Contents

Two paradigms for the design of artificial collectives

Kristina Lerman and Aram Galstyan
Information Sciences Institute
Univ. of Southern California
4676 Admiralty Way
Marina del Rey, CA 90292
{`lerman,galstyan`}`@isi.edu`

January 6, 2003

ABSTRACT Artificial collectives are systems composed of multiple autonomous information or software agents, mobile robots, or nodes in a sensor or communication networks. In the future, such systems will be responsible for many important tasks, such as highway traffic control, disaster response, toxic spill monitoring and clean-up, as well as exploration of other planets. Because such systems will have to function in environments with unreliable communication channels, where agents are likely to fail, they will have to be reliable, scalable, robust, adaptable, and amenable to quantitative mathematical analysis. The last property is important because analysis is crucial to understanding the issues of the design, control, adaptability and dynamics of collective behavior. We describe two approaches to distributed control of artificial collectives and study them quantitatively. The first, biologically-based control, relies on local interactions among many simple agents to create desirable collective behavior. The second approach allows collectives to maximize their world utility using market-based mechanisms. We present two applications — foraging in a group of robots and resource allocation in dynamic environments — that utilize these control paradigms and perform an analysis of each problem.

0.1 Introduction

Artificial collectives, also known as multi-agent systems, are systems composed of multiple autonomous agents, *e.g.*, software agents, mobile robots, or nodes in a sensor network, and have become a focus of intense study by the Artificial Intelligence and Networks communities. The interest is easy to understand: in the future, many tasks such as directing traffic flow on roads and highways, coordinating a response to a disaster or emergency, toxic spill monitoring and clean-up, exploration of other planets, to name just a few applications, will be delegated to swarms of simple (and cheap) units, be they information or software agents, mobile embodied agents, or nodes in a sensor or communications network. These multi-agent swarms will have to function in dynamic environments where communication channels are unreliable and agents are likely to fail. Therefore, to be useful, they will have to satisfy the following criteria: they will have to be (1) reliable: show good performance in uncertain dynamic environments; (2) scalable: work equally well for systems composed of tens or thousands of agents; (3) robust: be tolerant of individual agent error or failure; (4) adaptable: be tolerant of changing environment or task requirements.

Distributed control schemes in systems composed of simple agents, which *collectively* accomplish some desired task, satisfy the above requirements and are preferable to alternative system architectures. Here by simple agents we mean agents that are not deliberative, *i.e.*, that do not have the capacity to reason, plan action or negotiate with other agents. Simple agents are less likely to fail or produce errors than more complex deliberative agents, making the system more robust. Although deliberative agents are also capable of collective action in the absence of central control, these agents require detailed knowledge about the capabilities and states of other agents with which they may interact. Acquiring the knowledge necessary to coordinate collective behavior and carrying out the requisite computation may be expensive and impractical, especially for systems containing more than a dozen or so agents.

Several tools are available for studying the behavior of artificial collectives. In the robotics domain, experiment is the most direct way to observe the behavior of the system. Sensor-based simulators that attempt to realistically model the environment and the robots imperfect sensing of and interaction with it, offer another investigation tool. However, experiments and even grounded simulations, are very costly and time-consuming to implement, and often do not scale well as the size of the system grows. Numeric simulations and analytic models are examples of a mathematical approach that offers an alternative to experiment and grounded simulations. Analysis allows us to efficiently study the behavior of collectives, even very large ones, and gain insight into system design, (*e.g.*, what parameters optimize performance or prevent instability, *etc.*). Despite its power, mathematical approach has so far not often been used to study artificial collectives (see [24], [52] and our own [35] and [34] for exceptions).

One the obstacles to wider implementation of artificial collectives is the difficulty of inverse problem, *i.e.*, specifying the local rules that will lead to desirable collective behavior, by which we mean the behavior that will maximize the value of the collective's global utility. This problem has been exacerbated by lack of analysis tools to evaluate different designs. In this paper we discuss two approaches to the design of distributed control for multi-agent systems and present two analysis tools used to study them. The first approach, referred to as biologically-based control (Section 0.2), uses local interactions among many simple agents to create desirable collective behavior. The second approach to distributed control uses market-based mechanisms (Section 0.3) for coordination and adaptation between agents. We describe two applications that utilize these control paradigms and perform a quantitative analysis of each problem. We use two different analysis techniques: the biologically-inspired system is studied analytically, while the market-based system is studied numerically. The former is a general analytic approach that can be applied to other artificial collectives satisfying simple criteria which we define in the paper.

0.2 Biologically-Based Control

Biologically-based control [9], sometimes referred to as Swarm Intelligence [6], is modeled on social insects (and other species), in which complex collective behavior arises out of local interactions among simple agents [48]. This approach takes advantage of the minimalist design [8, 36] of the individual agents and the inherent distributedness of the collective exploiting either direct (agent-to-agent) or indirect (through the environment) local interactions. In the last few years, this control paradigm has been successfully applied in the collective robotics domain: aggregation [5, 52, 39, 36] and segregation [23], beacon and odor localization [20, 21], collaborative mapping [7], collaborative transportation [27, 33], work division and task allocation [32, 1], flocking and foraging [42, 18]. All these works have been performed using groups of simple, autonomous reactive or behavior-based [2] robots or embodied simulated agents, exploiting local communication forms among teammates (implicit, through the environment, or explicit, wireless communication), and fully distributed control. The simplicity of agents and their interactions makes the collective behavior of the system amenable to quantitative mathematical analysis. In this paper we present a methodology for creating a mathematical model of an artificial collective using biologically-based control. We use the well-known (outside of computer science) theory of stochastic systems, and show how this analysis applies to a general class of Markov, or memoryless, agents. We illustrate the approach by using it to study collaboration in a group of robots and qualitatively compare the predictions of the mathematical model to the results of experiment and sensor-based simulations.

0.2.1 Mathematical Analysis of Collectives

A mathematical model of a collective comes in two flavors: *microscopic* or *macroscopic*. A microscopic description treats individual agent as the fundamental unit of the model and describes the agent's interactions with other agents and the environment. There are several variations of the microscopic approach. A common method employed by physicists consists of writing down the microscopic equations of motion for each agent and solving them to study the behavior of the system. For large systems, however, solving equations with many degrees of freedom is often impractical. Microscopic simulations, such as molecular dynamics, cellular automata [57] and particle hopping models [14], are popular tools for studying dynamics of large multi-agent systems. In these simulations, agents change state stochastically or depending on the state of their neighbors. Another example of the microscopic approach are the probabilistic models developed by Martinoli and coworkers [37, 38, 25] to study collective behavior in groups 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, allow them to study the group behavior of the multi-robot system.

