots detect each other and commence avoiding maneuvers or when a searching robot detects another robot in any of the remaining states. The number of available pucks is just the number of pucks in the arena less the pucks held by the homing robots. The last two terms in the equation require more explanation. We assume that it takes on average τ_h time for a robot to reach home after grabbing a puck. Then the average number of robots that deliver pucks during a short time interval dt and return to the searching state can be approximated as dtN_h/τ_h . Likewise, after a period of time τ , dtN_{av}^s/τ robots leave the avoiding state and resume searching. We assume that interference will increase the homing time for each robot: if τ_h^0 is the average homing time in the absence of collisions with other robots, then the effective homing time can be estimated as $\tau_h = \tau_h^0[1 + \alpha'_r \tau N_0]$.

The remaining equations have similar interpretations. We can take advantage of the conservation of the total number of robots to compute $N_{av}^h(t)$. These equations are solved numerically under the conditions that initially, at $t = 0$, there are M_0 pucks and N_0 searching robots.

Figure 3 shows the time evolution of the fraction of searching robots and pucks for $M_0 = 20$, $N_0 = 5$, $\tau = 3$ s, $\tau_h^0 = 16$ s. The number of searching robots (solid line) first quickly decreases as robots find pucks and carry them home, but then it increases and saturates at some steady state value as the number of undelivered pucks approaches zero (dashed line). The fraction of the searching robots in the steady state depends on the avoiding time parameter, which determines the fraction of robots in the avoiding state.




Figure 3: Time evolution of the fraction of searching robots (solid line) and undelivered pucks (dashed line) for $\tau = 3$ s, $\alpha_p = 0.015$, $\alpha_r = 0.04$, and $\alpha'_r = 0.08$.

To validate results of the model, we ran foraging simulations using Player/Stage simulator [23] for groups of one to ten robots and twenty pucks randomly scattered around the arena. Details of the simulations are presented in [40].

Figure 4(a) shows the total time required to complete the task for two differ-

ent interference strengths, as measured by the avoiding time parameter τ . The solid line is the result of the model's prediction for $\tau = 3$ s, and the dotted line for $\tau = 1.5$ s³, and $\tau_h^0 = 16$ s, $\alpha_p = 0.015$, $\alpha_r = 0.04$, and $\alpha'_r = 0.08$. The simulations data shows that the average avoiding time per collision increases as the group size grows. This is due to multiple avoidance moves per each attempt to avoid collision, caused by an increase in the local density of robots. Therefore, in the model, we take the avoiding time parameter τ as a linearly increasing function of N_0 , with the initial value of $\tau^0 = 3$ s (or 1.5 s). The agreement between the model and simulations is good.

The final plot (Fig. 4(b)) shows that interference causes deterioration in relative performance. The per robot efficiency is a monotonically decreasing function of the group size. Thus, adding one new robot to the group decreases the performance of every robot, though if the initial group size was less than the optimal size, adding a robot will increase the overall efficiency of the group.

4.2 Foraging in Communicating Robots

Sugawara and coworkers [65, 66, 67] carried out quantitative studies of foraging in groups of communicating robots in different environments. In their system when a robot finds a puck, it broadcasts a signal for a duration of time x . If other robots detect the signal, they turn and move towards it. After the interaction period, the broadcasting robot turns off the signal, and makes a transition to the homing state. The macroscopic state diagram for this system is shown in Fig. 5 and is composed of the following states: searching (S), broadcasting (B), homing (H), moving to the signal source (M), and avoiding (A), with the corresponding dynamic variables representing the average number of robots in each state.

$$\frac{dN_h}{dt} = \frac{1}{(x+1)}N_b - \frac{1}{\tau}N_h \quad (13)$$

$$\frac{dN_b}{dt} = -\frac{1}{(x+1)}N_b + \alpha N_s + \frac{\gamma}{a+N_a}N_a \quad (14)$$

$$\frac{dN_s}{dt} = -\alpha N_s + nN_m + \frac{1}{\tau}N_h - al(x)N_sN_b \quad (15)$$

$$\frac{dN_a}{dt} = \frac{v}{d}N_m - \frac{\gamma}{a+N_a}N_a \quad (16)$$

$$\frac{dN_m}{dt} = -\frac{v}{d}N_m - bN_m + al(x)N_sN_b \quad (17)$$

where α is the probability to find a puck, b probability to lose signal source, τ time to return home, x interaction duration, a is the probability of catching the broadcast signal, $l(x)$ turn angle, d average distance between interacting robots, v velocity of the robot, and γ probability to find a puck by following other robots.

³Although in the simulations we specified the avoidance time to be 1 s, multiple collisions caused the average avoiding time per collision to be slightly higher.

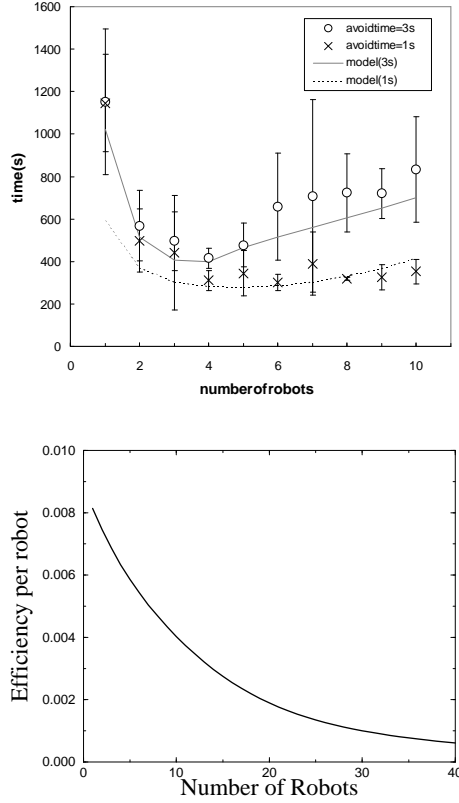


Figure 4: The time it takes the group to collect and deliver pucks home for two interference strengths, $\tau = 3 s$, and $\tau = 1.5s$ and $\tau_h^0 = 16 s$, $\alpha_p = 0.015$, $\alpha_r = 0.04$, $\alpha'_r = 0.08$. (b) Efficiency per robot




Figure 5: State diagram of foraging robots from Sugawara *et al.* [65]

Interference between robots due to collision avoidance is assumed to be negligible except during crowding near a broadcasting robot; therefore, avoiding terms appear only in the equations describing broadcasting and moving robots. The strength of interference is phenomenologically described by a simple function that is inversely proportional to the density of robots.

