MODELING ARTIFICIAL, MOBILE SWARM SYSTEMS

Thesis by

William Agassounon

In Partial Fulfillment of the Requirements

for the Degree of

Doctor of Philosophy

California Institute of Technology

Pasadena, California

2003

(Defended May 20, 2003)

Acknowledgements

I am grateful to Willy Behrens and the Caltech professors of my candidacy, steering, and defense committees, for securing a financial support for my research, following my progress, and sharing their comments: Yaser Abu-Mostafa, Chris Adami, Richard Murray, and Pietro Perona.

I am grateful to Erik Antonsson for suggesting me the idea of having a Ph.D. steering committee and especially to Robert McEliece for accepting to chair this committee and for his "Let me know if I can help you in any way."

I owe many thanks to Rodney Goodman for sparking my interest in the topic and giving me the opportunity, freedom, and independence to accomplish this work and in particular to Alcherio Martinoli whose guidance and objective opinion made this work possible.

I would like also to thank some former and present members of the micro-systems laboratory and the collective robotics group, and many friends at Caltech who have given me fond memories of my time in Pasadena. I cannot list them all but to name a few: Joseph Chen, Nikolaus Correll, Jeff Dickson, Kjerstin Easton, Cédric Florens, Adam Hayes, Sarah Heilshorn, Possu Huang, Cindy Hunt, Sanza Kasadi, Amit Kenjale, Vincent Koosh, Olga Lassak, Marissa Mock, Anne-Laure Le Ny, Daniel Lieberman, Shervin Taghavi, Philip Tsao, Paul Wiggins, and Yizhen Zhang.

Finally, I am grateful to my parents, my siblings, and the Lemaire and Vigneron families for providing unconditional love and support, for fostering my imagination, nurturing my intellectual interests as long as I can remember, and for encouraging me to always follow my intuition.

My work was supported by grants from the TRW Foundation and TRW Space and Electronics. Further funding was provided by the Caltech Engineering Research Centers program of the National Science Foundation under award number EEC-9402726. The experimental data reported in Chapter 4 are courtesy of Kjerstin Easton.

Abstract

Swarm intelligence is a new research paradigm that offers novel approaches for studying and solving distributed problems using solutions inspired by social insects and other natural behaviors of vertebrates. In this thesis, we present methodologies for modeling artificial mobile systems within the swarm intelligence framework. The proposed methodologies provide guidelines in the study and design of artificial swarm systems for the following two classes of experiments: distributed sensing and distributed manipulation.

Event discovery and information dissemination through local communication in artificial swarm systems present similar characteristics as natural phenomena such as foraging and food discovery in insect colonies and the spread of infectious diseases in animal populations, respectively. We show that the artificial systems can be described in similar mathematical terms as those used to describe the natural systems. The proposed models can be classified in two main categories: non-embodied and embodied models. In the first category agents are modeled as mobile bodiless points, whereas the other models take into account the physical interference between agents resulting from embodiment. Furthermore, within each category, we distinguish two subcategories: spatial and nonspatial models. In the spatial models we keep track of the trajectory of each agent, the correlation between the positions occupied by the agents over consecutive time steps, or make use of the spatial distribution resulting from the movement pattern of the agents. In the nonspatial models we assume that agents hop around randomly and occupy independent positions over consecutive time steps.

In our description of distributed manipulation in swarm robotic systems we present two case studies of non-collaborative and collaborative manipulations, respectively. The general approach proposed here consists of first representing the group behavior of the active agents with a Finite State Machine (FSM) then describing mathematically the dynamics of the group. The first case study is the aggregation experiment that consists of collecting and gathering objects scattered around an enclosed arena. We present a macroscopic model that accurately captures the dynamics of the experiment and a suite of threshold-based, scalable, and fully distributed algorithms for allocating the workers to the task optimally. The second case study is that of the stick-pulling experiment in which a group of robots is used to pull sticks from the ground. This task requires the collaborative effort of two robots to be successful. Here, we present a discrete-time macroscopic model that helps

us uncover counter-intuitive behaviors that result from collaboration between the agents.

We complete each proposed modeling methodology by showing how the parameters of the models can be calculated using solely the characteristics of the environment and those of the agents and by analyzing the constraints and limitations of the different models. Finally, we use different tools (simulations and real robots) to validate the proposed models.