Unlike microscopic models, macroscopic models directly describe collective behavior. A macroscopic description offers several advantages over the microscopic approach. It is more computationally efficient, because it uses many fewer variables. The macroscopic descriptions also tend to be more universal and, therefore, more powerful: the same mathematical description can be applied to other systems governed by the same abstract principles. Of course, the two descriptive levels are related, and it may be possible in some cases to exactly derive the parameters of the macroscopic model from microscopic theory. Schweitzer and coworkers [50] have done just that in their analytic study of trail formation by ants and people.

Our approach is based on viewing collectives as stochastic systems, therefore, it is inherently probabilistic. Though the microscopic description of such systems may be very complex, we will show that a macroscopic description takes a very simple (probabilistic) form.

Stochastic Systems Approach: a Tutorial

The behavior of individual agents in a collective has many complex influences, even in a controlled laboratory setting. Agents are influenced by external, often unanticipated, forces. For robots, external forces include friction, which may vary with the type of surface the robot is moving on, battery power, sound and ambient light. 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 robots is for them to turn a random angle and move forward. In summary, the behavior of agents in a collective is so complex, the collective itself may best described probabilistically, as a stochastic system.

Below we present a tutorial on mathematical analysis of stochastic systems. For details we refer the reader to an excellent text by Van Kampen [29]. We begin by defining some concepts. *State* labels a set of related agent behaviors required to accomplish some task. Each of the high level states may consist of a single action or behavior, or a set of behaviors. For example, when the 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 one state to another, and upon completion of the task it may return to the initial state to repeat the cycle once again. It is clear that during a sufficiently short time interval, each agent in a collective is in exactly one of a finite number of states. 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 often useful to *coarse-grain* the system by choosing a smaller number of states, each incorporating a set of related agent actions or behaviors. Coarse-graining is even more desirable because we are often interested in finding the *minimal* model that explains observed features of the collective.

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

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

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. The time evolution of this probability distribution is described by the Stochastic Master Equation.

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; however, many artificial collectives, specifically those based on reactive and behavior-based robots, satisfy the Markov property. The Markov property can be restated more precisely: the configuration of a system at time $t + \Delta t$ depends only on the configuration of the system at time t. This fact allows us to rewrite the marginal probability density $P(\vec{n}, t + \Delta t)$ in terms of conditional probabilities:

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

We can, therefore, write the change in probability density as

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

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

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

with transition rates defined as

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

Equation 3 is known as the Master Equation, and it fully determines the evolution of a stochastic system. Once the probability distribution $P(\vec{n}, t)$

is found, we can calculate system properties. Sometimes, however, it is more useful to study the system's average quantities. The Rate Equation, which can be derived from the Master Equation (see, *e.g.*, [29]), governs the time evolution of average occupation numbers:

$$\frac{\partial}{\partial t}\langle n_k \rangle = \sum_j w_{jk}(\langle \vec{n} \rangle)\langle n_j \rangle - \langle n_k \rangle \sum_j w_{kj}(\langle \vec{n} \rangle) \tag{5}$$

where $w_{ij}(\langle \vec{n} \rangle)$ is the transition rate from state j to state i (for simplicity, we allow only individual transitions between states). This equation has the following interpretation: occupation number n_k will increase in time (first term in Eq. 5) 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).

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 [17]; in physics, the growth of semiconductor surfaces [4]; in ecology to study population dynamics including predator-prey systems [19]; in biology to model the behavior of ant colonies [46]. The rate equation has also found applications in the social sciences [22]. However, with the exception of the work by Huberman and Hogg [24] on computational ecologies and Sugawara and coworkers [53, 54] on foraging in a group of communicating robots, the rate equation approach has not been used in the robotics and AI communities.

The rate equation is usually derived from the phenomenological finite difference equation describing the change in the instantaneous value of a dynamic variable (*e.g.*, US population) over some time interval Δt (*e.g.*, a decade is used by the Census Bureau). By taking the limit $\Delta t \to 0$, one recovers the differential form of the rate equation. For stochastic Markov systems, the rate equation simply describes the evolution of the first moment (mean) of the probability distribution. How closely the mean tracks the behavior of a dynamic variable depends on the magnitude of its fluctuations (higher moments of the distribution). The larger the system, the smaller the (relative) fluctuations, the more accurately the rate equation describes the evolution of dynamic variables. In a small system, the experiment may be repeated many times to average out the effect of fluctuations; therefore, the (continuous) occupation number in the rate equation can be thought to represent an (integer) dynamic variable averaged over repeated experiments. Pacala *et al.* [46] 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.

Rate Equation and Artificial Collectives

The rate equation approach presented above 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 and sensors, 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 swarm-like systems of simple agents. 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 dilute and 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. Crowding effects will also make calculation of transition rates difficult. 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. Another potential limitation of the approach is that, as a mean-field type approach, it is better suited for larger systems, where fluctuations are relatively less important. However, as shown by work of Pacala *et al.*, the rate equations approach becomes a good quantitative description of systems as small as ten agents. We have obtained excellent quantitative agreement with simulations data for systems containing as few as two to ten agents [34].

Despite the limitations outlined above, we believe that the rate equation is a useful tool for mathematical analysis of macroscopic dynamics of artificial collectives. 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, specified by its controller (often represented as a Finite State Machine). Clearly, in the worst case, the macroscopic diagram will be represented by the same Finite State Machine (FSM) as the microscopic controller. However, it is often useful to coarse-grain the system by merging related blocks into a single state, thereby reducing the complexity of the macroscopic diagram.

