Université Libre de Bruxelles
Faculté des Sciences Appliquées
CoDE – Computers and Decision Engineering

Modeling swarm intelligence systems with Markov chains

IRIDIA
Institut de Recherches Interdisciplinaires
et de Devéloppments en Intelligence Artificielle

Promoteur:
Prof. Marco DORIGO
Co-promoteur:
Dr. Mauro BIRATTARI

Mémoire de:
Gabriele VALENTINI

Rapport d'avancement de recherche
Année Accadémque 2012-2013

Summary

We detail a modeling methodology developed to study and analyze distributed, self-organized swarm intelligence systems. Swarm intelligence systems are composed of a multitude of agents that interact in a self-organized manner without using any sort of centralized control. These systems are appealing for their properties of robustness, scalability and flexibility that allow engineers to address real world problems that cannot be tackled with classical centralized approaches. However, the design of swarm intelligence systems is non-trivial. Indeed, a desired collective behavior of the swarm is achieved through the definition of local agent's control rules. The local interactions among the agents are characterized by strong non-linearities and it is hard to understand a priori the swarm-level effect of a particular agent-level control rule. Mathematical analysis is thus an essential step in the design of swarm intelligence systems. Mathematical models allow designers to study swarm intelligence systems without the need of time-consuming physics-based simulations and real robot experiments. We propose a modeling methodology based on the well-known formalism of time-homogeneous Markov chains. In particular, we employ absorbing Markov chain to define macroscopic models of swarm intelligence systems. The applicability of the proposed methodology is illustrated using two different case studies: First, we consider opinion dynamics in a collective decision making scenario for a swarm of robots. Second, we study the performance of a distributed communication protocol that allows a heterogeneous swarm of robots to achieve spatially targeted communication.

Contents

Chapter 1

Introduction

Artificial swarm intelligence systems are distributed, decentralized and self-organizing systems formed by a collective of agents (Bonabeau et al., 1999; Kennedy and Eberhart, 2001; Şahin, 2005). The primary source of inspiration of artificial swarm intelligence is nature. When defining collective strategies, designers often consider in their work the behavior of social insects (e.g., colonies of ants or termites, swarms of bees or wasps) and gregarious animals (e.g., schools of fish, flocks of birds). In a swarm intelligence system, agents can be either virtual or embodied in a physical body. In the first case, swarms of virtual agents are used to address optimization and other soft computing tasks (Kennedy and Eberhart, 2001). In the second case, swarms of embodied agents correspond to distributed robotic systems and are more commonly referred to as swarms of robots (Şahin, 2005). In this thesis, we focus on the study of swarm robotics systems without taking into consideration systems consisting of virtual agents.

The interest of the scientific community in the study of swarms of robots is motivated by their primary properties: scalability, flexibility and robustness (Brambilla et al., 2013). Control strategies for swarms of robots are scalable and flexible because they are independent from the number of agents in the swarm and from the particular environmental conditions. Control strategies only focus on the individual agent and its local perception of the system resulting in behaviors that adapt to the environment. Swarm of robots are robust because none of the members of the swarm play a critical role. In case of failure of an agent its neighboring swarm-mates obviate to the loss. Such properties are a result of the local interactions of individual components of the swarms

— the robots — and make swarm robotics systems very appealing to engineers for applications where a centralized approach is unfeasible.

In contrast to a centralized (robotic) system, in swarm robotics no particular robot plays the role of the leader. The overall behavior of the swarm is the result of local interactions between its components. A robot in the swarm is unaware of what is happening throughout the rest of the swarm and it solely relies on local information to decide the actions to perform or the goals to pursue. This peculiarity allows systems composed of thousands of robots to be scalable, flexible and robust. However, it also gives rise to the main challenge in swarm robotics: How to design an appropriate set of control rules for the individual robots that allows the system to achieve a desired goal at the swarm level?

In swarm robotics, as well as in other research fields that study distributed systems, mathematical models are in general defined at two different levels (cf. Brambilla et al., 2013). The individual level, or *microscopic* level, considers all the characteristics of the individual agents and their interactions. The swarm level, or *macroscopic* level, models instead the behavior of the swarm as a whole disregarding details of the individual agents. The mapping between microscopic and macroscopic levels is a particularly tough problem in swarm robotics and no general technique exists to solve it. That is, there is no general approach that allows designers to define desired macroscopic properties of the system as a whole and to translate them into a set of control rules for the individual robot. Usually, designers needs to proceed through a trial and error process in which the control strategy of the single robot is repeatedly refined until the desired collective behavior is achieved (Brambilla et al., 2013). Therefore, analysis becomes an essential phase during the design of a swarm robotics system.

The aim of this report is to improve the modeling techniques available to analyze swarm robotics systems. Indeed, through modeling and analysis, designers can check whether a certain property of the systems holds, for which range of control parameters the property is present, and up to which level of noise the property still manifests itself. Furthermore, the ability to define proper mathematical models of swarm robotics systems allows designers to study their properties avoiding the burden of extensive physics-based simulations and experiments with real robotic platforms, tasks that in general are very time

consuming. In this work, we focus on modeling swarm robotics systems from a macroscopic viewpoint. Our approach is built on the formalism of Markov chains (Kemeny and Snell, 1976; Norris, 1997). We put particular emphasis in the definition of models based only on agents' control parameters, i.e, models that do not require to estimate or to assume the values of free parameters. The proposed technique is applied to two different case studies. Firstly, we analyze a collective decision-making system in which a swarm of robots needs to collectively agree on the shortest path between the two available paths in the classical double bridge problem (Goss et al., 1989; Montes de Oca et al., 2011). Secondly, we analyze a distributed communication protocol for heterogeneous robotics swarms that allows the swarm to achieve spatially targeted communication (Mathews et al., 2010a).

The reminder of this report is organized as follows. In Chapter 2, we first review the primary theoretical results of Markov chains used throughout the report; then, we introduce the modeling methodologies developed in this study. In Chapter 3, we apply the proposed techniques to the first case study of a collective decision system where robots are required to collectively agree on the most favorable action to perform. In Chapter 4, we move to the second case study, spatially targeted communication, and we define a Markov chain model to study the performance of an existing communication protocol for heterogeneous swarms of robots. Finally, we conclude with Chapter 5 by discussing the contributions of this report and providing future lines of research.

Chapter 2

Methodology

In this chapter, we briefly review the concepts underlying the theory of time-homogeneous Markov chain that are relevant for the understanding of this work. The material presented here focuses on finite absorbing Markov chains and summarizes the presentation of Norris (1997) and that of Kemeny and Snell (1976); for a deeper introduction to the topic of Markov chain the reader can refer to (Kemeny and Snell, 1976; Norris, 1997). We formally introduce the methodology utilized throughout the report to define macroscopic models of swarm robotics systems. We clarify the conditions in terms of agents' control strategy necessary to the application of this modeling approach and we summarize the primary macroscopic quantities that can be studied in this way.

2.1 Time-homogeneous Markov chain

Let $\mathbb{N}$ represent the set of naturals. Each element $i \in \Omega$ is called a *state* and the finite set $\Omega \subset \mathbb{N}$ is called the *state-space*. We say that the vector $\lambda = (\lambda_i : i \in \Omega)$ is a probability distribution on Ω if $0 \leqslant \lambda_i \leqslant 1$ for all $i \in \Omega$ and $\sum_{i \in \Omega} \lambda_i = 1$. We consider the *random variable* X with values in Ω and we set

$$\lambda_i = \mathbb{P}(X = i),$$

where the function $\mathbb{P}$ gives the probability of any event. The random variable X models a random state which takes value i with probability λ_i and λ defines the distribution of X.

Let $P = (p_{ij} : i, j \in \Omega)$ be a stochastic matrix, i.e., a matrix in which every row $(p_{ij} : j \in \Omega)$ is a probability distribution. We say that

the sequence of random variables $\{X_n : n \in \mathbb{N}\}$ is a *Markov chain* with *initial distribution* λ and *stochastic transition matrix* P if

i. X_0 has probability distribution λ;

ii. for $n \in \mathbb{N}$, conditional on $X_n = i$, X_{n+1} has probability distribution $(p_{ij} : j \in \Omega)$ and is independent of $X_0, \ldots, X_{n-1}$.

The first condition states that the probability distribution λ defines the distribution of the initial state X of the chain at step $n = 0$. The second condition defines the *Markov property* of a stochastic process: the lack of memory. More explicitly, the future state of a Markov process depends only on the present state and is independent of its past (Kemeny and Snell, 1976; Norris, 1997). That is, given $X_n = i$, the entry p_{ij} of the stochastic transition matrix P gives the probability that the next state X_{n+1} will be equal to j. Notice that P is independent of the time n, which results in the time-homogeneous property of the chain. In general, the probability that the state of the chain at step n will be equal to $j \in \Omega$ is

$$\mathbb{P}(X_n = j) = (\lambda P^n)j,$$

where the function $(\cdot)j$ returns the j-th element of a vector.

The state-space Ω of the Markov chain is divided into *equivalence* classes of states — also known as *communicating* classes — in which one can go from any state to any other state in the same class (not necessarily in one step). If, from a state i of the equivalence class $\mathcal{A}$ it is not possible to go to a state j of any other equivalence class, $\mathcal{A}$ is said to be *closed* and their states are called *ergodic* states, otherwise the states are considered *transient*. Under this conditions, we have that, if a process leaves a transient class it can never return to this class, while if it once enters a closed class, it can never leave it. In particular, an ergodic class that consists of a single state i results in a state that cannot be left and that is called *absorbing* state. An absorbing state i is thus characterized by $p_{ii} = 1$ and $p_{ij} = 0, \forall j \neq i$. A Markov chain that has at least one absorbing chain is called absorbing and the process on the chain will eventually be trapped in one absorbing state.