Contents

List of Figures

List of Tables

List of Symbols

Modeling Information Flow

D: diffusion coefficient.

D_m: displacement after m consecutive moves.

N: total number of mobile agents/sensors.

$P(R)$: probability of successful transmission.

P_w: probability to encounter the surrounding wall.

P_e: probability to encounter the event.

P_s: probability to encounter a susceptible agent.

P_i: probability to encounter an informed agent.

Φ_N: probability that at least one of N agents finds the event at a particular time step (event with random disappearance).

R: communication range of a mobile agent.

T_0: constant duration (event with constant duration).

T_{const}: Expected residency time of event with constant duration and disappearance after been discovered.

T_{rand}: Expected residency time of event with random disappearance and disappearance after been discovered.

W_k: two-dimensional vector representing the kth move in the correlated random walk model.

$a(\theta)$: overlapped area between two consecutive moves.

α: agent-event encounter rate.

d: constant length of a correlated random walk.

d_0: radius of the arena.

$f(\cdot)$: distribution function of the agents.

$g(\theta)$: probability distribution of θ.

γ: fraction of total area effectively explored per agent per time step.

γ_0: probability to discover an event at a particular time step.

$h(l)$: probability distribution of l.

i: fraction of informed agents.

l: path length of the correlated random walk model.

λ: agent-agent encounter rate.

p: probability of spontaneous disappearance (event with random disappearance).

r: distance to origin.

ρ_N: probability that at least one of N agents finds the event at a particular time step (event with constant duration).

$\rho(r)$: density of agents at location r.

τ_e: duration of the event detection maneuver.

τ_i: duration of the information reception from an informed agent.

τ_{obs}: duration of the agent-agent collision avoidance maneuver (embodied models).

τ_s: duration of the susceptible agent avoidance maneuver.

τ_w: duration of the wall avoidance maneuver.

θ: direction of the correlated random walk model.

s: fraction of susceptible agents.

(ξ,η): spatial coordinates.

w: detection width of an agent.

Modeling Non-Collaborative Distributed Manipulation

L^*: average cluster size at steady-state.

M: total number of seeds.

Γ_f: fraction of disappointed unloaded robots.

Γ_c: fraction of disappointed loaded robots.

W_0: total number of robots.

γ_w: rate of agent-wall encounter.

γ_r: rate of agent-agent encounter.

γ_k^{dec}: rate of encounter with a cluster of size k from its seed-removal angle.

γ_k^{inc}: rate of encounter with a cluster of size k from its seed-release angle.

n_k: number of clusters of size k.

τ_{search}: timeout in the search for work.

w_f: number of unloaded robots.

w_c: number of loaded robots.

w_g: number of robots gripping a seed.

w_d: number of robots releasing a seed.

w_q: number of robots quitting the task.

w_i: number of robots idle after quitting.

Modeling Collaborative Distributed Manipulation

$C(k)$: number of successful collaborations at iteration k.

Γ: fraction of disappointed robots.

M_0: total number of sticks.

R_g: collaboration parameter.

T_g: gripping time parameter (in number of iterations).

W_a: number of robots in collision avoidance state.

W_d: number of robots in dance state.

W_g: number of robots in gripping state.

W_i: number of robots in interference state.

W_s: number of robots in search state.

p_{g1}: probability to find an isolated stick.

p_{g2}: probability to find a stick half-pulled out by another robot.

τ_g: gripping time parameter (in time units).

List of Acronyms

CDMA: code division multiple access.

CR: communication range.

DDE: differential difference equation.

DE: difference equation.

FSM: finite state machine.

IC: integrated cost.

n-D: n-dimensional (for different positive integer values of n).

PFSM: probabilistic finite state machine.

PrFT: private, fixed threshold.

PrVT: private, variable threshold.

PuFT: public, fixed threshold.

PuVT: public, variable threshold.

SI: swarm intelligence.

Chapter 1

Introduction

In this introductory chapter, we give a problem formulation and name previous and related works. The chapter is organized as follows: In section 1.1 we give a historic overview of the field of Swarm Intelligence. Section 1.2 discusses the special significance and importance of modeling in the study of swarm-based systems, which is the motivation behind this research. Related work is presented in section 1.3. The contributions of this thesis are summarized in section 1.4. This thesis is written such that each chapter is self-contained and can be read independently from the others.