Sugawara *et al.* found that for non-communicating robots ($x = 0$), the time to complete the task was proportional to the inverse of the number of robots — $T \approx N^{-1}$. This is true when robots work independently of one another. Interaction through signal broadcast improves the efficiency of group behavior — $T \approx N^\beta$, $-1 < \beta < -2$ — for most durations x of interaction. This result is independent of whether the puck distribution was homogeneous or localized (specified by parameter γ). These findings were confirmed by simulations and experiments with physical robots.

At first glance, the results of Sugawara *et al.* appear to contradict those presented in Fig. 4(a), which shows that performance time T decreases with group size N up to some critical value and then starts to increase. In fact, the results of both works support the same basic conclusion that avoiding needs to be included in the model when crowding conditions occur (the exact density threshold remains to be determined). We believe the apparent discrepancy in results can be explained by the difference in systems being modeled. In Sugawara *et al.* work, “home” is located at the center of the arena. Avoiding is not taken into account in the no-interaction model because there is less crowding near the home region than near a broadcasting robot. In the system we studied, on the other hand, “home” is located at the edge of the arena, and crowding is more pronounced. Both studies construct a minimal model required to explain the observed behavior of the system. The differences in the models and conclusions can be traced back to the differences in the systems being modeled and their behavior.

5 Collaboration in Robots

Collaboration can significantly increase the performance of a multi-agent system. In some systems collaboration is an explicit requirement, because a single agent cannot successfully complete the task on its own. Such “strictly collaborative” [45] systems are common in insect and human societies, *e.g.*, in transport of an object too heavy or awkward to be lifted by a single ant, flying the space shuttle, playing a soccer match, *etc.* Collaboration in a group of robots has been studied by several groups [51, 39, 47, 35, 38]. We will focus on one group of experiments initiated by Martinoli and collaborators [47] and studied by Ijspeert *et al.* [35] that take a swarm approach to collaboration. In this system collaboration in a homogeneous group of simple reactive agents was achieved entirely through local interactions, *i.e.*, without explicit communication or coordination among the robots. Because they take a purely swarm approach, their system is a compelling and effective model of how collaboration may arise in natural systems, such as insect societies.

5.1 Stick-pulling Experiments in Groups of Robots

The stick-pulling experiments were carried out by Ijspeert *et al.* to investigate the dynamics of collaboration among locally interacting simple reactive robots. Figure 6 is a snapshot of the physical set-up of the experiments. 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 (pull the stick out) on its own — a collaboration between two robots is necessary for the task to be successfully completed. In a more general case, a collaboration between an arbitrary number of robots may be required to successfully complete the tasks (because sticks may be of varying length). The collaboration occurs in the following way: one robot finds a stick, and waits for a second robot to find it, lifting the stick partially out of its hole. When a second robot finds it, it will grip the stick and pull it out of the ground, successfully completing the task. (In the general case, a group of some size has to accumulate at the site of the stick before the required number of robots necessary to complete the task is present.)




Figure 6: Physical set-up of the stick-pulling experiment showing six Khepera robots (courtesy of A. Martinoli).

The actions of each robot are governed by a simple controller, outlined in Figure 7. 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 half way 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, 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.

Ijspeert *et al.* studied the dynamics of collaboration in stick-pulling robots

on three different levels: by conducting experiments with physical robots; with a sensor-based simulator of robots; and using a probabilistic microscopic model. The physical experiments were performed with 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 [53], was used to systematically explore the parameters affecting the dynamics of collaboration. Webots simulator attempts to faithfully replicate the physical experiment by reproducing the robots’ (noisy) sensory input and the (noisy) response of the on-board actuators in order to compute the trajectory and interactions of each robot in the arena. The probabilistic microscopic model, on the other hand, does not attempt to compute the trajectories of individual robots. Rather, 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. 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, each representing a single 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.* 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 300 times faster than Webots for these experiments.

5.1.1 Experimental Results

Ijspeert *et al.* systematically studied the collaboration rate (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. They found very good qualitative and quantitative agreement between the three different levels of experiments, as shown in Figure 8. Their main observation was 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 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.* was that when there are fewer robots than sticks, there is an optimal value of the gripping time parameter which maximizes the collaboration rate. In the other regime, the collaboration rate appears to be independent of the gripping

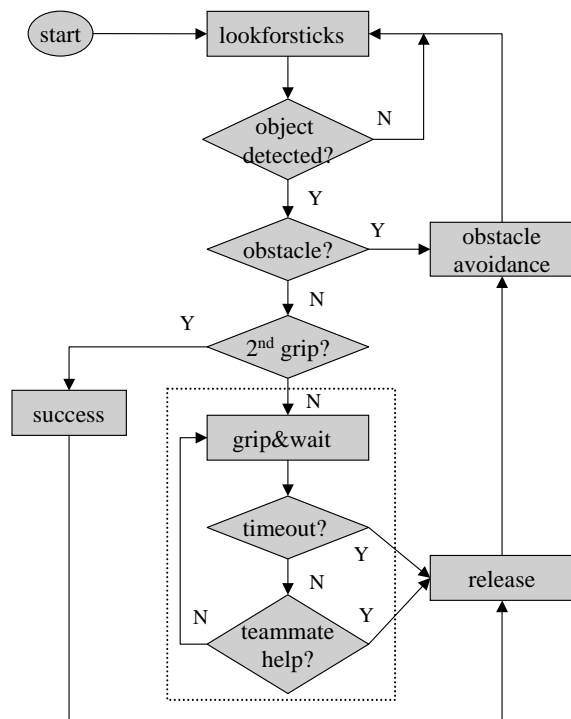


Figure 7: Flowchart of the robots' controller (from Ijspeert et al [35]).

time parameter above a specific value, so the optimal strategy is for the robot to grip a stick and hold it indefinitely.




Figure 8: Collaboration rate vs. the gripping time parameter for groups of two to six robots and four sticks (from Ijspeert *et al*). Heavy symbols represent experimental results, symbols connected by lines are the results of sensor-based simulations, while the smooth heavy lines are the results of the probabilistic microscopic model.