2.2 Macroscopic models of swarm of robots

Let us consider a system composed by a swarm of N agents that are driven by some set of control rules. In order to apply the theory of

Markov chains, either from a microscopic or macroscopic viewpoint, it is essential that the system fulfills the Markov property, i.e., the lack of memory, and thus, that the agent's control strategies do not require memory[1].

Among the different design approaches for swarm intelligence systems (see Brambilla et al., 2013), the behavior-based design approach results in control strategies that are reactive, often represented by probabilistic finite state machines (PFSMs), and which are well-suited to Markov theory (Martinoli et al., 2004). Control strategies designed in such a way are characterized by a finite set of agent's states and a set of transition probabilities among pairs of states. At any time, an agent driven by a (purely) reactive strategy decides the next action to perform on the basis of its current state and sensory inputs. We thus focus on the task of modeling such behavior-based strategies.

We consider each agent of the swarm to be characterized by a finite set of states possibly with several attributes. The combinations of state and attributes define the set of quantities that we need to keep track of in order to define a mathematical model. We denote by $\mathcal{S} = \{s_1, \ldots, s_m\}$ the set of agent's attribute-state combinations, henceforth simply state, and we assume the individual agent to be at any time in a unique state $s \in \mathcal{S}$. We refer to the state at time t of the agent i with the notation $x_i(t)$. In principle, we would like to study the time evolution of the states of the whole swarm of N agents and thus of $x_1(t), \ldots, x_N(t)$. This task implies the definition of a microscopic model of the system whose state-space is given by the Cartesian product of all agents' state-spaces. However, as soon as the size N of the swarm and the set of states $\mathcal{S}$ increase, such a microscopic model is subject to a state-space explosion that compromises its mathematical tractability. That is, the model becomes so complex that any analysis would be computationally unfeasible. In contrast, macroscopic models are less affected by the state-space explosion problem and allow designers to broaden the class of systems that can be studied following this approach.

A macroscopic model of a system composed of N agents, each of which characterized by m states, is equivalent to a model of a process in which N balls are repeatedly redistributed into m different bins. Such

[1]It is always possible, even if sometimes may not be computationally tractable, to model a stochastic process with finite memory by means of a Markov chain of order m.

a model needs to keep track of the number of balls that are in each bin, which in terms of the swarm corresponds to $N_1(t), \ldots, N_m(t)$. In other terms, a macroscopic model completely disregards which agent is in a particular state and only focus on the number of agents in a given state. Assuming $N \gg m$, the macroscopic viewpoint on the system strongly delays the state-space explosion effect and allows designers to study swarm intelligence systems with several different states and thousands of agents. As we will show in Chapter 3 and Chapter 4, a further reduction of the state-space can be achieved in some cases through intuition and insights into the dynamics of the system by eliminating unfeasible attribute-state configurations. Nonetheless, for increasing values of N and m, the macroscopic model will eventually become intractable as well and the study of the system will demand different mathematical approaches.

In order to study swarm intelligence systems, we propose the use of absorbing Markov chains to model the time evolution of the vector $(N_1, \ldots, N_m)$ of the number of agents in each state. We employ absorbing models for two different classes of scenarios:

i. Processes that end in a finite time with a single known outcome (e.g., spatially targeted communication protocols end when a communication channel is established) are modeled by chains with a single absorbing state. The process ends when it reaches the absorbing state and the focus of the model is on the time to absorption.

ii. Processes that continue for an infinite amount of time but that reach a particular absorbing macroscopic state among different possibilities (e.g., the opinion dynamics in a collective decision system evolves for an infinite time but in finite-size system consensus is reached in finite time) are modeled by chains with more absorbing states. The process ends when it reaches one of the possible absorbing states and the focus of the model is to predict both the time to absorption and the absorption probabilities for each state.

The analysis and definition of a macroscopic model as an absorbing Markov chain do not depend on the number of absorbing states[2]. In

[2]Indeed, the only difference between a chain with a single absorbing state and a chain with more absorbing states lies in the absorption probabilities that can be computed only for the second chain.

the reminder of this section we thus consider a process with more than one absorbing state.

We take into account an absorbing, time-homogeneous Markov chain $\{X_n : n \in \mathbb{N}\}$ whose state space $\Omega = \{1, \ldots, w\}$ enumerates the subset of feasible configurations of the vector $(N_1, \ldots, N_m)$. The conditions under which a configuration is considered feasible depend on the particular case study. Nonetheless, we can in general assume the conservation of the number of agents, i.e., $\sum_{i=1}^{m} N_i = N$, which results in a strong reduction of the state space Ω. At any time step n, a process on the chain in the state $X_n = i$ corresponds to a particular macroscopic configuration $(N_1^i, \ldots, N_m^i)$ of the swarm intelligence system under consideration. The state-space Ω includes a number r of absorbing states. These states identify macroscopic configurations of the swarm that are relevant to the purposes of the model (e.g, establishment of a communication link, consensus on a particular opinion). Finally, by properly setting the entries of the stochastic transition matrix P, the Markov chain model is completely defined and it can be used to study the system.

The definition of the transition probabilities between pairs of states of the chain depends on the particular swarm intelligence system under consideration. Nonetheless, we can still provide some general assumptions and guidelines on the definition of P. We consider the nature of the interactions among the agents of the swarm. In the most trivial case, the agents in the swarm do not interact with each other, as will be the case of the model in Chapter 4, and are thus considered completely independent. In more complex scenarios, agents in the swarm have pairwise interactions or even higher order interactions. An example of this scenario is provided by the model developed in Chapter 3. In this case, to define a valid Markov chain it is essential to assume a uniform distribution of the states of the agents that may interact with each other. That is, an agent has the same probability to interact with any other agent in the same physical region (e.g., in the nest, along a path). Therefore, only the proportions of the different states are of interest when defining a model. This assumption is called the *well-mixed* assumption — also known as the *well-stirred* assumption — and is also considered one of the primary mechanisms underlying natural swarms (Nowak, 2006).

Once the Markov chain is completely defined, the stochastic transi-

tion matrix P can be used to answer a number of questions regarding the performance of the system. To this end, the first step of the analysis consists in finding the *canonical* form of P (see Kemeny and Snell, 1976). By a reordering of the states in Ω it is possible to write the stochastic transition matrix P as

$$P = \left(\begin{array}{c|c} I & O \\ \hline R & Q \end{array} \right).$$

The matrix Q is a $(w-r) \times (w-r)$ matrix of transition probabilities between transient states, R concerns the probabilities to go from a transient state to an absorbing state, O consists entirely of 0's, and the identity matrix I whose size is $r \times r$ defines the absorbing states. Notice that, as a consequence of the absorbing nature of the Markov chain, the entries of Q^n tend to 0 as $n \to \infty$. That is, the probability to find the process in a transient state vanishes as the time passes. This result provides sufficient conditions for the existence of the inverse of $I - Q$, called *fundamental* matrix, and given by

$$F = (I - Q)^{-1} = I + Q + Q^2 + \cdots = \sum_{k=0}^{\infty} Q^k.$$

Each entry f_{ij} of the fundamental matrix F gives the mean time that a process started in the transient state i spends in the transient state j.

By means of the canonical decomposition of P we can derive a number of interesting macroscopic quantities regarding the dynamics of the swarm. At first, we are interested in the probability that a system initially started in the state $X_0 = i$ will eventually be absorbed in the absorbing state $X_n = j$ for some time step n. Considering all possible initial states i and all possible absorbing states j, we have the set of absorption probabilities that is given in matrix form by

$$B = FR.$$

The entry b_{ij} thus provides the absorption probability for the pair of initial and absorbing states (i, j). Besides absorption probabilities, we may want to study the time to absorption, i.e., the number of steps necessary to the process on the chain to enter one absorbing state. Let us denote with ϑ the random variable that counts the number of steps to absorption. The expectation and the variance of ϑ are given by

$$\mathbb{E}[\vartheta] = F\xi,$$

$$\mathbb{V}[\vartheta] = (2F - I)\mathbb{E}[\vartheta] - \mathbb{E}_{sq}[\vartheta].$$

In the above formulas, the term ξ identifies a column vector of all 1's and the vector $\mathbb{E}_{sq}[\vartheta]$ corresponds to a version of $\mathbb{E}[\vartheta]$ with squared entries. The i-th entry of the vectors $\mathbb{E}[\vartheta]$ and $\mathbb{V}[\vartheta]$ gives, respectively, the expectation and the variance of the time to absorption for a system initially started in state i. Finally, we can derive the cumulative distribution function $G(\vartheta)$ of the time to absorption, and thus the probability mass function $g(\vartheta)$ as well, for a given initial state $X_0 = i$ as the infinite series

$$G(\vartheta) = 1 - \sum_{j \in \Omega} Q_{i,j}^n, \text{ for } n \to \infty. \tag{2.1}$$

The subcomponent $\sum_{j \in \Omega} Q_{i,j}^n$ of the above equation gives the probability that the process will be in a transient state at step n. The complement of this value provides the probability of entering the absorbing state prior to step n.