For example, we may take the searching state of robots to consist of the actions *wander in the arena, detect objects* and *avoid obstacles.* When it is necessary to explicitly take obstacle avoidance into account, *e.g.*, when the density of robots becomes high, the searching state may be split into three states, one for each behavior. In most cases, however, we are 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 the modeler's intuition, a simple rule of thumb can be used as a guide for proper coarse-graining — merge only contiguous states of the FSM. This is easily done (programmatically, if necessary) by drawing a box around these states. To complete the model, 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 describe how the number of agents in those states changes in time. Every transition will be accounted for by a term in each equation, with transition rates specified by the details of the interactions between agents. Note that the macroscopic state diagrams bare a resemblance to Markov chains [45]. Indeed, our models can be considered instances of continuous time Markov chain models. However, we are not interested in studying the characteristics of Markov chains *per se* (*e.g.*, identifying recurrent states, communicating classes, *etc.*); rather, we use them as a guide for constructing a model of the collective and studying its dynamics. In the next section we illustrate our approach by applying it to study collaboration in a group of robots.

0.2.2 Collaboration in Robots

Collaboration can significantly improve the performance of a collective. In "strictly collaborative" systems [36], collaboration is an explicit requirement, because no single agent can successfully complete the task on its own. Such systems are common in insect as well as human societies, *e.g.*, transport of objects 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 [41, 40, 54, 33, 55, 25]. We will focus on a specific case study initiated by Martinoli and collaborators [40] and studied in detail by Ijspeert *et al.* [25]. In this system collaboration in a group of reactive robots was achieved entirely through local interactions, *i.e.*, without explicit global communication or coordination among robots. This system, therefore, is a compelling and effective model of how collaboration may arise in natural systems, such as insect societies. In addition, the simplicity of the robots' interactions lends itself to mathematical analysis. In this section we study an analytical model of collaboration in a group of robots.

Stick-pulling Experiments in Groups of Robots

The stick-pulling experiments were carried out by Ijspeert *et al.* to study the dynamics of collaboration among locally interacting simple reactive robots. Figure 1 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 on its own (because the stick is too long) — a collaboration between two robots is necessary for the task to be successfully completed. In a general case, a collaboration between an arbitrary number of robots may be required to successfully complete the task (because sticks may be of varying length). A collaboration occurs in the following way: one robot finds a stick, and waits for a second robot to find it, partially lifting the stick 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 1. 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 2. 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.

Real Robots, Embodied Simulations and Microscopic Modeling

Ijspeert *et al.* [25] studied the dynamics of collaboration in the stick-pulling experiment at three different levels: by conducting experiments with physical robots; using a sensor-based simulator of robots; and using a microscopic probabilistic model. The physical experiments were carried out in groups of two to six Khepera robots in an arena containing four sticks. Because experiments with physical robots are very time consuming, Webots [43], the sensor-based simulator of Khepera robots, was used to systematically explore parameters affecting the dynamics of collaboration. The Webots simulator attempts to faithfully model the environment and replicate the 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 all the robots in the arena. The probabilistic microscopic model, on the other hand, does not attempt to compute trajectories of individual robots. Rather, it is a numerical model in which the robot's actions — encountering a stick, a wall, another robot, a robot gripping a stick, or wandering around the arena — are represented as a series of stochastic events, with probabilities based on simple geometric considerations and systematic tests with one or two real robots. For example, the probability of a robot encountering a stick is equal to the product of the number of ungripped sticks, and the detection area of the stick normalized by the arena area. Probabilities of other interactions can be similarly calculated. The microscopic simulation consists of running several processes in parallel, one for each robot, while keeping track of the global state of the environment, such as the number of gripped and ungripped sticks. According to Ijspeert *et al.* the acceleration factor for Webots and real robots can vary between one and two orders of magnitude for the experiments presented here. Because the probabilistic model does not require calculations of the details of the robots' trajectories, it is about 300 times faster than Webots for this experiment.

Experimental and Simulation 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 3. 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 collabora-

FIGURE 2. Flowchart of the robots' controller (from Ijspeert *et al.*).

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

A Mathematical Model of Collaboration

In the following sections we present a macroscopic analytical model of the stick-pulling experiments in a homogeneous multi-robot system. Such a model is useful for the following reasons. First, the complexity of a macroscopic model is independent of the system size, *i.e.*, the number of robots: therefore, the time required to obtain solutions for a system of $5,000$ robots is as long as that to obtain solutions for a system of five robots, whereas for a microscopic description the time required for computer simulation scales

FIGURE 3. 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.

at least linearly with the number of robots. Second, our approach allows us to derive analytic expressions for certain important parameters, (*e.g.*, those for which the performance is optimal). It also enables us to study the stability properties of the system, and see whether solutions are robust under external perturbation or noise. These capabilities are important for the design and control of large multi-agent systems.

In order to construct a mathematical model of stick-pulling experiments, it is helpful to draw the 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, shown in Fig. 2, 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. 2, 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 [35] and found that it does not qualitatively change the results of the simpler model. For now, we are interested in the *minimal* model that reproduces main experimental results.

In addition to states, we must also specify all possible transitions between states. When it finds a stick, the robot makes a transition from the search state to the gripping state. After both a successful collaboration and when it times out (unsuccessful collaboration) the robot releases the stick and makes a transition into the searching state, as shown in Fig. 4. These arrows correspond to the arrow entering and the two arrows leaving the dotted box in Fig. 2. 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 system is a semi-Markov system, because the transition from gripping to the searching state depends not only on the present state (gripping) but also on how long the robot has been in the gripping state, *i.e.*, whether the waiting has timed out. This property of the system is captured by time-dependent transition rates.

FIGURE 4. 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 time out.

The dynamic variables of the model are $N_s(t)$, $N_g(t)$, the number of robots in the searching and gripping states respectively. Also, let $M(t)$ be the number of unextracted sticks at time t. The latter variable does not represent a macroscopic state, rather it tracks the state of the environment. The rate equations governing the dynamics of the system read:

$$\begin{aligned}
\frac{dN_s}{dt} &= -\alpha N_s(t)\Big(M(t) - N_g(t)\Big) + \tilde{\alpha} N_s(t) N_g(t) \\
&\quad + \alpha N_s(t-\tau)\Big(M(t-\tau) - N_g(t-\tau)\Big)\Gamma(t;\tau) && (6) \\
N_g &= N_0 - N_s && (7) \\
\frac{dM}{dt} &= -\tilde{\alpha} N_s(t) N_g(t) + \mu(t) && (8)
\end{aligned}$$

where α, $\tilde{\alpha}$ are the rates at which a searching robot encounters a stick and a gripping robot respectively, τ is the gripping time parameter, and $\mu(t)$ is the rate at which new sticks are added. The parameters α, $\tilde{\alpha}$, and τ 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. $\Gamma(t;\tau)$, the fraction of failed collaborations at time t, is the probability no robot came "to help" during the time interval $[t-\tau, t]$. It corresponds to the time-dependent transition rates (marked "u" in Fig. 4) 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

$$\begin{aligned} \Gamma(t;\tau) &= \prod_{i=1}^{K} [1 - \tilde{\alpha}\delta t N_s(t-\tau+i\delta t)]\Theta(t-\tau) \\ &\equiv \exp\left[\sum_{i=1}^{K} \ln[1 - \tilde{\alpha}\delta t N_s(t-\tau+i\delta t)]\right]\Theta(t-\tau) \qquad (9) \end{aligned}$$