In the following section we present a macroscopic mathematical model of the stick-pulling experiments in a homogeneous multi-robot system. Such a model is useful for the following reasons. First, the model is independent of the system size, *i.e.* the number of robots; therefore, solutions for a system of 5,000 robots take just as long to obtain as solutions for 5 robots, whereas for a microscopic description the time required for simulation scales at least linearly with the number of robots. Second, our approach allows us to directly estimate certain important parameter values, (*e.g.*, those for which the performance is optimal) without having to resort to the time consuming simulations or experiments. 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.

5.2 The Rate Equation Model of Collaboration

In order to construct a mathematical model of collective behavior in stick-pulling experiments, it is helpful to draw the macroscopic state diagram of the system. On a macroscopic level, during a sufficiently short time interval, each robot will be in one of two states: *searching* or *gripping*. Using the flowchart of the robots' controller (Fig. 7) as reference, we include in the search state the set of behaviors associated with the looking for sticks mode, such as wandering around the arena ("look for sticks" action), detecting objects and avoiding obstacles;

while the gripping state is composed of decisions and an action inside the dotted box. We assume that actions “success” (pull the stick out completely) and “release” (release the stick) take place on a short enough time scale that they can be incorporated into the search state. Of course, there can be a discrete state corresponding to every action depicted in Fig. 7, but this would complicate the mathematical analysis without adding much to the descriptive power of the model. 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 studied the effect of avoidance in [41] and found that it does not qualitatively change the results of the simpler model that does not include avoidance; therefore, we will leave it out for clarity.

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 Fig. 9. These arrows correspond to the arrow entering and the two arrows leaving the dotted box in Fig. 7. We will use the macroscopic state diagram as the basis for writing down the rate equations that describe the dynamics of the stick-pulling experiments. Note that the experimental system was a semi-Markov system, because the robot’s transition from gripping to the searching state depended not only on its present state (gripping) but also on how long the robot has been in the gripping state.




Figure 9: Macroscopic state diagram of the multi-robot 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, *i.e.*, when the robots releases the stick without a successful collaboration taking place.

We can simplify analysis by considering a modified version of the robot controller presented in Fig. 7, where instead of waiting a specified period of time, each robot releases the stick with some probability per unit time. Note that this simplification restores the Markovian property of the system. We also construct and analyze a more complex model (Sec. 5.3) that explicitly includes the gripping time parameter. We show that the system based on the simplified controller produces qualitatively the same macroscopic behavior as the original system, which is modeled by the second and more complex model. Both systems are described by the same macroscopic state diagram and differ only in the

details of the transition rate between the gripping and searching states.

Each box in Fig. 9 corresponds to a macroscopic state and therefore requires a dynamic variable to describe it. Thus, the variables of our 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 uncollected sticks at time t . The latter variable does not represent a macroscopic state, rather it tracks the state of the environment. The mathematical model of the stick-pulling experiments consists of a series of coupled rate equations, each describing how the dynamic variables evolve in time:

$$\frac{dN_s}{dt} = -\alpha N_s(t) \left(M(t) - N_g(t) \right) + \tilde{\alpha} N_s(t) N_g(t) + \gamma N_g(t), \quad (18)$$

$$N_0 = N_s + N_g, \quad (19)$$

$$\frac{dM}{dt} = -\tilde{\alpha} N_s(t) N_g(t) + \mu(t), \quad (20)$$

where α , $\tilde{\alpha}$ are the rates at which a searching robot encounters a stick and a gripping robot respectively, γ is the rate at which robots release sticks and $\mu(t)$ is the rate at which new tasks are added. The parameters α , $\tilde{\alpha}$, and γ connect the model to the experiment: 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 Eq. 18 correspond to the three arrows between the states in Fig. 9. The first term accounts for the decrease in the number of searching robots because some robots find and grip sticks; the second term describes the successful collaborations between two robots, and the third term accounts for the failed collaborations (*i.e.*, when a robot releases a stick without a second robot present), both of which lead to an increase the number of searching robots. We do not need a separate equation for N_g , since this quantity may be calculated from the conservation of robots condition, Eq. 19. The last equation states that the number of sticks, $M(t)$, decreases in time at the rate of successful collaborations, Eq. 20. The equations are subject to the initial conditions that at $t = 0$ the number of searching robots is N_0 and the number of sticks is M_0 .

To proceed further, let us introduce

$$n(t) = N_s(t)/N_0, \quad (21)$$

$$m(t) = M(t)/M_0, \quad (22)$$

$$\beta = N_0/M_0, \quad (23)$$

$$R_G = \tilde{\alpha}/\alpha, \quad (24)$$

$$\tilde{\beta} = R_G \beta \quad (25)$$

and dimensionless time $t \rightarrow \alpha M_0 t$. Here $n(t)$ is the fraction of robots in the search state and $m(t)$ is the fraction of uncollected sticks at time t . Due to the conservation of number of robots, the fraction of robots in the gripping state is simply $1 - n(t)$. Equations 18–20 can be rewritten in dimensionless form as:

$$\frac{dn}{dt} = -n(t)[m(t) + \beta n(t) - \beta] + \tilde{\beta} n(t)[1 - n(t)] + \gamma[1 - n(t)] \quad (26)$$

$$\frac{dm}{dt} = -\beta\tilde{\beta}n(t)[1 - n(t)] + \mu' \quad (27)$$

Equations 26–27 together with initial conditions $n(0) = 1$, $m(0) = 1$ determine the dynamical evolution of the system. Note that only two parameters, β and γ , 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 paper we will use $R_G = 0.35$).⁴

We assume that new sticks are added to the system at the rate that the robots pull them out; therefore, the number of sticks does not change with time ($m(t) = m(0) = 1$). This situation may be realized experimentally by replacing the sticks in their holes after they are pulled out by robots. 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 Eq. 26 to zero:

$$-n[1 + \beta n - \beta] + \tilde{\beta}n[1 - n] + \gamma[1 - n] = 0. \quad (28)$$

This quadratic equation can be solved to obtain steady state values of $n(\beta, \gamma)$.

Figure 10(a) shows the dependence of the fraction of searching robots in the steady state on the parameters β and γ . The x-axis has units of time, although the scale is different than in Fig. 8. Note, that for small enough β 's $n(\gamma) \rightarrow 0$ as $1/\gamma \rightarrow \infty$. The intuitive reason for this is the same one given in Section 5.1.1: when there are fewer robots than sticks, and each robot holds the stick indefinitely (vanishing release probability), after a while every robot is holding a stick, and no robots are searching. For larger values of β , however, $n(\gamma) \rightarrow \text{const} \neq 0$ as $1/\gamma \rightarrow \infty$. Figure 10(b) shows how a typical solution $n(t)$ relaxes to its steady state value.