Chapter 3

Case study I: Collective Decision Making

In this chapter[1], we apply the methodology proposed in Chapter 2 to a collective decision making problem. This case study consists in a swarm of robots that has to collectively agree on the fastest path within the boundaries of the classical double bridge problem (Goss et al., 1989). We first introduce the swarm robotics system and discuss previous modeling approaches. Next, we formally introduce a Markov chain model of the collective decision system and we provide an analysis of its dynamics. Finally, we conclude with a discussion of the contribution and with a proposal for future directions of research.

3.1 Majority Rule with Differential Latency

We analyze a swarm robotics system originally proposed by Montes de Oca et al. (2011). Robots in the swarm need to collectively decide between two possible actions to perform, henceforth referred to as action A and action B. Actions have the same outcome but different execution times. The goal of the swarm is to reach consensus on the action with the shortest execution time. In particular, Montes the Oca et al. studied this system in a collective transport scenario where robots in the swarm need to transport objects from a *source* area to a *destination* area. To this end, robots can choose between two possible paths. This corresponds to perform action A or action B. The two paths differ in

[1]This chapter is based on the author's article (Valentini et al., 2013) published in the Proceedings of the European Conference on Complex Systems.

length and thus in the traversal time. Each robot in the swarm has an *opinion* for a particular path. Moreover, an object is too heavy for a single robot to be transported. A team of 3 robots is needed. The team collectively decides which path to take considering the opinion favored by the majority.

Opinion formation models, such as the majority-rule model by Galam (1986), allow us to study and analyze this kind of systems. Krapivsky and Redner (2003) provided an analytical study of the majority-rule model under the assumption of a well mixed[2] population of agents. Later, Lambiotte et al. (2009) extended the work of Krapivsky and Redner introducing the concept of *latency*. In the model of Lambiotte et al., when an agent switches opinion as a consequence of the application of the majority rule, it turns in a latent state for a latency period that has stochastic duration. A latent agent may still participate in voting, thus influencing other agents, but its opinion does not change as a result of the decision. This extension gives rise to a richer dynamics depending on the duration of the latency period. Based on these works, Montes de Oca et al. (2011) proposed the *differential latency* model where the duration of the latency period depends on the particular opinion adopted. After a decision, differently from the model of Lambiotte et al., the agents in the team become latent with a common latency period and are not involved in further voting until the end of the latency period. Montes de Oca et al. showed that the differential latency in the majority-rule model steers the agents towards consensus on the opinion associated to the shortest latency. These results have been applied to the study of the swarm robotics system described above by modeling actions of the robots as opinions and their execution times as the latency periods of different duration.

In the context of swarm robotics, a number of works have been devoted to the differential latency model. Montes de Oca et al. (2011) first proposed a fluid-flow analysis of this model — using a system of ODEs — aimed at studying the dynamics leading to consensus. This analysis, derived in the thermodynamic limit, deterministically predicts consensus as a function of the initial configuration of the system. However, in a finite population, random fluctuations may drive the system to converge to the long path, even when the fluid-flow model predicts

[2]In a well mixed population each agent has the same probability to interact with each other agent (Nowak, 2006).

that it should converge to the short one. Later, Scheidler (2011) extended the previous analysis using methods from statistical physics — e.g., master equation and Fokker-Planck equation — to derive continuous approximations of a system with a finite population size. With this approach, Scheidler was able to study the exit probability, i.e., the probability that the system eventually reaches consensus on the opinion associated to the shortest latency, and the expectation of the time necessary to reach consensus. Finally, Massink et al. (2012) provided a specification of the system using a stochastic process algebra. On the basis of this specification, the authors obtained a statistical model checking and a fluid-flow analysis.

Continuous approximations provide reliable predictions only when the number of robots is relatively large — e.g., thousands of robots. However, swarm robotics aims to design scalable control policies that operate for swarms of any size, ranging from tens to millions of robots. These models cover only the upper part of this range. Besides, using a continuous approximation model it is usually hard to derive statistics different from the expectation of a metric, which in turn, often gives a poor representation of the underlying distribution — e.g., when the variance is large compared to the expectation or when the distribution is not symmetric.

The aim of this work is to study the majority rule with differential latency with an approach able to cope with the limitations of previous approaches. Inspired by the work of Banisch et al. (2011), we use the formalism of absorbing, time homogeneous Markov chains with finite state space (Kemeny and Snell, 1976) to model the dynamics of the system. This approach allows us to consider swarms of finite size and to derive reliable estimations of both the exit probability and the distribution of the number of decisions necessary to reach consensus.

3.2 Markov Chain Model

We model the majority rule with differential latency in a system of M robots as an absorbing Markov chain (Kemeny and Snell, 1976). Robots can be *latent* or *non-latent*. Only non-latent robots, once grouped in a team of 3 members, take part to the decision-making mechanism. As in Montes de Oca et al. (2011), we consider a scenario where the number k of latent teams is constant, and where the latency

period follows an exponential distribution whose expectation depends on the team's opinion. Without loss of generality, we consider the expected latency periods to be 1 for opinion A and $1/\lambda$ with $0 \leqslant \lambda \leqslant 1$ for opinion B. Moreover, we are interested in the number ϑ of applications of the majority rule, thus, we consider each application of the decision-making mechanism as one step of the process along the chain.

At each step ϑ we consider 3 stages:

i. A latent team becomes non-latent (it finishes its latency period).

ii. A new team of 3 robots is randomly formed out of the set of non-latent robots.

iii. The team applies the majority rule to decide the team's opinion. Next, it turns in a latent state.

We are interested in the evolution over ϑ of the number of robots with opinion A — the opinion associated to the shortest latency. Let $\mathbb{N}$ be the set of naturals. The state of the Markov chain is a vector $s = (s^l, s^n)$, where $s^l \in \{l : l \in \mathbb{N}, 0 \leqslant l \leqslant k\}$ is the number of latent teams with opinion A and $s^n \in \{n : n \in \mathbb{N}, 0 \leqslant n \leqslant M - 3k\}$ is the number of non-latent robots with opinion A. The state space of the Markov chain consists of m states, where $m = (k + 1)(M - 3k + 1)$ is the cardinality of the Cartesian product of the domains of s^l and s^n. By s_i and s_j we refer to two generic states. By s_a and s_b we refer to the consensus states in which the whole swarm agrees on opinion A and B, respectively. Notice that s_a and s_b are the *absorbing* states of the chain, that is, states that once reached can never be left (cf. Kemeny and Snell, 1976).

At the generic step ϑ, the process moves from $s(\vartheta) = s_i$ to $s(\vartheta + 1) = s_j$ following the aforementioned 3 stages. At stage i., a latent team finishes its latency period, becomes non-latent and disbands. The probability p_i that this team has opinion A is:

$$p_i = \frac{s_i^l}{s_i^l + \lambda(k - s_i^l)}. \tag{3.1}$$

The set of non-latent robots with opinion A increases of $c = 3$ units, if the disbanding team has opinion A; and of $c = 0$, otherwise. At stage ii., 3 random robots form a new team in the set of non-latent robots. We are interested in the probability q_i that the new team has a number

$0 \leqslant d \leqslant 3$ of preferences for opinion A. This probability is given by the hyper-geometric distribution

$$q_i(d;c) = \frac{\binom{s_i^n+c}{d}\binom{M-3k-s_i^n+3-c}{3-d}}{\binom{M-3k+3}{3}} \tag{3.2}$$

of the $M-3k+3$ preferences in the current set of non-latent robots, composed of s_i^n+c votes for opinion A and $M-3k-s_i^n+3-c$ votes for opinion B. At stage iii., the majority rule is applied and the outcome is determined by the value of d. Eventually, the process moves to the next state $s(\vartheta+1) = s_j$.

Equations (3.1) and (3.2) allow us to define the transition probabilities between each possible pair of states s_i and s_j. These probabilities are the entries of the stochastic transition matrix P, which completely defines the dynamics of a Markov chain (Kemeny and Snell, 1976). However, not all pairs of states define a feasible move of the process along the chain according to the rules of the system, i.e., not all pair of states are *adjacent*. Two states s_i and s_j are adjacent if $\Delta_{ij}s = (\Delta_{ij}s^l, \Delta_{ij}s^n) = s_j - s_i$ appears in the first column of the following table. The corresponding transition probability P_{ij} is given in the second column:

$(\Delta_{ij}s^l, \Delta_{ij}s^n)$	P_{ij}	stage 1)	stage 2)
$(-1,3)$	$p_iq_i(3-\Delta_{ij}s^n;3)$	A	$3B$
$(-1,2)$		A	$A2B$
$(0,1)$		A	$2AB$
$(0,0)$	$p_iq_i(3-\Delta_{ij}s^n;3)$	A	$3A$
	$+(1-p_i)q_i(\lvert\Delta_{ij}s^n\rvert;0)$	B	$3B$
$(0,-1)$	$(1-p_i)q_i(\lvert\Delta_{ij}s^n\rvert;0)$	B	$A2B$
$(1,-2)$		B	$2AB$
$(1,-3)$		B	$3A$

Columns three and four provide the corresponding events observed in stages i. and ii., respectively: the opinion of the robots in the next latent team finishing its latency period and the opinions of the robots that randomly form a new team in the set of non-latent robots. For values of $\Delta_{ij}s$ not included in column one, the transition probability is $P_{ij} = 0$.