The step function $\Theta(t-\tau)$ ensures that $\Gamma(t;\tau)$ is zero for $t < \tau$. Finally, expanding the logarithm in Eq. (9) and taking the limit $\delta t \to 0$ we obtain

$$\Gamma(t;\tau) = \exp[-\tilde{\alpha} \int_{t-\tau}^{t} dt' N_s(t')]\Theta(t-\tau) \qquad (10)$$

The three terms in Eq. 6 correspond to the three arrows between the states in Fig. 4. 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, 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. 7. Equation 8, states that the number of sticks, $M(t)$, decreases in time at the rate of successful collaborations but also increases at the rate new sticks are added. The equations are subject to the initial conditions that at $t = 0$ the number of searching robots in N_0 and the number of sticks is M_0.

To proceed further, let us introduce $n(t) = N_s(t)/N_0$, $m(t) = M(t)/M_0$, $\beta = N_0/M_0$, $R_G = \tilde{\alpha}/\alpha$, $\tilde{\beta} = R_G\beta$ and a dimensionless time $t \to \alpha M_0 t$, $\tau \to \alpha M_0 \tau$. μ' is the dimensionless rate at which new sticks are added. $n(t)$ is the fraction of robots in search state and $m(t)$ is the fraction of unextracted 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 6–8 can be rewritten in dimensionless form as:

$$\begin{aligned} \frac{dn}{dt} &= -n(t)[m(t) + \beta n(t) - \beta] + \tilde{\beta} n(t)[1-n(t)] + n(t-\tau)[m(t-\tau) \\ &\quad +\beta n(t-\tau) - \beta] \times \gamma(t;\tau) \qquad (11) \\ \frac{dm}{dt} &= -\beta\tilde{\beta} n(t)[1-n(t)] + \mu' \qquad (12) \end{aligned}$$

$$\gamma(t;\tau) \quad = \quad \exp[-\tilde{\beta}\int_{t-\tau}^{t} dt' n(t')] \qquad (13)$$

Equations 11–13 together with initial conditions $n(0) = 1$, $m(0) = 1$ determine the dynamical evolution of the system. Note that only two parameters, β and τ, appear in the equations and, thus, determine the behavior of solutions. The third parameter $\tilde{\beta} = R_G\beta$ is fixed experimentally and is not independent. Note that we do not need to specify α and $\tilde{\alpha}$ — they enter the model only through R_G (throughout this paper we will use $R_G = 0.35$, the value reported in [25]).[1] Below we provide a detailed analysis of the equations.

Results

In [35] we studied the steady state properties of the system (Eq. 11–13) for the case of a static environment, $m(t) = const = 1$. Experimentally this was realized by replacing sticks in their holes after they were pulled out, $\mu(t) = \tilde{\alpha}N_s(t)N_g(t)$. Particularly, it was found that for $\beta < \beta_c \equiv 2/(1 + R_G)$ the steady state collaboration rate, given by $R(\tau, \beta) = \beta\tilde{\beta}n(\tau, \beta)(1 - n(\tau, \beta))$, has a maximum for a certain value of the gripping time parameter τ.

Below we provide an analysis for the case of a dynamically changing environment. There are no new sticks added to the system, $\mu(t) = 0$; therefore, the number of unextracted sticks decreases monotonically. First, we consider the case of $\tau = \infty$, which corresponds to robots gripping the sticks and holding them part-way out of the ground indefinitely. Solving Eq. 11–13 yields the number of robots in the search state, Fig. 5(a), and the number of unextracted sticks, Fig. 5(b), as a function of time for two values of β. A qualitatively different solution is obtained for small and large values of the parameter β. For small β's, the fraction of searching robots exponentially decreases to zero and the fraction of unextracted sticks "saturates" at $m(t) \to m \neq 0$ as $t \to \infty$. For sufficiently large β's, however, all the robots are in the searching mode and all the sticks are collected in the long time limit. This different behavior is illustrated in Fig. 5. For ($\beta = 0.5$), the number of searching robots drops to zero as all robots end up gripping a stick and only a small fraction of the sticks is collected (solid lines). For $\beta = 1$, however, the number of searching robots first decreases as robots find sticks, but then it increases because successful collaborations return

---

[1] The parameter α can be easily calculated from experimental values quoted in [25]. 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_RW_Rdt, where $V_R = 8.0\ cm/s$ is robot's speed, and $W_R = 14.0\ cm$ is robot's detection width. If the arena radius is $R = 40.0\ cm$, a robot will detect sticks at the rate $\alpha = V_RW_R/\pi R^2 = 0.02\ s^{-1}$. According to [25], 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.

the gripping robots back to the searching state. Eventually, all the sticks are collected (dashed lines), and all the robots are in the searching state.

FIGURE 5. (a) Fraction of robots in searching state and (b) fraction of unextracted sticks as a function of time for $\tau = \infty$, $\beta = 0.5$ (solid), and $\beta = 1$ (dashed), where τ is the gripping time parameter and β is the ratio of robots to sticks.

When the gripping time parameter τ is finite, the solution to Eq. 11 displays characteristic oscillations (Fig. 6(a)) which die out as the solution approaches its steady state value $n(t \rightarrow \infty) = 1$, $m(t \rightarrow \infty) = 0$. Note that the steady state solution for this case is trivial in the sense that all the sticks are eventually collected and all the robots are in the searching state after some transient time. Consequently, we modify the measure of collaboration as the inverse of the time it takes to collect certain fraction of sticks (*e.g.*, 0.9 in our analysis). Collaboration rate *vs* the gripping time parameter for different values of β is shown in Fig. 6(b). As for the static case studied in the experiments, an optimal behavior is seen for small β, *i.e.*, for these cases there exists a gripping time parameter that maximizes the collaboration rate. Note also, that the kinks in the graph are not numeric artifacts, but are due to our definition of the collaboration rate. Because $n(t)$ oscillates in the transient regime, the time to collect $100f\%$ of the sticks will vary significantly if it falls within this regime. We have checked that these kinks disappear as $f \rightarrow 1$.