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

$$R(\gamma, \beta) = \beta\tilde{\beta}n(\gamma, \beta)[1 - n(\gamma, \beta)], \quad (29)$$

where $n(\gamma, \beta)$ is the steady-state number of searching robots for a particular value of γ and β , and $(1 - n(\gamma, \beta))$ is the steady-state number of gripping robots. Figure 11 depicts the collaboration rate as a function of $1/\gamma$. For $\beta > \beta_c$ the collaboration rate increases monotonically with $1/\gamma$. However, for $\beta < \beta_c$ there is an optimal stick release rate which maximizes the collaboration rate. The optimal value of γ which maximizes the collaboration rate can be computed

⁴The parameter α can be easily calculated from experimental values quoted in [35]. 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 $V_R = 8.0 \text{ cm/s}$ is robot's speed, and $W_R = 14.0 \text{ cm}$ is robot's detection width. If the arena radius is $R = 40.0 \text{ cm}$, a robot will detect sticks at the rate $\alpha = V_R W_R / \pi R^2 = 0.02 \text{ s}^{-1}$. According to [35], 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 10: (a) Steady state solution vs inverse stick release rate $1/\gamma$. (b) Typical relaxation to the steady state of the fraction of searching robots for $\gamma = 0.2$, $\beta = 0.5$.




Figure 11: Collaboration rate per robot vs inverse stick release rate $1/\gamma$ for $\beta = 0.5$, $\beta = 1.0$, $\beta = 1.5$. These values of β correspond, respectively, to two, four, and six robots in the experiments with four sticks (*cf.* Fig. 8).

from the condition $dR(\gamma, \beta)/d\gamma = \beta\tilde{\beta}d(n - n^2)/d\gamma = 0$, with n given by roots of Eq. 28. Another way to compute the optimal release rate is by noting that for a given value of β below some critical value, the collaboration rate is greatest when half of the robots are gripping and the other half are searching. Substituting $n = 1/2$ into Eq. 28, leads to $\gamma_{opt} = 1 - (\beta + \tilde{\beta})/2$. γ_{opt} vanishes as β exceeds its critical value, $\beta_c = 2/(1 + R_G)$; therefore, for $\beta > \beta_c$, $n > 1/2$, and no optimal release rate exists.

The three curves in Fig. 11 are qualitatively similar to those in Fig. 8 for 2 robots ($\beta = 0.5$), 4 robots ($\beta = 1.0$) and 6 robots ($\beta = 1.5$). Even the grossly simplified model reproduces the following conclusions of Ijspeert *et al.*: the different dynamical regimes depending on the value of the ratio of robots to sticks (β) and the optimal gripping time parameter for $\beta < \beta_c$.

In the next section, we construct a model of the stick pulling experiments that explicitly includes the gripping time parameter τ . The collective behavior predicted by the more accurate model is both qualitatively and quantitatively similar to that predicted by the simplified model.

5.3 Model with Gripping Time Parameter

A more accurate mathematical description of the stick pulling experiments should explicitly include the gripping time parameter τ . Note that the system is now a semi-Markov system, because transitions from gripping to the searching state depend not only on the present state (gripping) but also on how long the robot has been in the gripping state, *i.e.*, whether or not it has timed out. This property of the system will be captured by time-dependent transition rates. The system is described by the same macroscopic state diagram, Fig. 9, and the same set of equations, Eq. 18–20, with only the last term in Eq. 18 different. We write the new equation below:

$$\begin{aligned} \frac{dN_s}{dt} = & -\alpha N_s(t) \left(M(t) - N_g(t) \right) + \tilde{\alpha} N_s(t) N_g(t) \\ & + \alpha N_s(t - \tau) \left(M(t - \tau) - N_g(t - \tau) \right) \Gamma(t; \tau). \end{aligned} \quad (30)$$

All the parameters have the same meaning as before. $\Gamma(t; \tau)$, the fraction of failed collaborations at time t , is the probability no robot came “to help” during the time interval $[t - \tau, t]$. This is a time-dependent parameter, and it describes unsuccessful transitions from the gripping state in this semi-Markov system. To calculate $\Gamma(t; \tau)$ 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

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

$$\equiv \exp\left[\sum_{i=1}^K \ln[1 - \tilde{\alpha}\delta t N_s(t - \tau + i\delta t)]\right] \Theta(t - \tau) \quad (31)$$

The step function $\Theta(t - \tau)$ ensures that $\Gamma(t; \tau)$ is zero for $t < \tau$. Finally, expanding the logarithm in Eq. (31) and taking the limit $\delta t \rightarrow 0$ we obtain

$$\Gamma(t; \tau) = \exp\left[-\tilde{\alpha} \int_{t-\tau}^t dt' N_s(t')\right] \Theta(t - \tau) \quad (32)$$

We rewrite the model in dimensionless form using variable transformations given by Eqs. 21–25 and dimensionless gripping time parameter $\tau \rightarrow \alpha M_0 \tau$:

$$\begin{aligned} \frac{dn}{dt} &= -n(t)[m(t) + \beta n(t) - \beta] + \tilde{\beta} n(t)[1 - n(t)] \\ &\quad + n(t - \tau)[m(t - \tau) + \beta n(t - \tau) - \beta] \times \gamma(t; \tau) \end{aligned} \quad (33)$$

$$\frac{dm}{dt} = -\beta \tilde{\beta} n(t)[1 - n(t)] + \mu' \quad (34)$$

$$\gamma(t; \tau) = \exp\left[-\tilde{\beta} \int_{t-\tau}^t dt' n(t')\right] \quad (35)$$

Equations 33–35 are solved subject to initial conditions $n(0) = 1$ and $m(0) = 1$ to determine the dynamic evolution of the system. Parameters β and τ alone appear in the equations and thus, determine the behavior of the system.