The probabilistic interpretation of P is straightforward: at any step ϑ, if the process is in state $s(\vartheta) = s_i$ it will move to state $s(\vartheta+1) = s_j$ with probability P_{ij}. It is worth noticing that, being the consensus

states s_a and s_b two absorbing states, the probability mass of s_a and s_b is concentrated in the corresponding diagonal entries of P, that is: $P_{aa} = 1$ and $P_{bb} = 1$.

3.3 Analysis of Opinion Dynamics

In order to analyze the dynamics of the majority rule with differential latency, we consider the matrices Q, R, and N of Kemeny and Snell (1976) derived from the canonical decomposition of P (see Section 2.1). Q describes transitions between transient states, R gives the probability to move from a transient state to an absorbing state, and $N = (I-Q)^{-1}$ is the *fundamental* matrix with I being the identity matrix. From matrices Q, R, and N of the Markov chain model we study the behavior of the system. We validate the predictions of the model with the results of Monte Carlo simulations[3] averaged over 1000 independent runs for each choice of the parameters of the system.

First, we derive the exit probability $E(s_i)$, i.e., the probability that a system of M robots that starts in the initial configuration $s(\vartheta_0) = s_i$ reaches consensus on the opinion associated to the shortest latency — opinion A. This probability is given by the entries associated to the consensus state s_a of the product NR, which corresponds to the matrix of the absorption probabilities (Kemeny and Snell, 1976).

Figure 3.1 reports the predictions of the exit probability over the initial density $\delta = (3s_i^l + s_i^n)/M$ of robots favoring opinion A for several configurations of the system. As found by Scheidler (2011), the larger is the expected latency period $1/\lambda$ associated to opinion B, the smaller is the initial number of preferences for opinion A such that the exit probability is $E(s_i) > 0.5$. Moreover, when the number of robots M increases, the exit probability approaches a step function around the critical density. It is worth noticing that the Markov chain model predicts the outcome of the simulations with great accuracy, regardless of the number of robots. A result that in general cannot be achieved

[3]We simulated two sets of robots: latent teams characterized by an opinion and a latency, and non-latent robots described only by their opinions. The simulation proceeds as follow until consensus is reached: 1) the latent team having minimum latency is disbanded and its component robots are added to the set of non-latent robots, 2) 3 robots are randomly sampled from the set of non-latent robots and the majority rule is applied among them, 3) the new team is added to the set of latent teams and its latency is drawn from the exponential distribution according to the team's opinion.

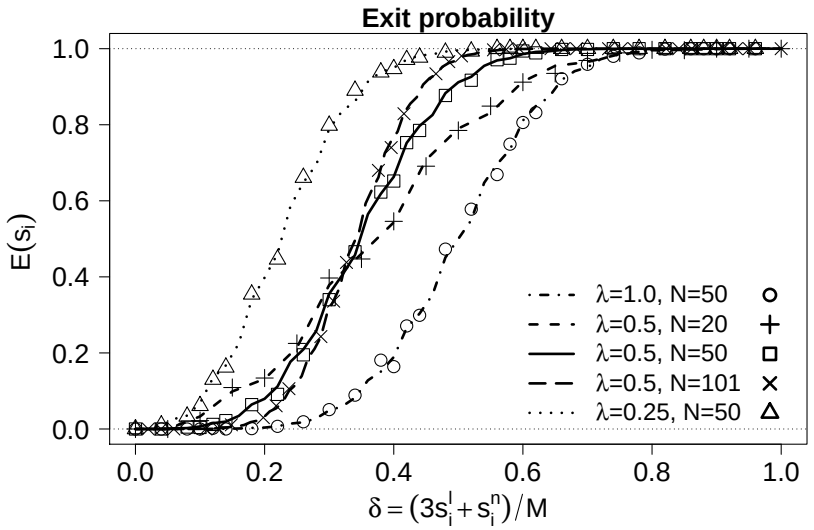

Figure 3.1: Probability $E(s_i)$ to reach consensus on opinion A versus initial proportion δ of robots favoring that opinion for different settings of the system: $(M = 20, k = 6, \lambda = 0.5)$, $(M = 50, k = 16, \lambda \in \{1, 0.5, 0.25\})$ and $(M = 101, k = 33, \lambda = 0.5)$. Lines refer to the predictions of the Markov chain model, symbols refer to the average results of 1000 Monte Carlo simulations for each initial configuration of the system.

using continuous approximations approaches.

Next, we analyze the number τ of applications of the majority rule necessary to reach consensus. As stated above, we consider each step along the chain as one application of the decision-making mechanism. The expected value of τ is given by $\hat{\tau} = \xi N$, where ξ is a column vector of all 1s. The entries of $\hat{\tau}$ correspond to the row sums of the fundamental matrix N. In turn, N gives the mean sojourn time for each transient state of a Markov chain (Kemeny and Snell, 1976), that is, the expected number of times that a process started in state $s(\vartheta_0) = s_i$ passes from state s_j. The variance of τ is given by $\hat{\tau}_2 = (2N - I)\hat{\tau} - \hat{\tau}_{sq}$, where I is the identity matrix and $\hat{\tau}_{sq}$ is $\hat{\tau}$ with squared entries (Kemeny and Snell, 1976).

Figures 3.2 and 3.3 show the predictions of the expectation $\hat{\tau}$ and the variance $\hat{\tau}_2$ of the number of decisions necessary before consensus for a system with $M = 50$ robots. Again, the Markov chain model predicts the Monte Carlo simulations with great accuracy. Similarly to the findings of Scheidler (2011) for the consensus time, the value of $\hat{\tau}$ is maximum near the critical density of the initial number of robots favoring opinion A. However, the expected number of decisions, which

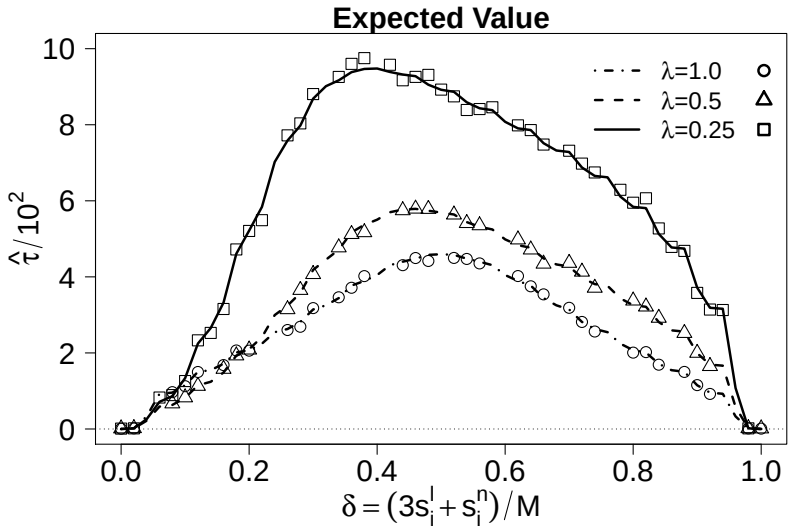

Figure 3.2: Expectation $\hat{\tau}$ of the number τ of applications of the majority rule necessary to reach consensus for a system of $M = 50$ robots, $k = 16$ teams and $\lambda \in \{1, 0.5, 0.25\}$. Lines refer to the predictions of the Markov chain model, symbols refer to the average results of 1000 Monte Carlo simulations for each initial configuration of the system.

is related to the consensus time in (Scheidler, 2011), is not a reliable statistics for this system. Indeed, the variance $\hat{\tau}_2$ is about three orders of magnitude larger than $\hat{\tau}$.

Finally, we derive the cumulative distribution function $P(\tau \leqslant \vartheta; s_i)$ of the number of decisions before consensus as well as its probability mass function $P(\tau = \vartheta; s_i)$. From a swarm robotics perspective, we are interested in the dynamics of a system initially unbiased, i.e., a system that starts with an equal proportion of preferences for the opinions A and B. Let $s(\vartheta_0) = s_u$ represents this initial unbiased configuration. Recalling that Q is the matrix of the transition probabilities for the transient states, we have that the entries Q^{ϑ}_{uj} of the ϑ^{th} power of Q give the probabilities to be in the transient state s_j at step ϑ when starting at s_u. Thus, the row sum of the u^{th} row of Q^{ϑ} gives the probability to still be in one of the transient states. From this probability, we can derive the cumulative distribution function $P(\tau \leqslant \vartheta; s_u)$ simply by computing the series $\{1 - \sum_j Q^{\vartheta}_{uj}\}$ for values of ϑ such that $Q^{\vartheta} \to 0$.

Figure 3.4 shows the cumulative distribution function $P(\tau \leqslant \vartheta; s_u)$ for a system of $M = 50$ robots that starts unbiased. Obviously, the longer is the latency period $1/\lambda$ of opinion B, the larger is the number of applications of the majority rule necessary to reach consensus.