0.3 Market-Based Control

Previous research has shown that market-based (or game–dynamical) control strategy, in which agents make "economically" motivated decisions with the goal of maximizing individual payoff but which result in the optimization of some global system property, *e.g.*, world utility, is a feasible distributed control mechanism for artificial collectives (see, *e.g.*, [56] for an overview).

Below we present a model of distributed coordination of agents in nonstationary environments. Specifically, we consider a network of interconnected boolean agents that compete for a resource with a limited capacity using a form of Minority Games. At each time step the agents face a choice of whether to use the resource or not, and the agents in the winning group are rewarded while those in the loosing group are punished. The winning group is that whose size is less than resource capacity: thus, this game penalizes overcrowding. We will consider the case where the resource capacity changes in time and demonstrate that under certain conditions our model shows globally adaptive and coordinated behavior, resulting in efficient resource utilization.

0.3.1 Minority Games

The Minority Game [12] (MG) is a simple model of competing agents yet it has a very rich dynamics. It was introduced by Challet and Zhang as a simplification of Arthur's El Farol Bar attendance problem [3]. The MG consists of N agents with bounded rationality that repeatedly choose between two alternatives labelled 0 and 1 (*e.g.*, staying at home or going to the bar). At each time step, agents who made the minority decision win. In the Generalized Minority Game [26], the wining group is 1 (0) if the fraction of the agents who chose "1" is smaller (greater) than the capacity level η, $0 < \eta < 1$. For $\eta = 0.5$, the game reduces to the the traditional MG. Each agent uses a set of S strategies to decide its next move and reinforces strategies that would have predicted the winning group. A strategy is simply a lookup table that prescribes a binary output for all possible inputs. In

FIGURE 6. (a) Fraction of searching robots for $\tau = 5$ and $\beta = 0.4$. (b) Collaboration rate per robot vs gripping time parameter τ for $\beta = 0.25, 0.5, 0.75, 1$, for the bottom to top curves respectively. Collaboration rate is defined as the inverse of the time needed to collect 90% of the sticks.

the original version of the game, the input is a binary string containing the last m outcomes of the game, so the agents interact by sharing the same global signal. If the agents choose either action with probability 1/2 (the random choice game), then, on average, the number of agents choosing "1" (henceforth referred to as utilization) is $(N-1)/2$ with standard deviation $\sigma = \sqrt{N}/2$ in the limit of large N. The most interesting phenomenon of the minority model is the emergence of a coordinated phase, where the standard deviation of utilization, the volatility, becomes smaller than in the random choice game. The coordination is achieved for memory sizes for which the dimension of the reduced strategy space is comparable to the number of agents in the system, $2^m \sim N$ [13, 49]. It was later pointed out that the dynamics of the game remains mostly unchanged if one replaces the string with the actual histories with a random one [10], provided that all the agents act on the same signal. Analytical studies based on this simplification has revealed many interesting properties of the minority model [11, 44].

In addition to the original MG, different versions of the game where the agents interact using local information only (cellular automata [28], evolving random boolean networks [47], personal histories [15]), have been studied. In particular, it was established that coordination still arises out of local interactions, and the system as a whole achieves "better than random" performance in terms of the utilization of resources.

In all previous studies the capacity level has been fixed as an external parameter, so the environment in which the agents compete is stationary. As we mentioned above, however, we are interested in a situation where the environment is changing. It is interesting to see if a coordinated behavior still emerges, and to what degree agents can adapt to the changing environment. Namely, we study a system of boolean agents playing a generalized minority game, assuming that the capacity level is not fixed but varies with time, $\eta(t) = \eta_0 + \eta_1(t)$, where $\eta_1(t)$ is a time dependent perturbation. The framework of the interactions is based on Kauffman NK random boolean nets [30], where each agent gets its input from K other randomly chosen agents, and maps the input to a new state according to a boolean function of K variables, which is also randomly chosen and quenched throughout the dynamics of the system. The generalization we make is that agents are allowed to adapt by having more than one boolean function, or strategy, and the use of a particular strategy is determined by an agent based on how often it predicted the winning group throughout the game. Note that this approach is very different from adaptation through evolution studied previously in the context of the minority model [47].

0.3.2 The Model

We consider a set of N boolean agents described by "spin" variables $s_i = \{0, 1\}$, $i = 1, \ldots, N$. Each agent gets its input from K other randomly chosen agents, and maps the input to a new state:

$$s_i(t+1) = F_i^j(s_{k_1}(t), s_{k_2}(t), ..., s_{k_K}(t)) \tag{14}$$

where s_{k_i}, $i = 1, \ldots, K$ are the set of neighbors, and $F_i^j, j = 1, \ldots, S$ are randomly chosen boolean functions (called strategies hereafter) used by the i-th agent. For each strategy F_i^j, the agent keeps a score that monitors the performance of that strategy, adding (subtracting) a point if the strategy predicted the winning (loosing) side. Let the "utilization" $A(t)$ be the cumulative output of the system at time t, $A(t) = \sum_{i=1}^N s_i(t)$. Then the winning choice is "1" if $A(t) \leq N\eta(t)$, and "0" otherwise. Those in the wining group are awarded a point while the others loose one. Agents play the strategies that have predicted the winning side most often, and the ties are broken randomly.

As a global measure of performance, the world utility, we define efficient utilization as one that has the smallest cumulative resource "waste". If $\delta(t) = A(t) - N\eta(t)$ describes the deviation from the optimal resource utilization, cumulative "waste" over a certain time window is

$$\sigma = \sqrt{\frac{1}{T} \sum_{t=t_0}^{t_0+T} \delta(t)^2} \tag{15}$$

For $\eta_1(t) = 0$ this quantity is simply the volatility as defined in the traditional minority game.