Equation 33 has a non-trivial steady-state solutions which satisfy the following transcendental equation:

$$-1 + (\beta + \tilde{\beta})(1 - n) + (1 - \beta(1 - n))e^{-\tilde{\beta}\tau n} = 0 \quad (36)$$

Figure 12 shows the dependence of the steady state fraction of searching robots on the gripping time τ for different values of β . Note, that for small enough β 's $n(\tau) \rightarrow 0$ as $\tau \rightarrow \infty$, *i.e.*, 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, $n(\tau) \rightarrow \text{const} \neq 0$ as $\tau \rightarrow \infty$. The inset in Fig. 12 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 steady-state collaboration rate $R(\tau; \beta)$, the rate at which robots pull sticks out of their holes, is given by: $R(\tau, \beta) = \beta \tilde{\beta} n(\tau, \beta)[1 - n(\tau, \beta)]$ 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 13 depicts collaboration rate as a function of τ . For $\beta > \beta_c$ the collaboration rate increases monotonically with τ . However, for $\beta < \beta_c$ there is an optimal gripping time, $\tau = \tau_{opt}$, which maximizes the collaboration rate. We use the same arguments as before to understand this behavior: maximum collaboration rate for a given β is achieved for $n(\tau, \beta) = 1/2$. For $\beta > \beta_c$,




Figure 12: Steady state solution vs (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 13: Collaboration rate per robot vs (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.* Fig. 8).




Figure 14: Optimal gripping time parameter and inverse of the optimal release rate vs β in the two models

however, the solution of Eq. 36 is always greater than $1/2$, so an optimal solution does not exist. For $\beta < \beta_c$, simple analysis gives

$$\tau_{opt} = \frac{2}{\beta} \ln \frac{1 - \beta/2}{1 - 1/2(\beta + \tilde{\beta})}, \quad \beta < \beta_c = \frac{2}{1 + R_G} \quad (37)$$

This dependence on R_G was quantitatively confirmed through embodied and microscopic simulations [49]. Figure 14 compares the optimal gripping time parameter and inverse of the optimal release rate predicted by the two models. For $\beta < 1$ the two models give *quantitatively* similar results, in addition to predicting the same β_c . This example illustrates our claim that in many cases a minimal model is sufficient to explain and predict interesting system properties. Ref. [41] presents more details of the analysis of the collaboration task, including the effects of interference.

5.4 Difference Equation Model of Collaboration

Martinoli and Easton [48, 49] consider a more fine-grained macroscopic model than the one described above that takes into account more of the individual robot behaviors shown in Fig. 7. Their model consists of coupled finite difference equations, one for each state: searching (N_s), avoiding (N_a), interference (N_i), success dance (N_d), and gripping (N_g). The equation for how the number for searching robots changes in time is presented below; equations for other variables are similar.

$$\begin{aligned} N_S(k+1) &= N_s(k) - (\alpha_w + \alpha_r)N_s(k) - \tilde{\alpha}N_g(k)N_s(k) \\ &\quad - \alpha(M_0 - N_g(k))N_s(k) + \alpha_w N_s(k - T_a) \\ &\quad + \alpha_r N_s(k - T_{ia}) + \tilde{\alpha}N_g(k - T_{ca})N_s(k - T_{ca}) \\ &\quad + \tilde{\alpha}N_g(k - T_{cda})N_s(k - T_{cda}) \\ &\quad + \alpha(M_0 - N_g(k - T_{cda}))N_s(k - T_{cga})\Gamma(k, T_{ga}). \end{aligned}$$

Here, α_w and α_r are the rates at which robots encounter a wall or another robot. The current time step is k , $k = 0, 1, 2, \dots$, and $T_{xyz} = T_x + T_y + T_z$ are the number




Figure 15: Collaboration rate as a function of the gripping time parameter for several robot group sizes. Results are for embodied simulations (dotted lines), the microscopic model (dashed lines), and the macroscopic model (solid lines) in the 16 stick 80 cm radius arena. (Figure courtesy of A. Martinoli)

of time steps required to complete actions such as avoidance (T_a), success dance (T_d) or stick centering (T_c). $\Gamma(k, T_{ga}) = \prod_{j=k-T_{ga}}^k [1 - \tilde{\alpha} N_s(j) \Theta(k - T_{ga})]$ is the discrete time analog of Eq. 32.

Figure 15 shows collaboration rate as a function of the gripping time parameter for different robot group sizes. Collaboration rate is computed from the number of extracted sticks per unit time, as before: $C(k) = \tilde{\alpha} N_s(k - T_{cd}) N_g(k - T_{cd})$. We can see that for groups as small as 8 robots, the results of the macroscopic model quantitatively agree with embodied and microscopic simulations.

6 Limitations of this approach

The rate equation approach we presented is valid for Markov and semi-Markov systems, in which the agent's future state depends only on its present state and, for semi-Markov systems, on how much time it has spent in that state, and not on any of the past states. While many systems, including reactive and behavior-based robotics and some software agent systems, clearly obey the Markov property, other systems composed of agents with memory, learning or deliberative capabilities do not, and therefore, cannot be described by the simple models presented here. However, the rate equations are useful for studying systems of simple agents. Moreover, with some additional complexity, this approach can be extended to cover a broader range of agent behaviors. For instance, in simple learning scenarios, agents could adjust their behaviors based on the number of

times they enter certain states. For instance, to estimate the density of objects or other robots in their environment based on the number of encounters. Thus the transition rates would depend on the agent histories. Although the rate equations wouldn't model the learning process per se, it could include the effect on system behavior of improved discrimination ability. Such a model could indicate the likely tradeoff between exploring the environment to improve parameter estimates and exploit those estimates for the task at hand.

Another potential limitation of the approach is that it is best as a description of large systems. The benefit of working with large systems is that, usually, one may safely ignore fluctuations. Although it is possible to use the Master Equation to derive the equation for the fluctuations about the average quantities, in practice it is too algebraically messy. Fortunately, there exists empirical evidence [57, 49] that approximate average models can provide a good quantitative description of systems as small as a dozen agents. Large collections of very simple agents are also seen to be reasonably described by relations among a few aggregate variables [9].

Developing a suitable stochastic model for a multi-agent system requires an understanding of the environment and agent actions sufficient to determine the appropriate state description and resulting transition rates. More generally, this issue can be viewed as identifying an appropriate statistical ensemble, i.e., set of states with associated probabilities for their occurrence [30]. Even oversimplified models can be useful in giving qualitative understanding of the likely design tradeoffs prior to a more detailed evaluation via simulation or experiments. Nevertheless, it can be difficult to determine the transition rates between states, especially if they involve correlated activities among several agents. Along these lines is the question of to what extent a detailed model of the agent behaviors must be included. Complex agent programs could give rise to complex transition rates that make the analysis described here impractical. However, even in those cases, there may be some aggregate behaviors that can be approached by suitable coarse-graining. An extreme example is economics, which predicts various aggregate human behaviors without requiring detailed psychological models.