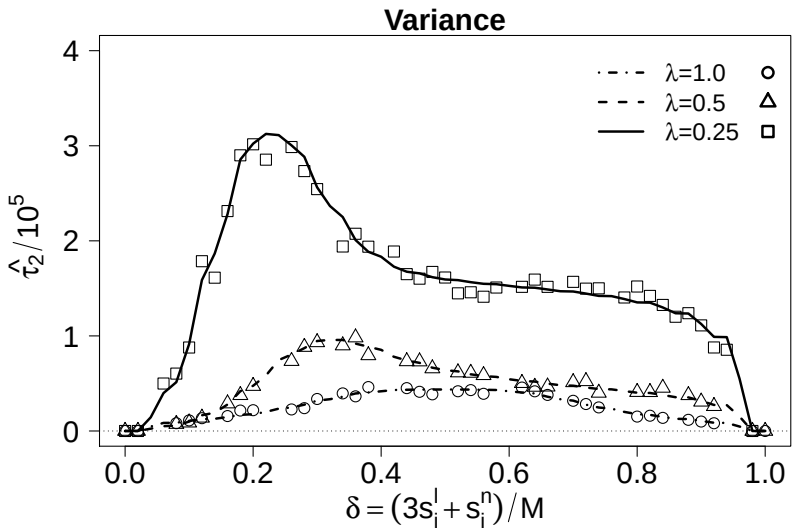

Figure 3.3: Variance $\hat{\tau}_2$ of the number τ of applications of the majority rule necessary to reach consensus for a system of $M = 50$ robots, $k = 16$ teams and $\lambda \in \{1, 0.5, 0.25\}$. Lines refer to the predictions of the Markov chain model, symbols refer to the average results of 1000 Monte Carlo simulations for each initial configuration of the system.

Figure 3.5, provides the probability mass function $P(\tau = \vartheta; s_u)$, together with details of mode, median, and mean values of τ. As can be seen, the values of the mode, median and mean statistics diverge for increasing values of the ratio $1/\lambda$ of the two expected latency periods. Moreover, when $1/\lambda \to \infty$ the shape of the distribution $P(\tau = \vartheta; s_u)$ tends to a flat function, thus revealing that the variance dominates the system.

3.4 Conclusion

We designed an absorbing Markov chain model for collective decisions in a system with a finite number of robots based on the majority rule with differential latency. Using our model, we derived the probability that a system of M robots reaches consensus on the opinion associated to the shortest latency period, and the distribution of the number of applications of the majority rule necessary to reach consensus. This latter reveals that the system is characterized by a large variance of the number of decisions necessary before consensus, and thus, that its expected value, which was mainly adopted in previous studies, is a poor statistic for this system. In contrast to continuous approximations, we

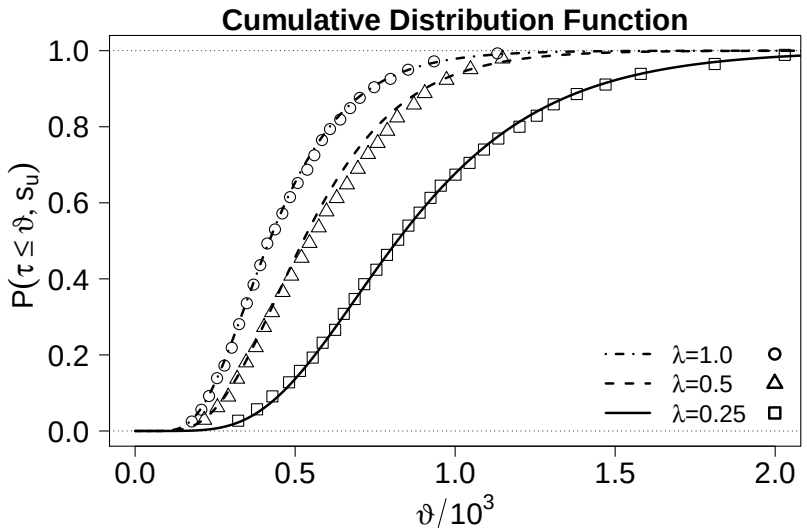

Figure 3.4: Cumulative distribution function $P(\tau \leqslant \vartheta; s_u)$ of the number τ of applications of the majority rule necessary to reach consensus for a system of $M = 50$ robots, $k = 16$ teams and $\lambda \in \{1, 0.5, 0.25\}$. Lines refer to the predictions of the Markov chain model, symbols refer to the average results of 1000 Monte Carlo simulations for each initial configuration of the system.

explicitly model the state space of the system — which is discrete — and the transition probabilities governing its dynamics. This approach allows us to always derive reliable predictions of a system regardless of its size.

Our contribution is relevant from a swarm robotics perspective not only for the reliability of its predictions; but also, because it allows us to perform a deeper analysis of the system. The analysis of our Markov chain model, with particular regard to the distribution of the number of decisions necessary to consensus, gives the possibility to perform statistical inference on certain interesting aspects of the system. Moreover, the approach can be easily extended to other voting schemata, allowing the comparison at design time of different choices for the decision-making mechanism.

In real-robot experiments, latency periods are unlikely to be exponentially distributed. Moreover, it is hard to ensure a constant number of teams in time. These assumptions represent hard constraints for a swarm robotics system. Massink et al. (2012) propose to cope with these constraints modeling the latency period with an Erlang distribution. We plan to extend our approach in a similar way and to validate

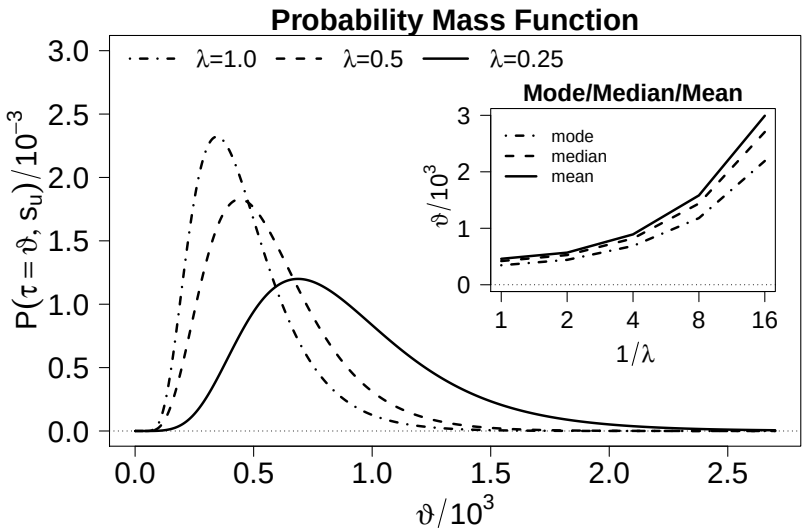

Figure 3.5: Probability mass function $P(\tau = \vartheta; s_u)$ with details of mode, median and mean values of the number τ of applications of the majority rule necessary to reach consensus for a system of $M = 50$ robots, $k = 16$ teams and $\lambda \in \{1, 0.5, 0.25\}$. Lines refer to the predictions of the Markov chain model, symbols refer to the average results of 1000 Monte Carlo simulations for each initial configuration of the system.

the resulting model with physics-based simulations and real-robot experiments.

Chapter 4

Case study II: Communication Protocols

In this chapter[1], we show the application of the methodology proposed in Section 2 to a different case study: *spatially targeted communication* (Mathews et al., 2010a). We first review the mechanism underlying the considered swarm robotics system. Next, we formally introduce a Markov chain model of the communication protocol. We provide an analysis of the performance of the systems through the Markov chain model by comparing its predictions with both physics-based simulations and real robot experiments. Finally, we conclude the chapter with a discussion of the contribution and with a proposal for future directions of research.

4.1 Spatially Targeted Communication: One-to-One link

We analyze spatially targeted communication — distributed and situated communication protocols originally proposed by Mathews et al. (2010a). These protocols have been specifically designed to permit situated communication in heterogeneous robotic swarms composed of ground based wheeled robots and aerial robots in which members of the swarm cannot be uniquely identified. In such scenarios, members of the swarm of wheeled robots can be supervised by aerial robots that, exploiting their superior perception of the environment, can precisely

[1]This chapter is based on an author's ongoing research in collaboration with the authors of (Mathews et al., 2010a).

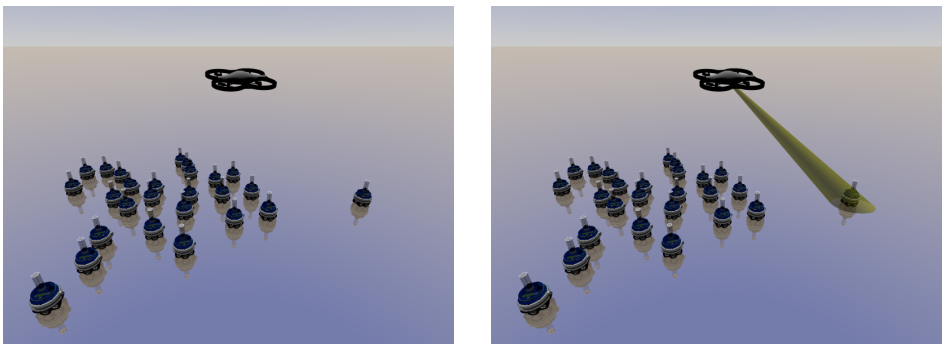

(a) (b)

Figure 4.1: Illustration of a possible scenario of a one-to-one spatially targeted communication link. (a) the aerial robot perceive a wheeled robot to far from the swarm and begins the establishment of a communication link with the one-to-one selection algorithm. (b) the one-to-one communication link has been established: the aerial robot can now instruct the wheeled robot on the position of the swarm. (Images reproduced courtesy of N. Mathews)

instruct specific robots of the swarm. Spatially target communication allows an aerial robot to establish a communication channel with a particular robot on the ground and to extend this channel to a cohesive group of robots (Mathews et al., 2010a,b). In Figure 4.1 for example, after establishing a one-to-one communication link, the aerial robot instructs a lost wheeled robot about the current position of the swarm.