1.1 Swarm Intelligence: An Historic Perspective

Albert Einstein once said: "Things should be made as simple as possible, but not any simpler." Nature seems to have been using this principle for millions of years. Haven't we all been amazed by the individually simple but collectively complex behavior exhibited by social insects such as termites, bees, and wasps? For two decades, scientists have been trying to understand the underlying principles of the robust, scalable, and distributed self-organization observed in some social animal societies. The attempt to apply the insight gained through this research to artificial systems (e.g., massively distributed computer systems and robotics) has given rise to a new research topic called *Swarm Intelligence* (SI) [12, 16].

Thus, swarm intelligence is an innovative computational and behavioral metaphor for solving distributed problems that takes its inspiration from the biological examples provided by social insects [16] such as ants, termites, bees, and wasps and by swarming, flocking, herding, and shoaling phenomena in vertebrates [85]. The abilities of such natural systems appear to transcend the abilities of the constituent individual agents. In most biological cases studied so far, the robust and capable high-level group behavior (e.g., division of labor [45, 102] and sophisticated nest architecture construction [97]) has been found to be mediated by nothing more than a small set of simple low-level interactions between individuals, and between individuals and the environment. The SI approach emphasizes distributedness, parallelism, exploitation of direct (peer-to-peer) or indirect (via the en-

vironment) local communication mechanisms among relatively simple agents, and self-organization.

1.2 From Natural to Artificial Systems: The Importance of Modeling

Related to having a single powerful agent, designing, building, and using several minimalist agents to accomplish a given task may be easier because they use, for example, simpler sensor and actuator apparatuses, cheaper because they are simple, more flexible without the need to reprogram the agents, and more reliable and fault-tolerant because one or several agents may fail without affecting task completion (although completion time may be affected by such a perturbation). However, understanding natural systems and designing useful artificial swarm systems are quite different endeavors. On the one hand, understanding biological systems requires observing, carrying out experiments, and making models (of the natural systems) whose predictions are consistent with observation and/or experiments. Designing an artificial swarm-based system, on the other hand, is better done when there are models that not only accurately describe the (artificial) system but also take into account all its parameters of importance and allow to study their (individual) influence on the behavior and performance of the active constituents of the system.

Swarm systems are stochastic and this results in a high variability of their behavior and performance [72, 55]. Moreover, most contributions in swarm systems, especially in swarm robotics, are in the forms of statistical analysis of experimental results or simulations [11, 62, 61, 43]. While experiments allow researchers to study the system's behavior in real environments, they are costly and time-consuming, both to set up and execute, this prevents experimenters from studying swarm-like systems and constrains them to rely on simple extrapolation of the results obtained with small size systems. In addition, although sensor-based simulations, such as Webots [74] and Player/Stage [31], recreate the experiment under realistic conditions and provide statistically significant results much more rapidly than experiments [68, 41], they usually do not scale either (as well as the size of the system grows), and are, therefore, impractical tools for a detailed investigation of the properties of large multi-agent systems.

Finally, models of swarm systems can help to identify the parameters that influence the individual and/or the team performance, to answer some fundamental questions (e.g., for a given task, is there an optimal number of agents that will complete the task in the shortest period of time?), to find and understand the boundaries of the systems (e.g., what is the maximum number of agents for any given environmental setup?), to predict the long-term behavior of the systems, and to gain insight into system design (e.g., what parameters determine group behavior, optimize performance, prevent instabilities). However, with the exception of a limited number of research works [4, 68, 57, 93], few models have been proposed to study swarm systems in general, and mobile swarm robotic systems

in particular.

1.3 State of the Art in Modeling Artificial Swarm Systems

Few accurate and computationally efficient tools exist for studying artificial swarm systems. The existing models can be classified in two categories: microscopic and macroscopic models. While microscopic models describe the individual agent's interactions with its teammates and the environment (see for instance [63, 64]), macroscopic models offer a direct (mostly mathematical) description of the collective group behavior. In general, the macroscopic models are more computationally efficient than their microscopic counterparts although the latter often offer more detailed information about the dynamics of the same systems to which they are applied. In the microscopic model proposed by Martinoli and coworkers [63], several simulation runs are required to obtain a statistically significant (mean) behavior of the dynamic variables of the system.