Beyond these general limitations of the stochastic approach described here, we introduced several simplifying assumptions. While these are not strictly necessary for the validity of the overall approach, they are important for producing analytically simple models. In particular we suppose the interesting system behavior is governed by averages and the transition rates are spatially uniform so it is not necessary to include position as part of the state. We also extensively use the mean-field approximation. In cases where this is not sufficient, more accurate approximations of our statistical models are possible [55], but are more difficult to evaluate.

When evaluating the suitability of stochastic models, in addition to these technical limitations, it is important to note the kinds of results the models can deliver. In particular, they address properties of the distribution of outcomes, *e.g.*, average and variance, over a set of repeated experiments. This is often appropriate for evaluating how a multi-agent system will likely perform for a

class of problems. However, if one is interested in worst-case bounds, the behavior in exceptional situations or extreme values of the distribution (*e.g.*, time required to find the first or last object in a search scenario), these stochastic techniques are unlikely to provide much insight unless the distributions happen to be limited to a narrow range. Thus the usefulness of these models depends not only on the complexity of the agent behaviors but also on the nature of the collective properties of interest.

7 Conclusion

We have presented an overview of mathematical approaches for modeling and analysis of multi-agent systems, focusing on macroscopic analytic models. Moreover, we have described a general methodology for mathematical analysis of such systems. Our analysis applies to a class of systems known as stochastic Markov systems. They are stochastic, because the behavior of each agent is inherently probabilistic in nature and unpredictable, and they are Markovian, because the state of an agent at a future time depends only on the present state (and perhaps on much time the agent has spent in this state) and not on any past state. Though each agent is unpredictable, the probabilistic description of the collective behavior is surprisingly simple. Our mathematical approach is based on the stochastic Master Equation, and on the Rate Equations derived from it, that describe how the average macroscopic, or collective, system properties change in time. In order to create a mathematical model, one needs to account for every relevant state of the multi-agent system as well as for transitions between states. For each state there is a dynamical variable in the mathematical model and a rate equation that describes how the variable changes in time.

We illustrated the mathematical formalism by applying it to study collective behavior of robotic systems. Even the simplest type of dimensional analysis of the equations yields important insights into the system, such as what are the important parameters that determine the behavior of the system.

In the applications, we focused on the simplest mathematical models (*i.e.*, those using the smallest possible number of states) that capture the salient features of each system. These simple models provide a good description of the *qualitative* behavior of the system, but in order to also obtain good *quantitative* agreement with experiment or simulations, it is often necessary to include more states in the model. The approach presented here can be easily extended to describe heterogeneous agent systems. As a simple example, consider two, possibly interacting, populations of foraging robots, each described by different physical parameters. The model of a heterogeneous foraging system will consist of two sets of coupled differential equations, one for each sub-population, possibly with couplings between them to represent interactions between the two populations. It is trivial to extend the analysis to more than two populations.

The models also provide design guidelines: choosing the agent behaviors to closely match the simplifying assumptions of this approach allows evaluating the resulting collective behavior via the rate equations.

Acknowledgements

The authors thank Dani Goldberg, Bernardo Huberman, Jeff Kephart, Alcherio Martinoli, Maja Matarić, Richard Ross and Onn Shehory for useful discussions.

References

- [1] Harold Abelson et al. Amorphous computing. *Communications of the ACM*, 43:74–82, May 2000.
- [2] William Agassounon and Alcherio Martinoli. A macroscopic model of an aggregation experiment using embodied agents in groups of time-varying sizes. In *Proc. of the IEEE Conf. on System, man and Cybernetics SMC-02, October 2002, Hammamet, Tunisia. To appear.* 2002.
- [3] Ronald C. Arkin. *Behavior-Based Robotics*. The MIT Press, Cambridge, MA, 1999.
- [4] A.-L. Barabasi and H.E. Stanley. *Fractal Concepts in Surface Growth*. Cambridge University Press, 1995.
- [5] R. Beekers, O. E. Holland, and J. L. Deneubourg. From local actions to global tasks: Stigmergy and collective robotics. In Rodney A. Brooks and Pattie Maes, editors, *Proceedings of the 4th International Workshop on the Synthesis and Simulation of Living Systems ArtificialLifeIV*, pages 181–189, Cambridge, MA, USA, July 1994. MIT Press.
- [6] Carl M. Bender and Steven A. Orszag. *Advanced Mathematical Methods for Scientists and Engineers*. McGraw Hill, NY, 1978.
- [7] G. Beni. The concept of cellular robotics. In *Proc. of the 1988 IEEE Int. Symp. on Intelligent Control*, pages 57–62, Los Alamitos, CA, 1988. IEEE Computer Society Press.
- [8] Andrew A. Berlin and Kaigham J. Gabriel. Distributed MEMS: New challenges for computation. *Computational Science and Engineering*, 4(1):12–16, January-March 1997.
- [9] Karl F. Bohringer et al. Computational methods for design and control of MEMS micromanipulator arrays. *Computational Science and Engineering*, 4(1):17–29, January-March 1997.
- [10] Hristo Bojinov, Arancha Casal, and Tad Hogg. Multiagent control of modular self-reconfigurable robots. *Artificial Intelligence*, 142:99–120, 2002. Preprint available at Los Alamos archive cs.RO/0006030.
- [11] Eric Bonabeau. From classical models of morphogenesis to agent-based models of pattern formation. *Artificial Life*, 3:191–211, 1997.