Mathews et al. (2010a) designed two different protocols for spatially targeted communication that differ in the nature of the communication channel established: on the one hand, a one-to-one link between an aerial robot and a wheeled robot; on the other hand, a one-to-many link between an aerial robot and a group of wheeled robots. In both cases, robots in the swarm achieve spatially targeted communication by iteratively displaying different colors through their LEDs and by perceiving them with their on-board cameras. Depending on the color displayed at each iteration by the aerial robot, some wheeled robots on the ground may understand that are not the targets of the communication, and therefore, they exit the process.

Here, we focus on the study of the performance of the one-to-one selection algorithm that allows a heterogeneous swarm of robots to establish a one-to-one communication link between a pair of members. We formally describe the distributed selection process from a macroscopic viewpoint by means of an absorbing, time-homogeneous Markov

chain (Kemeny and Snell, 1976). The Markov chain model is subject to validation with both physics-based simulations and real robot experiments. The robotic platforms employed in this study are the commercial aerial robot *AR.Drone* (Bristeau et al., 2011), and the ground-based wheeled robot *marXbot* (Bonani et al., 2010). The physics-based simulator chosen for the study is the open-source, multi-robot simulator *ARGoS* (Pinciroli et al., 2012).

Given a set $C := \{c_1, \ldots, c_s : s \geq 3\}$ of signals available to both robot types, the aerial robot can establish a spatially targeted communication with a particular wheeled robot through an iterative selection process. Henceforth, we refer to this particular wheeled robot as the *target* robot. The set C consists of the designated signal c_1 used to initialize and terminate the elimination process and the subset $C_s := \{c_2, \ldots, c_s\}$ that is allocated to the iterative component of the elimination process. The nature of the signals depends on the particular robotic platform adopted. In our case the robots exploit their on-board LEDs and cameras to visually communicate by means of colors. Moreover, we assume that the aerial robot has already selected a target robot with which it wishes to communicate.

Let consider a scenario where a one-to-one communication link has to be formed in a swarm of robots having signals $C := \{red, blue, green\}$ and thus $C_s := \{blue, green\}$. The aerial robot first attracts the attention of all wheeled robots in visual range by signaling *red*, the SOS signal. All wheeled robots able to perceive the SOS signal register to the iterative selection process by replying with *blue*. The aerial robot responds to this initial registration with a matching handshake using *blue*. Next, the iterative selection process starts. At each iteration, every wheeled robot that is still part of the selection process randomly chooses and illuminates a color from the set C_s. The aerial robot reacts by illuminating its LEDs to match the color chosen by the target robot. At the end of the iteration, only those robots whose color matches that of the aerial robot remain part of the selection process. The wheeled robots which are not part of the selection process do not illuminate any color. The iterative selection process continues until the target robot is the only illuminated robot. At this point, the aerial robot indicates the termination of the selection process to the target robot by repeating *red* signal. Finally, the target robot acknowledges this signal by matching the aerial robot's color. The aerial and the target robots

have now established a spatially targeted communication link.

4.2 Markov Chain Model

Given a swarm of N wheeled robots, we are interested in the time necessary for an aerial robot to establish a one-to-one communication link with a particular wheeled robot. Here we consider time to be discrete due to (i) the behavioral nature of the robot controllers that are discrete and defined using FSMs (cf. Martinoli et al., 2004), and (ii) because the actual time spent by the swarm in solving the task heavily depends on the particular robotic platform and communication medium adopted, details that we want to abstract from at this stage of the study while focusing on the performance of the algorithm.

In our model, each time step n corresponds to one iteration of the selection process. An iteration of the selection process comprises the completion by all the (remaining[2]) wheeled robots of the sub-tasks: choosing and displaying a random color; and assessing the match or mismatch of this with the color showed by the aerial robot. It is worth noticing that the cost in terms of time of each iteration is independent of the number of remaining robots involved in the process since the computation is done in parallel by each robot.

The one-to-one selection algorithm described so far gives rise to a *memory-less* process. Indeed, at each iteration, a wheeled robot acts only on the basis of its current state and of the color displayed by the aerial robot. Thus, the future state of the robot is independent of its past, fulfilling the Markov assumption (Kemeny and Snell, 1976; Norris, 1997). Moreover, the reactive nature of the communication protocol introduces a synchronization among the robots in the swarm and allows us to consider the system from a macroscopic point of view. Given that the robots are synchronized, we can abstract from behaviors performed by individual robots and focus on subsets of the swarm performing the same behavior. As a consequence, we employ a time-homogeneous Markov chain to define a macroscopic model of the one-to-one selection algorithm.

Let $\mathbb{N}$ be the set of naturals. We define our process as a Markov

[2]This because, as iterations lapse, the number of wheeled robot actively involved in the selection process decreases, thus we have to consider only those robots that are still part of the process.

chain $\{X_n, n \in \mathbb{N}\}$ with $N+1$ states. Hence, one state more than the number of robots N. The random variable $X \in \Omega := \{1, \ldots, N+1\}$ represents the state of the process. Each state $X = x$ is characterized by the number $\eta_x = N + 1 - x$ of wheeled robots that are still part of the selection process. All the states of the chain are transients except for $X = N + 1$ which is absorbing. That is, once the process enters the state $X = N + 1$ it will never leave it. Intuitively, in the state $X = N + 1$ the number of wheeled robot involved in the process is $\eta_x = 0$, which corresponds to the completion of the selection task.

At the generic time step n, the process will move from the current state $X_n = i$ to the next state $X_{n+1} = j$ with a probability π_{ij}. We define the stochastic transition matrix $\Pi := (\pi_{ij} : i, j \in \Omega)$ of the Markov chain as

$$\begin{cases} \pi_{ij} = C_{\eta_i - 1}^{\eta_j - 1} p^{\eta_j - 1} (1-p)^{\eta_i - \eta_j}, & \text{for } i < N, i \leqslant j \leqslant N, \\ \pi_{ij} = 1, & \text{for } i \geqslant N, j = N+1, \\ \pi_{ij} = 0, & \text{otherwise.} \end{cases} \tag{4.1}$$

The first equation in (4.1), which concerns the transition probability between transient states, gives the probability π_{ij} that $\eta_i - \eta_j$ wheeled robots will leave the selection process during the transition $i \to j$. This probability follows a binomial distribution with parameters $\eta_i - 1$, the number of wheeled robots still part of the process minus the target, and $p = 1/|C_S|$, the probability for one of them to select the same color as the target robot. The second equation in (4.1) says that the transition from $X_n = N$ to $X_{n+1} = N + 1$ will occur with probability $\pi_{ij} = 1$. This choice models the final handshake between the aerial robot and the target robot at the last iteration of the selection algorithm. Besides, it also tells that the state $X = N + 1$ is absorbing. Finally, the third equation in (4.1) says that all other possible transitions will never occur.

4.3 Model Analysis

The stochastic transition matrix Π formally describes the dynamics of the process and allows us to study the performance of the one-to-one selection algorithm. Let us introduce the random variable ϑ whose meaning is defined as the number of iterations necessary to select the target robot. To study the distribution of ϑ, we consider the matrices Γ and Σ derived from the canonical decomposition of Π (Kemeny and Snell, 1976). Γ provides the transition probabilities between transient

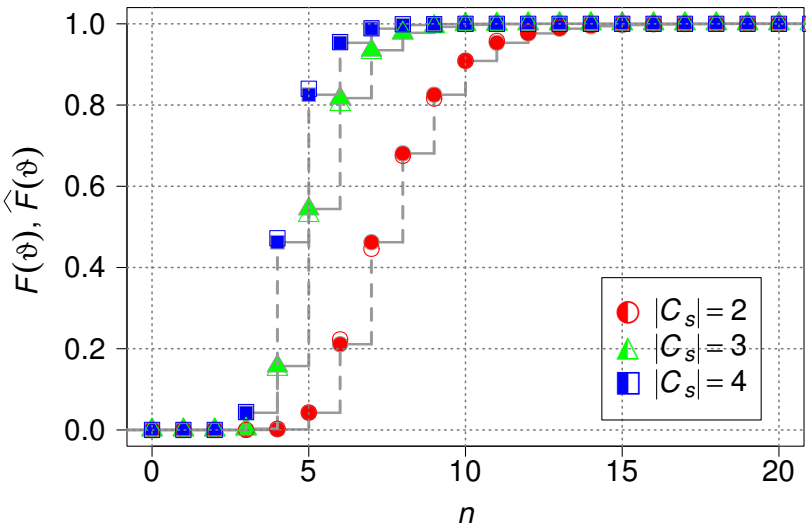

Figure 4.2: Cumulative distribution function $F(\vartheta)$ plotted against its empirical counterpart $\hat{F}(\vartheta)$ acquired from simulations for swarm size $N = 50$. Lines with filled symbols represent predictions of the model, empty symbols represent the result of 1000 physics-based simulations.

states. The fundamental matrix Σ, given by $(I - \Gamma)^{-1}$ with I being the identity matrix, gives the expectation of the number of visits to each transient state. In what follows, we first validate the accuracy of the Markov chain model by comparing model predictions with data acquired from physics-based simulations and real robots experiments. Then, we use the model to analyze the performance of the process through a scalability study of the one-to-one selection algorithm.