We compare the performance of our system to a default random choice game, defined as follows: assume that the agents are told what the capacity $\eta(t)$ is at time t, and they choose to go to the bar with probability $\eta(t)$. In this case the main utilization will be close to $\eta(t)N$ at each time step, and the fluctuations around the mean are given by the standard deviation

$$\sigma_0^2 = N \frac{1}{T} \int_{t_0}^{t_0+T} dt' \eta(t')[1 - \eta(t')] \tag{16}$$

0.3.3 Results

We performed intensive numerical simulations of the system described above, with the number of agents ranging from 100 to 10^4, and for network connectivity K ranging from 1 to 10. Although in our simulations we used different forms for the perturbation $\eta_1(t)$, in this paper we consider periodic perturbations only. For each K, a set of strategies was chosen for each agent randomly and independently from a pool of 2^{2^K} possible boolean functions, and was quenched throughout the game. In all simulations we used $S = 2$ strategies per agent. Starting from a random initial configuration, the system evolved according to the specified rules. The duration of the simulation T_0 was determined by the particular choice of $\eta(t)$.

Depending on the amplitude of the perturbation, we run the simulations for 10 to 20 periods, and usually used the data for the last two periods to determine σ.

Our main observation is that networks with $K = 2$ show a tendency towards self organization into a phase characterized by small fluctuations, hence, an effective utilization of the resource, even for relatively large variations in the capacity level $\eta(t)$. Note, that in the Kauffman nets with $K > 2$ the dynamics of the system is chaotic with an exponentially increasing length of attractors as the system size grows, while for $K < 2$ the network reaches a frozen configuration. The case $K = 2$ corresponds to a phase transition in the dynamical properties of the network and is often referred as the "edge of the chaos". We would like to reiterate, however, that our system is different from a Kauffman network since the agents have an internal degree of freedom, characterized by their strategies. Specifically, our system does not necessarily have periodic attractors, while in Kauffman nets periodic attractors are guaranteed to exist due to the finite phase space and quenched rules of updating.

Fig. 7 shows a typical segment of the time series of the utilization $A(t)$ for a system of size $N = 1000$, a sinusoidal perturbation $\eta_1(t)$, and different network connectivities. For $K = 1$ the agents react to the changes in the capacity level, however there are strong fluctuations in the utilization series. For $K = 4$, on the other hand, response of the system to the environmental dynamics is very weak, and, as our results indicate, becomes even weaker for larger K. Remarkably, the system with $K = 2$ adapts very efficiently to changes in the capacity level. The inset of Fig. 7 b) shows the time series of the deviation $\delta(t)$ for $K = 2$. Initially there are strong fluctuations, hence poor utilization of the resource, but after some transient time the system as a whole adapts and the strength of the fluctuations decreases. In fact, the standard deviation of the fluctuations is considerably smaller than in the random choice game as defined by Eq. 16. Note also, that the agents have information only about the winning choice, but not the capacity level. This suggests that the particular form of the perturbation may not be important as long as it meets some general criteria for smoothness.

For comparison, we also studied the effect of the changing capacity level in the traditional (generalized) minority model whit publicly available information about the last m outcomes of the game. We plot the utilization and deviation time series for a system with a memory size $m = 6$ (corresponding to the minimum of σ) in Fig. 8. One can see that in this case also the system reacts to the external change. However, the structure of adaptation is very different from the previous case. Indeed, an analysis of Fig. 8 shows that even though the total "wealth", i.e., the total points accumulated by the agents in the system, increases with time, the overall performance in terms of resource allocation as described by σ is much poorer compared to the previous case. Another important difference is that in the traditional system the distribution of wealth among the players is

FIGURE 7. A segment of the utilization time series for $\eta(t) = 0.5 + 0.15sin(2\pi t/T)$,$T = 1000$ and different network connectivity K.

FIGURE 8. Same as in Fig 7 for the traditional MG (global histories) with memory size m=5

much wider than in the system with local information exchange, i.e., the later mechanism of adaptation is socially more “fair”.

In Fig. 9 we plot the variance per agent versus network connectivity K, for system sizes $N = 100, 200, 500, 1000$. For each K we performed 32 runs and averaged results. Our simulations suggest that the details of this dependence are not very sensitive to the particular form of the perturbation $\eta_1(t)$, and the general picture is the same for a wide range of functions, provided that they are smooth enough. As we already mentioned, the variance attains its minimum for $K = 2$, independent of the number of agents in the system. For bigger K it saturates at a value that depends on the amplitude of the perturbation and on the number of agents in the system. We found that for large K the time series of the utilization closely resembles the time series in the absence of perturbation. This implies that for large K the agents do not “feel” the change in the capacity level. Consequently, the standard deviation increases linearly with the number of agents in the system, $\sigma \propto N$. For $K = 2$, on the other hand, σ increases considerably slower with the number of agents in the system, $\sigma \propto N^\gamma$, $\gamma < 1$ (see the inset in Fig. 9). Our results indicate that the scaling (i.e., the exponent γ) is not universal and depends on the perturbation.

We do not currently have an analytic theory for the observed emergent coordination. In contrast to the traditional minority game, where global interactions and the Markovian approximation allow one to construct a mean field description, our model seems to be analytically intractable due to explicit emphasis on local information processing. We strongly believe, however, that the adaptability of the networks with $K = 2$ is related to the peculiar properties of the corresponding Kauffman nets, and particularly, to the phase transition between the chaotic and frozen phases. It is known [16] that the phase transition in the Kauffman networks can be achieved by tuning the homogeneity parameter P which is the fraction of 1’s or 0’s in the output of the boolean functions (whichever is the major-

FIGURE 9. σ^2/N vs the network connectivity for different system sizes and $\eta(t) = 0.5 + 0.15 sin(2\pi t/T)$, $T = 1000$. Inset plot shows the scaling relationship between σ and N for $K = 2$. Average over 16 runs has been taken.

FIGURE 10. Standard deviation per agent vs homogeneity coefficient P for K=3 networks: N=1000, $\eta(t) = 0.5 + 0.15\sin(2\pi t/T)$, $T = 1000$

ity), with the critical value given by $P_c = 1/2 + 1/2\sqrt{1 - 2/K}$. To test our hypothesis, we studied the properties of networks with $K = 3$ for a range of homogeneity parameter P. In Fig. 10 we plot σ^2/N versus the homogeneity parameter P. One can see that the optimal resource allocation is indeed achieved in the vicinity of the $P_c \sim 0.78$, indicating that the properties of Kauffman networks at the phase transition might be responsible in this emergent coordination. Recent investigation of the coordinated phase using information theoretic analysis [51] suggests that emergence of the coordinated phase coincides with the growth of clusters of synchronous agents. More analysis is required to illuminate this insight.