- [12] Eric Bonabeau, Marco Dorigo, and Guy Theraulaz. *Swarm Intelligence: From Natural to Artificial Systems*. Oxford University Press, New York, 1999.
- [13] D. Chowdhury, L. Santen, and A. Schadschneider. Statistical physics of vehicular traffic and some related systems. *Physics Reports*, 329:199, 2000.
- [14] Scott H. Clearwater, editor. *Market-Based Control: A Paradigm for Distributed Resource Allocation*. World Scientific, Singapore, 1996.
- [15] P. J. Courtois. On time and space decomposition of complex structures. *Communications of the ACM*, 28(6):590–603, June 1985.
- [16] J. M. Cushing. *An Introduction to Structured Population Dynamics*, volume 71 of *CBMS-NSF Regional Conference Series in Applied Mathematics*. Society for Applied and Industrial Mathematics, Philadelphia, 1998.
- [17] Jean-Luis Deneubourg, Guy Theraulaz, and R. Beckers. Swarm-made architectures. in: Toward a practice of autonomous systems. In F.J. Varela and P. Bourgine, editors, *Proceedings of The First European Conference on Artificial Life*, pages 123–133, Cambridge, MA, 1992. MIT Press/Bradford Books.
- [18] O. Diekmann and J. A. P. Heesterbeek. *Mathematical Epidemiology of Infectious Diseases : Model Building, Analysis and Interpretation*. John Wiley & Sons, New York, mathematical and computational biology edition, April 2000.
- [19] R. Duncan et al. Statistical paradigms for robotic swarm modeling. <http://www.mhpcc.edu/research/ab98/98ab41.html>, 1998.
- [20] Miguel Schneider Fontan and Maja J Matarić. A study of territoriality: The role of critical mass in adaptive task division. In P. Maes, M. J. Matarić, J. A. Meyer, J. Pollack, and S. Wilson, editors, *From Animals to Animats 4: Proceedings of the 4th International Conference on Simulation of Adaptive Behavior*, pages 553–561, Cambridge, MA, 1996. MIT Press.
- [21] Robert A. Freitas Jr. *Nanomedicine*, volume 1. Landes Bioscience, Georgetown, TX, 1999. Available at www.nanomedicine.com.
- [22] C. W. Garnier. *Handbook of Stochastic Methods*. Springer, New York, NY, 1983.
- [23] Brian P. Gerkey, Richard T. Vaughan, Kasper Sty, Andrew Howard, Gaurav S. Sukhatme, and Maja J Matarić. 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*. 2001. <http://www-robotics.usc.edu/player/>.

- [24] Dani Goldberg and Maja J Matarić. Robust behavior-based control for distributed multi-robot collection tasks. Technical Report IRIS-00-387, USC Institute for Robotics and Intelligent Systems, 2000.
- [25] W. S. C. Gurney and R. M. Nisbet. *Ecological Dynamics*. Oxford University Press, New York, NY, 1998.
- [26] Richard Haberman. *Mathematical Models: Mechanical Vibrations, Population Dynamics, and Traffic Flow*. Society of Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1998.
- [27] Dirk Helbing. *Quantitative Sociodynamics: Stochastic Methods and Models of Social Interaction Processes*, volume 31 of *THEORY AND DECISION LIBRARY B: Mathematical and Statistical Methods*. Kluwer Academic, Dordrecht, 1995.
- [28] Dirk Helbing, Frank Schweitzer, Joachim Keltsch, and Peter Molnar. Active walker model for the formation of human and animal trail systems. *Physical Review*, E 56(3):2527–2539, 1997.
- [29] J. Hirshleifer. Competition, cooperation, and conflict in economics and biology. *The American Economic Review*, 68(2):238–243, May 1978.
- [30] T. Hogg and B. A. Huberman. Artificial intelligence and large scale computation: A physics perspective. *Physics Reports*, 156:227–310, 1987.
- [31] Tad Hogg, Bernardo A. Huberman, and Colin P. Williams, editors. *Frontiers in Problem Solving: Phase Transitions and Complexity*, volume 81, Amsterdam, 1996. Elsevier. Special issue of *Artificial Intelligence*.
- [32] Owen Holland and Chris Melhuish. Stigmergy, self-organization and sorting in collective robotics. *Artificial Life*, 5(2), 2000.
- [33] Bernardo A. Huberman and Natalie S. Glance. Evolutionary games and computer simulations. *Proceedings of the National Academy of Science USA*, 90:7716–7718, August 1993.
- [34] Bernardo A. Huberman and Tad Hogg. The behavior of computational ecologies. In B. A. Huberman, editor, *The Ecology of Computation*, pages 77–115, Amsterdam, 1988. Elsevier (North-Holland).
- [35] A. J. Ijspeert, A. Martinoli, A. Billard, and L. M. Gambardella. Collaboration through the exploitation of local interactions in autonomous collective robotics: The stick pulling experiment. *Autonomous Robots*, 11(2):149–171, 2001.
- [36] Sanza Kazadi, A. Abdul-Khaliq, and Ron Goodman. On the convergence of puck clustering systems. *Robotics and Autonomous Systems*, 38(2):93–117, 2002.

- [37] Jeffrey O. Kephart, Tad Hogg, and Bernardo A. Huberman. Collective behavior of predictive agents. *Physica, D* 42:48–65, 1990.
- [38] H. Kitano, M. Tambe, P. Stone, and M. Veloso. The RoboCup synthetic agent challenge 97. *Lecture Notes in Computer Science*, 1395:62–??, 1998.
- [39] C. R. Kube and E. Bonabeau. Cooperative transport by ants and robots. *Robotics and Autonomous Systems*, 30(1–2):85–101, 2000.
- [40] Kristina Lerman and Aram Galstyan. Mathematical model of foraging in a group of robots: Effect of interference. *Autonomous Robots*, 13(2):127–141, 2002.
- [41] Kristina Lerman, Aram Galstyan, Alcherio Martinoli, and Auke Ijspeert. A macroscopic analytical model of collaboration in distributed robotic systems. *Artificial Life Journal*, 7(4):375–393, 2001.
- [42] Kristina Lerman and Onn Shehory. Coalition Formation for Large-Scale Electronic Markets. In *Proceedings of the International Conference on Multi-Agent Systems (ICMAS'2000), Boston, MA, 2000.*, pages 167–174, 2000.
- [43] R. Mahnke and J. Kaupužs. Stochastic theory of freeway traffic. *Physical Review*, E59(1):117–125, 1999.
- [44] R. Mahnke and N. Pieret. Stochastic master-equation approach to aggregation in freeway traffic. *Physical Review*, E56(3):2666–2671, 1997.
- [45] A. Martinoli. *Swarm Intelligence in Autonomous Collective Robotics: From Tools to the Analysis and Synthesis of Distributed Control Strategies*. PhD thesis, PhD Thesis No 2069, EPFL, 1999.
- [46] A. Martinoli, A. J. Ijspeert, and L. M. Gambardella. A probabilistic model for understanding and comparing collective aggregation mechanisms. In Dario Floreano, Jean-Daniel Nicoud, and Francesco Mondada, editors, *Proceedings of the 5th European Conference on Advances in Artificial Life (ECAL-99)*, volume 1674 of *LNAI*, pages 575–584, Berlin, September 13–17 1999. Springer.
- [47] A. Martinoli and F. Mondada. Collective and cooperative group behaviors: Biologically inspired experiments in robotics. In O. Khatib and J. K. Salisbury, editors, *Proc. of the Fourth Int. Symp. on Experimental Robotics ISER-95*. Springer Verlag, June-July 1995.
- [48] Alcherio Martinoli and Kjerstin Easton. Modeling swarm robotic systems. In B. Siciliano and P. Dario, editors, *Proc. of the Eight Int. Symp. on Experimental Robotics ISER-02, Sant'Angelo d'Ischia, Italy*, Springer Tracts in Advanced Robotics 5, pages 297–306, New York, NY, July 2003. Springer Verlag.