Model Validation. We validate the accuracy of the model predictions by comparing the distribution of ϑ with the empirical data acquired from simulation. The cumulative distribution function $F(\vartheta) = P(\vartheta \leqslant n, x_0)$ can be obtained for a given initial state $X_0 = x_0$ as the infinite series

$$P(\vartheta \leqslant n, x_0) = 1 - \sum_{j \in \Omega} \Gamma^n_{x_0,j}, \text{ for } n \to \infty.$$

Indeed, $\sum_{j\in\Omega} \Gamma^n_{x_0,j}$ gives the probability that the process will still be in a transient state at step n. The complement of this value provides the probability to be absorbed prior to step n.

Figure 4.2 shows the cumulative distribution function $F(\vartheta)$ provided by the model together with the empirical distribution function $\hat{F}(\vartheta)$

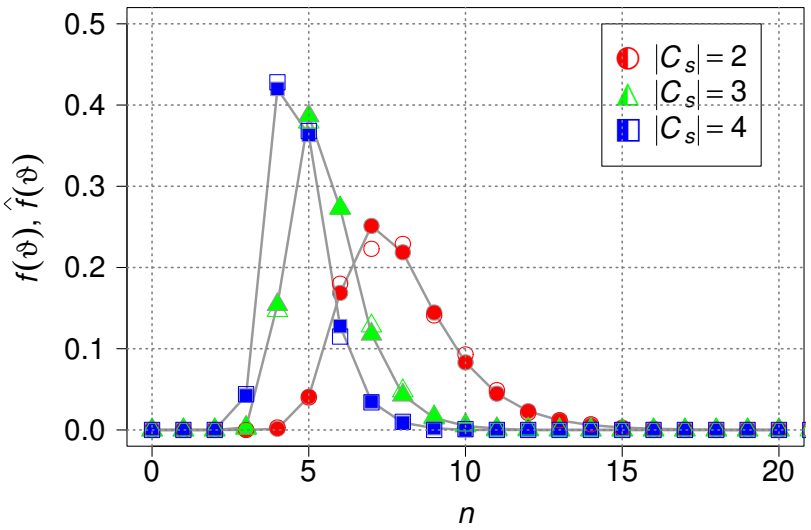

Figure 4.3: Probability mass function $f(\vartheta)$ plotted against its empirical counterpart $\hat{f}(\vartheta)$ acquired from simulations for swarm size $N = 50$. Lines with filled symbols represent predictions of the model, empty symbols represent the results of 1000 *physics-based simulations.*

obtained from 1000 simulation runs for $N = 50$ and different values of $|C_s|$. The theoretical predictions of the model are shown to closely match the empirical observations independently of the number of colors $|C_s|$ available to the swarm of robots.

Once computed $F(\vartheta)$ it is also interesting to have a look at $P(\vartheta = n, x_0)$, the probability mass function of ϑ, since it provides a better view of the performance of the system. As shown in Fig.4.3, the probability mass function is right-skewed, with the probability mass concentrated on the left side of the figure. This implies the possibility of long executions of the one-to-one selection algorithm. However, as we increase the number of colors available to the swarm, the variance of ϑ shrinks considerably reducing the probability mass under the right tail, and thus, the occurrence of long executions.

Real Robot Experiments. We performed a series of experiments with real robots using an AR.Drone placed on top of $N \in \{2, 4, 6, 8, 10\}$ marXbots distributed in a 1 m x 1.5 m arena with $|C_S| = 2$ colors. For each value of N, we have done 30 repetitions of the one-to-one selection algorithm. In each run, the marXbots are placed in the arena with random orientations and positions together with a light source identify-

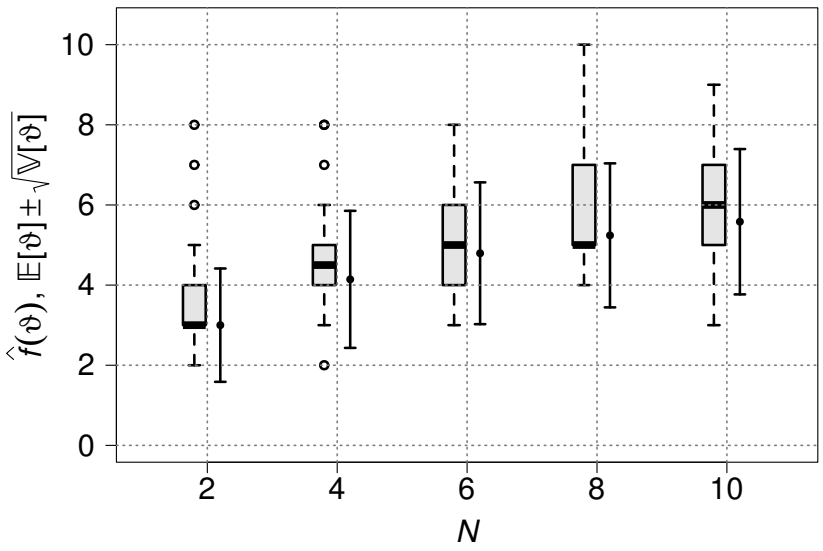

Figure 4.4: Empirical probability mass function $\hat{f}(\vartheta)$ computed from real robot experiments with varying sizes of N and $|C_s| = 2$. Boxplots depict 30 runs, errorbars correspond to predictions of the model ($\mathbb{E}[\vartheta] \pm \sqrt{\mathbb{V}[\vartheta]}$).

ing the point-of-interest (hereafter refered to as POI) to the AR.Drone. The AR.Drone chooses as target robot the closest marXbot to the POI. Fig. 4.4 shows the results of these experiments together with predictions of the Markov chain model. The latter, plotted as errorbars, show the expectation value $\mathbb{E}[\vartheta]$ and the standard deviation $\sqrt{\mathbb{V}[\vartheta]}$ of the number of iterations necessary to establish a one-to-one communication link.

From matrices Γ and Σ, we compute the expectation $\mathbb{E}[\vartheta]$ and the variance $\mathbb{V}[\vartheta]$:

$$\mathbb{E}[\vartheta] = \xi\Sigma, \tag{4.2}$$

and

$$\mathbb{V}[\vartheta] = (2\Sigma - I)\mathbb{E}[\vartheta] - \mathbb{E}_{sq}[\vartheta]. \tag{4.3}$$

In Eq. (4.2), the term ξ consists in a column vector of all 1's. That is, for every (transient) initial states, the expectation value of ϑ is given by the row sum of the fundamental matrix Σ. In Eq.(4.3), I corresponds to the identity matrix and $\mathbb{E}_{sq}[\vartheta]$ is $\mathbb{E}[\vartheta]$ with squared entries.

By means of Eq (4.2) and Eq. (4.3), we compare the empirical distribution of the results obtained from real robot experiments with the predictions of the Markov chain model defined in (4.1). As shown in Figure 4.4, the agreement between empirical observations and theoreti-

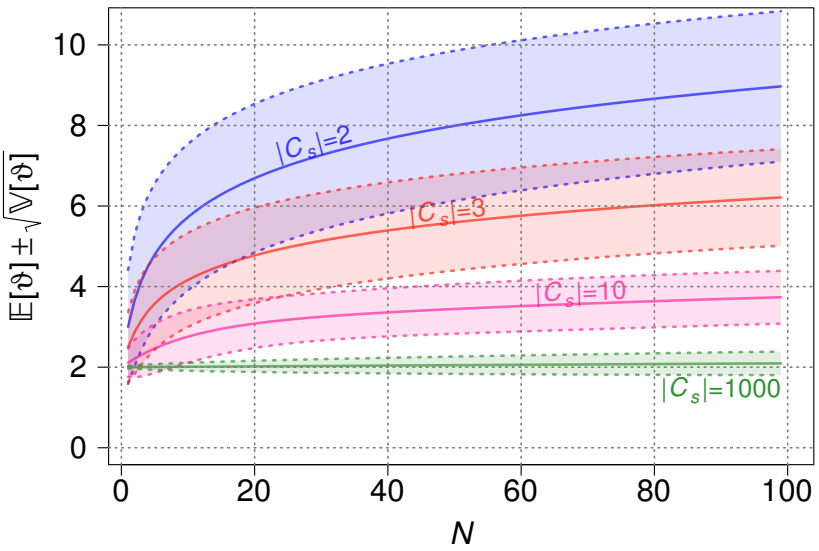

Figure 4.5: Scalability study based on data produced by the model for large N and $|C_s|$. The shaded areas correspond to $\mathbb{E}[\vartheta] \pm \sqrt{\mathbb{V}[\vartheta]}$.

cal predictions is rather high. In particular, the theoretical expectation of ϑ approximates well the median value showed in the boxplots of real robots experiments. Moreover, real robots experiments show a variance larger than the theoretical variance for the greatest values of N, (namely for $N \in \{6, 8, 10\}$). We argue that this discrepancy is due to the limited number of observations gathered during real robots experiments, and thus negligible. While the size of N in this set of experiments was increased by a factor of 5 (i.e., form 2 to 10 marXbots), the average number of iterations only experienced an increase by a factor of 1.6 (i.e., from 3.8 for 2 marXbots to 6.1 iterations for 10 marXbots). This may indicate promising scaling properties of the one-to-one selection algorithm. In the next section, we further investigate this property using the Markov chain model.