0.4 Related Work

With the exceptions noted below, there has been very little prior work on mathematical analysis of artificial collectives. The closest in spirit to our paper is the work by Huberman, Hogg and coworkers on computational ecologies [24, 31]. These authors mathematically studied collective behavior in a system of agents, each choosing between two alternative strategies. They derived a rate equation for the average number of agents using each strategy from the underlying probability distributions. Our approach is consistent with theirs — in fact, we can easily write down these rate equations directly from the macroscopic state diagram of the system, without having to derive them from probability distributions. Therefore, computational ecologies can be considered another application of the methodology we described in this paper.

Sugawara and coworkers [52, 54] carried out quantitative studies of cooperative foraging in a group of communicating and non-communicating robots in different environments. They have developed a simple state-based analytical model and analyzed it under different conditions. In their system when a robot finds a puck or a collection of pucks, it may broadcast a signal for a period of time to other robots, which move towards it. The robots pick up pucks and bring them home. Sugawara *et al.* did not take the interaction into account explicitly but in an approximate manner. In our model of collaboration, we include the duration of the interaction explicitly, resulting in a better description of the system. Another difference between their work and ours is that their system is not strictly collaborative — collaboration via signalling improves foraging performance but is not a requirement for task completion.

0.5 Conclusion

Distributed control is a superior paradigm for control of artificial collectives that has a number of advantages over alternative designs, namely, robustness, scalability, reliability, and adaptability. However, the “inverse” problem of distributed control is notoriously difficult to solve: it is hard to identify the local rules that will lead to desirable collective behavior. Lack of analysis tools to evaluate different designs has been one of the major obstacles to progress in this area.

We presented two approaches to the *design* of distributed control mechanisms for artificial collectives, as well as two *analysis* tools used to study these mechanisms. The first design approach is biologically-based control, that uses local interactions among many simple agents to create desirable collective behavior. We presented a general methodology for mathematical analysis of artificial collectives designed according to these principles. Our analysis applies to a class of systems known as stochastic Markov systems in which the state of an agent at a future time depends only on the present state (and perhaps on how much time the agent has spent in this state). Though each agent’s actions are unpredictable in stochastic systems, the probabilistic description of the collective behavior is surprisingly simple. It is given by the Rate Equation, that describes how the average macroscopic (collective) system properties change in time.

We illustrated our approach by applying it to study collaboration in a homogeneous group of reactive robots. The robots’ task was to pull sticks out of their holes, and it could be successfully achieved only through collaboration between two robots. The system we modeled was slightly different from the one studied by Ijspeert *et al.* [25] in that the number of sticks was not constant, but decreased as robots pulled them out. Detailed analysis of the Ijspeert *et al.* experiments and simulations is presented in a separate paper [35].

Mathematical analysis shows that the performance of the system where the number of sticks changes in time is qualitatively similar to the performance of the system in the static environment, where the number of sticks remains constant (see Fig. 3). Indeed, we were able to reproduce some of the conclusions of the Ijspeert *et al.* work: namely, the different dynamical regimes for different values of the ratio of robots to sticks (β) and the optimal gripping time parameter for β less than the critical value. Moreover, some conclusions, such as the importance of the parameter β, fall directly out of simple analysis of the model, without requiring any time consuming simulations or experiments. In [35] we reported qualitative agreement between the predictions of the model and results of experiment and simulation. In the simple state-based model we studied, the robot’s future state depends only on its present state (and on how much time it has spent in that state). While the reactive and behavior-based robots 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. Creating a mathematical model that addresses these issues is our next challenge.

The second approach to distributed control uses market-based mechanisms for adaptation of an artificial collective to a dynamically changing environment. We studied a network of boolean agents playing Minority Games and competing for use of a resource with a dynamic capacity. We studied the problem numerically, by simulating games in systems of different sizes, network connectivities K, and capacity functions. We established that networks with connectivity $K = 2$ can be extremely adaptable and robust with respect to capacity level changes. For $K > 2$ the coordination can be achieved by tuning the homogeneity parameter to its critical value. The adaptation happens without the agents knowing the capacity level. Remarkably, the system that uses local information is much more efficient in a dynamic environment than a system that uses global information.

Acknowledgements

The authors wish to thank the following people for insights and useful discussions: Alcherio Martinoli, Maja Matarić, Richard Ross, Dmitry Tsigankov, Jim Crutchfield and Cosma Shalizi. The research reported here was supported in part by the Defense Advanced Research Projects Agency (DARPA) under contracts 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.

0.6 REFERENCES

[1] W. Agassounon, A. Martinoli, and R. M. Goodman. A scalable, distributed algorithm for allocating workers in embedded systems. In *Proc. of the IEEE Conf. on System, man and Cybernetics SMC-01, October 2001, Tucson, AR, USA. To appear.* 2001.

[2] Ronald C. Arkin. *Behavior-Based Robotics.* The MIT Press, Cambridge, MA, 1999.

[3] Brian W. Arthur. Inductive reasoning and bounded rationality. *Am. Econ. Assoc. Papers Proc.*, 84:406, 1994.

[4] A.-L. Barabasi and H.E. Stanley. *Fractal Concepts in Surface Growth.* Cambridge University Press, 1995.

[5] R. Beckers, 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] G. Beni and J. Wang. Swarm intelligence. In *Proc. of the Seventh Annual Meeting of the Robotics Society of Japan, Tokyo, Japan*, pages 425–428, Tokyo, Japan, 1989. RSJ Press.

[7] A. Billard, A.J. Ijspeert, and A. Martinoli. A multi-robot system for adaptive exploration of a fast changing environment: Probabilistic modelling and experimental study. *Connection Science*, 11(3/4):359–379, 1999.

[8] K. Boehringer, R. Brown, B. Donald, J. Jennings, and D. Rus. Distributed robotic manipulation: Experiments in minimalism. In O. Khatib and J. K. Salisbury, editors, *Proc. of the Fourth Int. Symp. on Experimental Robotics, Stanford*, pages 11–25. Lecture Notes in Control and Information Sciences, Springer Verlag, 1995.

[9] Eric Bonabeau, Marco Dorigo, and Guy Theraulaz. *Swarm Intelligence: From Natural to Artificial Systems.* Oxford University Press, New York, 1999.