- [49] Alcherio Martinoli and Kjerstin Easton. Optimization of swarm robotic systems via macroscopic models. In A. C. Schultz, L. E. Parker, and F. E. Schneider, editors, *Proc. of the Second Int. Workshop on Multi-Robots Systems, March, 2003, Washington, DC*, pages 181–192. Kluwer Academic Publishers, 2003.
- [50] M. Matarić. Minimizing complexity in controlling a mobile robot population. In *Proceedings of the 1992 IEEE International Conference on Robotics and Automation*, pages 830–835, Nice, France, 1992.
- [51] M. J. Matarić, M. Nilsson, and K. Simsarian. Cooperative multi-robot box pushing. In *Proceedings of the 1995 IEEE/RSJ International Conference on Intelligent Robots*, 1995.
- [52] C. Melhuish, O. Holland, and S. Hoddell. Collective sorting and segregation in robots with minimal sensing. In R. Pfeifer, B. Blumberg, J.-A. Meyer, and S.W. Wilson, editors, *From Animals to Animats, Proceedings of the Fifth International Conference of The Society for Adaptive Behavior (SAB98)*, pages 465–470. MIT Press, 1998.
- [53] O. Michel. Webots: Symbiosis between virtual and real mobile robots. In J.-C. Heudin, editor, *Proc. of the First Int. Conf. on Virtual Worlds, Paris, France,*, pages 254–263. Springer Verlag, 1998. See also <http://www.cyberbotics.com/webots/>.
- [54] E. Nitz, R. C. Arkin, and T. Balch. Communication of behavioral state in multi-agent retrieval tasks. In Lisa Werner, Robert; O’Conner, editor, *Proceedings of the 1993 IEEE International Conference on Robotics and Automation: Volume 3*, pages 588–594, Atlanta, GE, May 1993. IEEE Computer Society Press.
- [55] Manfred Opper and David Saad, editors. *Advanced Mean Field Methods: Theory and Practice*. MIT Press, Cambridge, MA, 2001.
- [56] Esben H. Østergaard, Gaurav S. Sukhatme, and Maja J. Matarić. Emergent bucket brigading - a simple mechanism for improving performance in multi-robot constrained-space foraging tasks. In *Proceedings of the 5th International Conference on Autonomous Agents (AGENTS-01)*, 2001.
- [57] Stephen W. Pacala, Deborah M. Gordon, and H. C. J. Godfray. Effects of social group size on information transfer and task allocation. *Evolutionary Ecology*, 10:127–165, 1996.
- [58] Lynne Parker. Alliance: An architecture for fault-tolerant multi-robot cooperation. *IEEE Transactions on Robotics and Automation*, 14(2):220–240, 1998.
- [59] E. M. Purcell. Life at low reynolds number. *American Journal of Physics*, 45:3–11, 1977.

- [60] W.-J. Rappel, A. Nicol, A. Sarkissian, H. Levine, and W. F. Loomis. Self-organized vortex state in two-dimensional dictyostelium dynamics. *Physical Review Letters*, 83:1247–1250, 1999.
- [61] Ruud Schoonderwoerd, Owen Holland, and Janet Bruten. Ant-like agents for load balancing in telecommunications networks. In W. Lewis Johnson and Barbara Hayes-Roth, editors, *Proceedings of the 1st International Conference on Autonomous Agents*, pages 209–216, New York, February 5–8 1997. ACM Press.
- [62] Frank Schweitzer, Kenneth Lao, and Fereydoon Family. Active random walkers simulate trunk trail formation by ants. *BioSystems*, 41:163–166, 1997.
- [63] Herbert A. Simon. *The Sciences of the Artificial*. MIT Press, Cambridge, MA, 3rd edition, 1996.
- [64] Herbert A. Simon and Albert Ando. Aggregation of variables in dynamic systems. *Econometrica*, 29:111–138, 1961.
- [65] Ken Sugawara and Masaki Sano. Cooperative acceleration of task performance: Foraging behavior of interacting multi-robots system. *Physica*, D100:343–354, 1997.
- [66] Ken Sugawara, Masaki Sano, Ikuo Yoshihara, and K. Abe. Cooperative behavior of interacting robots. *Artificial Life and Robotics*, 2:62–67, 1998.
- [67] Ken Sugawara, Masaki Sano, Ikuo Yoshihara, K. Abe, and T. Watanabe. Foraging behavior of multi-robot system and emergence of swarm intelligence. In *Proc. of IEEE Int. Conf. on Systems, Man, and Cybernetics (SMC-99)*, pages 257–262, 1999.
- [68] Guy Theraulaz, Eric Bonabeau, and Jean-Luis Deneubourg. Response threshold reinforcement and division of labour in insect societies. *Proc. of Royal Society of London, Serie B*:327–332, 1998.
- [69] Richard T. Vaughan, Kasper Støy, Gaurav S. Sukhatme, and Maja J. Matarić. Blazing a trail: Insect-inspired resource transportation by a robot team. In *Proceedings of the 5th International Symposium on Distributed Autonomous Robotic Systems (DARS), Knoxville, TN, 2000*.
- [70] Richard T. Vaughan, Kasper Støy, Gaurav S. Sukhatme, and Maja J. Matarić. Whistling in the dark: cooperative trail following in uncertain localization space. In Carlos Sierra, Gini Maria, and Jeffrey S. Rosenschein, editors, *Proceedings of the 4th International Conference on Autonomous Agents (AGENTS-00)*, pages 187–194, NY, June 3–7 2000. ACM Press.
- [71] Daniel Walgraef. *Spatio-Temporal Pattern Formation and with Examples from Physics and Chemistry and Materials Science*. Springer, New York, NY, 1997.

- [72] Stephen Wolfram. *Cellular Automata and Complexity*. Addison-Wesley, Reading, Mass., 1994.