Altogether, these are either discrete-time or continuous-time models. Intuitively, since the physical setup operates in continuous real time, we might think that the (only) good way to obtain a faithful model would be with continuous-time models. However, if experiment is in continuous time, simulations require a time granularity often determined by a fixed iteration/time step. Moreover, each agent of an artificial swarm system (e.g., robots, sensors) is itself a hybrid system (in a the automatic control sense), in other words, it is made of diverse elements: continuous time elements (e.g., analog electronics parts, motors, sensors) but also discrete time elements (e.g., digital electronics parts like micro-controllers) or the agent is characterized by logical control mechanisms (the success dance in the stick-pulling experiment presented in Chapter 4 is a typical example: its duration is triggered by a sensory stimulus and lasts until an incremental timer of the robot reaches a fixed timeout). Finally, analytical (differential) equations of the continuous-time models describing such complex systems, often, do not have closed form solutions, thus, they are usually solved using numerical methods which require a fixed time granularity (e.g., Euler's method or Runge-Kutta's technique). This makes the case for discrete-time and continuous-time models to coexist under the same framework for describing the behavior of swarm systems and the dynamics of the variables of the systems.

1.3.1 Models of event discovery and information spread in natural swarm systems

Over the past several decades, theoreticians have described population dynamics in general and the spread of infectious diseases in particular, with simple macroscopic models [52, 77]. These are all systems in which distributedness and local interaction play a major role. Others, in biology, have proposed different mathematical models of flight, crawling, scouting, and foraging in insect colonies

(see, for instance, [49, 16]) that pave the way for describing artificial systems inspired by Nature. For instance, in [1], Adler and Gordon present a model of a network of scouting ants in information collection and spread. This was a pioneering work as it introduced a new approach in studying the outcome of local interaction in ant colonies different than the *spatial logistic equation* often used to study natural population dynamics (see section 2.5.1 for details) and provided a tool to build upon (and avoid starting from scratch) for describing artificial swarm sensor systems.

Sensor networks are being applied in areas including environmental monitoring, conditional base maintenance, surveillance, smart spaces, inventory tracking, etc., and many more applications will appear as the technology becomes available. Cooperative or distributed swarm sensing tasks require that the team of agents maintain an explicit or implicit flow of information about the environment, the progress in task accomplishment, etc. In large systems consisting of multiple autonomous agents, information exchange through a centralized communication medium creates a bottleneck which in turn renders the system fragile. Moreover, due to interference, e.g., presence of unnecessary or obsolete information and packets collision, and the limited communication capabilities of the agents (often due to power limitation), a local, peer-to-peer, concurrent communication system is more appropriate than a global communication medium as for instance, the former is flexible and reduces excessive information processing. The effectiveness of a network of mobile sensors in monitoring depends on how exhaustively it covers the territory, how sensitive it is to changes in the environment, and how well it disseminates information about such changes. Thus we are particularly concerned with modeling information collection and spread by information-gathering agents that do not search in isolation but encounter each other and share information.

In a more general scope, recent studies on simple mobile robots have shown that sophisticated and complicated tasks such as exploration [91, 69, 56, 13], map generation [89, 25] and transportation [24, 39, 92, 75] can be accomplished through cooperation. Most of these tasks require the existence of an information flow between the agents [9, 18] which can be either explicit, i.e., agent-to-agent information transmission or implicit, i.e., communication through the modification of the environment. The latter is not suitable for fast, dynamically changing environments as it relies on observation of other agents and the changes brought into the environment [11]. The explicit communication is appropriate for sensor networks and breaks in two categories: global and local inter-agent communication.

A global communication mechanism [105, 7, 80, 83] is often used in systems consisting of few agents where information exchanged is checked all the time, anywhere, assuming a full connectivity of the network. On the one hand, a common global communication channel can cause a decrease in transmission efficiency when the network size increases (failure to scale due to a limited bandwidth), a communication bottleneck and/or an insufficient fault tolerance when facing load increases. On the other hand, local communication reduces excessive information processing by allowing a distributed

and concurrent information transmission, and makes the whole system more robust to individual agent failure. Therefore the latter is used more frequently in multi sensing agent systems [47, 98, 3].