Scalability Study. We use Eq. (4.2) and Eq. (4.3) to study how the distribution of the expectation of ϑ scales for increasing values of N and $|C_s|$. As the results presented in Fig. 4.5 show, $\mathbb{E}[\vartheta]$ is characterized by a logarithmic trend indicating high scalability for increasing values of N. The same trend applies to the variance of ϑ: given a certain value for $|C_s|$, $\mathbb{V}[\vartheta]$ grows logarithmically for increasing N.

The results also show that as $|C_s|$ increases, the variance of ϑ decreases considerably, increasing the reliability of the expectation as

an aggregated indicator of the performance of the process (e.g., compare the width of the shaded areas for $N = 80$ between $|C_s| = 2$ and $|C_s| = 1000$ in Fig. 4.5). As it turns out, in the limit of $|C_s| \to \infty$, the probability for a wheeled robot to randomly choose the same color as the target robot tends to zero. The variance of ϑ vanishes and the one-to-one selection algorithm approaches a deterministic behavior that lasts exactly for 2 iterations: at the first iteration all wheeled robots except the target robot are eliminated, while in the second iteration the target robot performs the final handshake with the aerial robot.

4.4 Conclusion

In this chapter, we introduced an absorbing, time-homogeneous Markov chain model of a distributed communication protocol for a heterogeneous swarm of robots. We showed the applicability of the modeling methodology proposed in Chapter 2 to a different case study from that of collective decision-making provided in Chapter 3. The Markov chain model developed for the one-to-one spatially targeted selection algorithm has been extensively validated through both physics-based simulations and real robots experiments. The predictions of the model closely match the behavior of the robots in the experiments, and thus, can be reliably used at design time. Furthermore, the analysis of the Markov chain model tells us that the one-to-one selection algorithm has promising scalability properties. Indeed, the expectation of the number of iterations necessary to establish a one-to-one link is shown to scale logarithmically with the number of wheeled robots involved in the process.

In the near future, we plan to extend the modeling approach introduced here to study the performance of the one-to-many communication protocol introduced by Mathews et al. (2010a).

Chapter 5

Conclusions

In this report, we proposed a modeling methodology to define macroscopic mathematical models of swarm intelligence systems. The proposed approach is based on the use of the formalism of time-homogeneous Markov chains (Kemeny and Snell, 1976; Norris, 1997) and, in particular, through finite state-space, absorbing Markov chains. This approach allows designers to define tractable mathematical models of a swarm intelligence systems and to analyze in depth their dynamics and performance. By defining the Markov chain models from a macroscopic viewpoint, it is possible to describe systems with finite number of agents without demanding for continuous approximations. With respect to a microscopic model, the macroscopic model is characterized by a reduced state-space. This feature delays the state-space explosion effect and strongly broadens the family of swarm intelligence systems that can be studied with this approach. We have illustrated the applicability of the proposed modeling methodologies to two different case studies in swarm robotics.

In Chapter 3, we have analyzed the dynamics of the opinions in a collective binary decision-making system composed of a swarm of ground-based robots (Montes de Oca et al., 2011). We have defined an absorbing Markov chain with two absorbing states corresponding, respectively, to complete consensus of the swarm to opinion A and opinion B. We have derived the cumulative distribution function as well as the probability mass function of the number of applications of the majority rule necessary to reach consensus. With respect to previous studies (Massink et al., 2012; Montes de Oca et al., 2011; Scheidler, 2011), we have showed that the majority rule with differential latency

model is characterized by a variance of the time to consensus that is several orders of magnitude greater than the expectation. Hence, we remarked that the expectation of the time to consensus, primary used in previous studies, is not a reliable statistics of this system.

In Chapter 4, we analyzed the performance of a distributed communication protocol for heterogeneous robotic swarms. We defined an absorbing Markov chain of the one-to-one selection algorithm proposed in (Mathews et al., 2010a,b). In this case, the chain is characterized by a single absorbing state that identifies the completion of the selection algorithm. Thanks to the Markov chain model, we have provided a in-depth study of the number of iterations necessary to establish a one-to-one communication link between an aerial robot and a target wheeled robot. With respect to previous theoretical studies, we have showed that the one-to-one selection algorithm scales logarithmically with the number of wheeled robots involved in the process.

As future directions of research, we plan to extend the proposed modeling approach with mechanisms that allow designers to define models based on more complex assumptions on the spatial interactions of the agents in the swarm while maintaining the non-spatial characteristic of the models.

Bibliography

S. Banisch, R. Lima, and T. Araújo. Agent based models and opinion dynamics as markov chains. *ArXiv preprint arXiv:1108.1716v2*, 2011.

E. Bonabeau, M. Dorigo, and G. Theraulaz. *Swarm Intelligence: From Natural to Artificial Systems*. Oxford University Press, New York, 1999.

M. Bonani, V. Longchamp, S. Magnenat, P. Retornaz, D. Burnier, G. Roulet, F. Vaussard, H. Bleuler, and F. Mondada. The marXbot, a miniature mobile robot opening new perspectives for the collective-robotic research. In *Intelligent Robots and Systems (IROS), 2010 IEEE/RSJ International Conference on*, pages 4187–4193. IEEE Press, Piscataway, NY, 2010.

M. Brambilla, E. Ferrante, M. Birattari, and M. Dorigo. Swarm robotics: a review from the swarm engineering perspective. *Swarm Intelligence*, 7(1):1–41, 2013.

P.-J. Bristeau, F. Callou, D. Vissière, and N. Petit. The navigation and control technology inside the ar. drone micro uav. In *Proceedings of the 18th IFAC World Congress*, volume 18, pages 1477–1484. International Federation of Automatic Control, 2011.

S. Galam. Majority rule, hierarchical structures, and democratic totalitarianism: A statistical approach. *Journal of Mathematical Psychology*, 30(4):426 – 434, 1986.

S. Goss, S. Aron, J. L. Deneubourg, and J. M. Pasteels. Self-organized shortcuts in the argentine ant. *Naturwissenschaften*, 76(12):579–581, 1989.

J. G. Kemeny and J. L. Snell. *Finite Markov Chains*. Springer-Verlag New York, 1976.

J. Kennedy and R. C. Eberhart. *Swarm intelligence*. Morgan Kaufmann Publishers Inc., San Francisco, CA, 2001.

P. L. Krapivsky and S. Redner. Dynamics of majority rule in two-state interacting spin systems. *Phys. Rev. Lett.*, 90:238701, Jun 2003.

R. Lambiotte, J. Saramäki, and V. D. Blondel. Dynamics of latent voters. *Phys. Rev. E*, 79:046107, Apr 2009.

A. Martinoli, K. Easton, and W. Agassounon. Modeling swarm robotic systems: a case study in collaborative distributed manipulation. *The International Journal of Robotics Research*, 23(4-5):415–436, 2004.

M. Massink, M. Brambilla, D. Latella, M. Dorigo, and M. Birattari. Analysing robot decision-making with Bio-PEPA. In M. Dorigo, M. Birattari, C. Blum, A. L. Christensen, A. P. Engelbrecht, R. Groß, and T. Stützle, editors, *Swarm Intelligence*, volume 7461 of *Lecture Notes in Computer Science*, pages 25–36. Springer, Berlin, 2012.

N. Mathews, A. L. Christensen, E. Ferrante, R. O'Grady, and M. Dorigo. Establishing spatially targeted communication in a heterogeneous robot swarm. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: volume 1 - Volume 1*, AAMAS '10, pages 939–946, Richland, SC, 2010a. International Foundation for Autonomous Agents and Multiagent Systems.

N. Mathews, A. L. Christensen, R. OGrady, and M. Dorigo. Cooperation in a heterogeneous robot swarm through spatially targeted communication. In M. Dorigo, M. Birattari, G. A. Caro, R. Doursat, A. P. Engelbrecht, D. Floreano, L. M. Gambardella, R. Groß, E. Şahin, H. Sayama, and T. Stützle, editors, *Swarm Intelligence*, volume 6234 of *Lecture Notes in Computer Science*, pages 400–407. Springer, Berlin, 2010b.

M. Montes de Oca, E. Ferrante, A. Scheidler, C. Pinciroli, M. Birattari, and M. Dorigo. Majority-rule opinion dynamics with differential latency: a mechanism for self-organized collective decision-making. *Swarm Intelligence*, 5:305–327, 2011.

J. R. Norris. *Markov Chains*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, New York, NY, first edition, 1997.

M. A. Nowak. Five Rules for the Evolution of Cooperation. *Science*, 314(5805):1560–1563, 2006.

C. Pinciroli, V. Trianni, R. O'Grady, G. Pini, A. Brutschy, M. Brambilla, N. Mathews, E. Ferrante, G. Caro, F. Ducatelle, M. Birattari, L. M. Gambardella, and M. Dorigo. ARGoS: a modular, parallel, multi-engine simulator for multi-robot systems. *Swarm Intelligence*, 6(4):271–295, 2012.

E. Şahin. Swarm robotics: From sources of inspiration to domains of application. In E. Şahin and W. Spears, editors, *Swarm Robotics*, volume 3342 of *Lecture Notes in Computer Science*, pages 10–20. Springer, Berlin, 2005.

A. Scheidler. Dynamics of majority rule with differential latencies. *Phys. Rev. E*, 83:031116, 2011.

G. Valentini, M. Birattari, and M. Dorigo. Majority rule with differential latency: An absorbing markov chain to model consensus. In T. Gilbert, M. Kirkilionis, and G. Nicolis, editors, *Proceedings of the European Conference on Complex Systems 2012*, Springer Proceedings in Complexity. Springer, New York, 2013.