[10] A. Cavagna. Irrelevance of memory in the minority game. *Phys. Rev.*, E59:R3783, 1998.

[11] D. Challet and M. Marsili. Phase transition and symmetry breaking in the minority game. *Phys. Rev.*, E60:R6271, 1999.

[12] D. Challet and Y.-C. Zhang. Emergence of cooperation and organization in an evolutionary game. *Physica*, A246:407, 1997.

[13] D. Challet and Y.-C. Zhang. On the minority game: Analytical and numerical studies. *Physica*, A256:514, 1998.

[14] D. Chowdhury, L. Santen, and A. Schadschneider. Statistical physics of vehicular traffic and some related systems. *Physics Reports*, 329:199, 2000.

[15] M.A.R. de Cara, O. Pla, and F. Guinea. Learning, competition and cooperation in simple games. *Eur. Phys. J.*, B13:413, 2000.

[16] B. Derrida and Y. Pomeau. Random networks of automata: A simple annealed approximation. *Eur. Phys. Lett.*, 1:45, 1986.

[17] C. W. Garnier. *Handbook of Stochastic Methods.* Springer, New York, NY, 1983.

[18] 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.

[19] Richard Haberman. *Mathematical Models: Mechanical Vibrations, Population Dynamics, and Traffic Flow.* Society of Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1998.

[20] A. T. Hayes, A. Martinoli, and R. M. Goodman. Comparing distributed exploration strategies with simulated and real autonomous robots. In L.E. Parker, G. Bekey, and J. Bahren, editors, *Proc. of the Fifth Int. Symp. on Distributed Autonomous Robotic Systems DARS-00, October, 2000, Knoxville, TN*, pages 261–270. Springer Verlag, 2000.

[21] A. T. Hayes, A. Martinoli, and R. M. Goodman. Swarm robotic odor localization. In *Proc. of the IEEE Conf. on Intelligent Robots and Systems IROS-01, October-November 2001, Maui, Hawaii, USA.* 2001.

[22] 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.

[23] O. Holland and C. Melhuish. Stigmergy, self-organization, and sorting in collective robotics. *Artificial Life*, 5:173–202, 1999.

[24] 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).

[25] 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.

[26] N.F. Johnson, P.M. Hui, Dafang Zheng, and C.W. Tai. Minority game with arbitrary cuttoffs. *Physica*, A269:493, 1999.

[27] P.J. Johnson and J.S. Bay. Distributed control of simulated autonomous mobile robot collectives in payload transportation. *Autonomous Robots*, 2:43–63, 1995.

[28] T. Kalinowski, H.-J. Schulz, and M. Briese. Cooperation in the minority game with local information. *Physica*, A277:502, 2000.

[29] N. G. Van Kampen. *Stochastic Processes in Physics and Chemistry.* Elsevier Science, Amsterdam, revised and enlarged edition edition, 1992.

[30] Stuart A. Kauffman. *The Origins of Order.* Oxford University Press, New York, 1993.

[31] Jeffrey O. Kephart, Tad Hogg, and Bernardo A. Huberman. Collective behavior of predictive agents. *Physica*, D 42:48–65, 1990.

[32] M.J. B. Krieger and J.-B. Billeter. The call of duty: Self-organised task allocation in a population of up to twelve mobile robots. *Robotics and Autonomous Systems*, 30(1-2):65–84, 2000.

[33] C. R. Kube and E. Bonabeau. Cooperative transport by ants and robots. *Robotics and Autonomous Systems*, 30(1–2):85–101, 2000.

[34] Kristina Lerman and Aram Galstyan. Mathematical model of foraging in a group of robots: Effect of interference. *Autonomous Robots*, 13(2), 2002.

[35] 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.

[36] A. Martinoli. *Swarm Intelligence in Autonomous Collective Robotics: From Tools to the Analysis and Synthesis of Distributed Control Strategies.* PhD thesis, No 2069, EPFL, 1999.

[37] 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.

[38] 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.

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

[40] A. Martinoli and F. Mondada. Collective and cooperative group behaviors: Biologically inspired experiments in robotics. In O. Khatib and J. K. Salisbur, editors, *Proc. of the Fourth Int. Symp. on Experimental Robotics ISER-95.* Springer Verlag, June-July 1995.

[41] 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.

[42] Maja Matarić. *Interaction and Intelligent Behavior*. PhD thesis, Dept. of Electrical Engineering and Computer Science, MIT, Cambridge, 1994.

[43] 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/.

[44] See http://www.unifr.ch/econophysics/minority/ for an extensive collection of articles and refernces.

[45] J. R. Norris. *Markov Chains*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, Cambridge, UK, 1997.

[46] 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.

[47] M. Paczuski and K. E. Bassler. Self-organized networks of competing boolean agents. *Phys. Rev. Lett.*, 84:3185, 2000.

[48] J. K. Parrish and W. M. Hamner. *Animal Groups in Three Dimensions*. Cambridge University Press, 1997.

[49] R. Savit, R. Manuca, and R. Riolo. Adaptive competition, market efficiency, phase transition. *Phys. Rev. Lett*, 82(10):2203, 1999.

[50] F. Schweitzer, K. Lao, and F. Family. Active random walkers simulate trunk trail formation by ants. *BioSystems*, 41:153–166, 1997.

[51] Cosma Shalizi. private communication.

[52] Ken Sugawara and Masaki Sano. Cooperative acceleration of task performance: Foraging behavior of interacting multi-robots system. *Physica*, D100:343–354, 1997.

[53] Ken Sugawara, Masaki Sano, and Ikuo Yoshihara. Cooperative acceleration of task performance: Analysis of foraging behavior by interacting multi-robots. In *Proc. IPSJ Int. Symp. on Information Systems and Technologies for Network Society*, pages 314–317, Fukuoka, Japan, September 1997.

[54] Ken Sugawara, Masaki Sano, Ikuo Yoshihara, and K. Abe. Cooperative behavior of interacting robots. *Artificial Life and Robotics*, 2:62–67, 1998.

[55] 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.

[56] Michael P. Wellman. Market-Oriented Programming: Some Early Lessons. In S. H. Clearwater, editor, *Market-Based Control: A Paradigm for Distributed Resource Allocation*, pages 74–95. World Scientific, January 1996.

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