Moreover, recently, several studies on applying communication to distributed mobile robot systems have been proposed, some based solely on experimental results [46, 47, 99] and/or lacking a theoretical framework [71], others in the field of communication theory have not taken into account the embodiment of the agents and the transmission to a limited number of agents [95]. In the pioneering work of Yoshida *et al.* [104], for the first time the area of communication was designed using mathematical analysis and the authors also proposed an analysis of the diffusion time as how long to let the information diffuse so it can be transmitted to a desired number of agents. Nevertheless, none of these approaches has ever explicitly taken into account the movement pattern of the agents or the spatial distribution that can result from specific movement patterns.

Since the beginning of the 20th century, a lot of studies have been conducted on modeling and explaining the dynamics of biological phenomena of geographical spread. In particular, the characteristics of ecological and epidemics spread have been surveyed by chemists, biologists, and statisticians [60, 28, 53, 90, 52, 76, 77, 79]. For instance, population dynamics [28], infectious diseases spread in human populations [52, 76, 77], foraging, scouting, recruitment, and information exchange in ant colonies [17, 86, 35, 10, 1], etc., are all phenomena which characteristics are similar to those of event discovery and information dissemination in networks of mobile sensors, as they are all concerned with mobile autonomous agents that interact locally with each other. These studies inspire the theoretical work on swarm-based networks of sensors presented in this thesis.

The sensor systems we are concerned with present two main characteristics: local information exchange and mobility of the constituent agents. This is motivated by recent theoretical work on sensor networks showing that local communication reduces power consumption and communication interference and mobility improves the overall throughput[1] of such networks. Gupta and Kumar [37] proposed a model for studying the capacity of fixed ad-hoc networks, where nodes are randomly located but are immobile. In their system, each source node has a random destination to which to communicate. Their main result shows that as the number of nodes per unit area n increases, the throughput per source-destination pair decreases approximately like $1/\sqrt{n}$. —[2]. Their main insights are that communication must be limited to near neighbors in order to permit dense spatial channel reuse. Moreover, the importance of mobility in the sensor network is illustrated by Grossglauser and Tse [36]. The authors, examining the asymptotic throughput capacity of wireless ad-hoc sensor networks, have shown that the throughput of such networks is tremendously improved when the nodes are mobile.

Other advantages of mobility (that are not considered in this work) are self-deployment and self-

[1] The throughput of an ad-hoc sensor network is defined as the average number of successful transmissions per source-destination pair per time unit. Sometimes, this number is calculated per node.

[2] This was shown to be the optimal asymptotic behavior regardless of the routing and communication protocols.

configuration (e.g., formations) in response to the nature and the spatial distribution of the events to be monitored (e.g., dense, sparse, random, etc.)

1.3.2 Models of swarm robotic systems

As of today, few models of swarm robotic systems exist. Among those developed so far, some focus on task allocation mechanisms [2], others offer a reliable but mostly qualitatively accurate theoretical paradigm [57]. In [58], Lerman and Galstyan proposed a model of foraging robots to study the quantitative effects of (physical) interference on the ability of the team to find and collect objects.

However, recently, a few macroscopic models have been proposed to explain task allocation mechanisms in natural colonies (that could be applied to robotic systems): some of them based on threshold responses [15, 96], others focus only on task-switching probabilities [82]. However, none of these works has attempted to compare the macroscopic description of the system (which matches the experimental data) with a microscopic characterization, for instance, by investigating how workers gather the information necessary to decide whether or not to switch task or to engage in a task performance. More specifically, they have not taken into consideration the partial perception in time and space of the stimulus (e.g., demand for food) and the embodiment of the agents. For instance, partial perceptions of the global stimulus combined with real world uncertainties could strongly influence the optimal distribution of thresholds among teammates or the task-switching mechanism itself (i.e., probabilistic vs. deterministic). Thus, although these models of natural systems can be seen as good candidates for providing a macroscopic description of artificial systems, they fall short when it comes to explain how the global behavior of the system (i.e., what happens at the group level) derives from what happens locally at the individuals' level (so that the system can be designed). This latter characteristic is required from models of artificial systems, because, as opposed to models of natural systems, they are designed not just to accurately describe the systems but also to provide ways to control them and eventually optimize their performance.

1.3.3 A simulation tool: the nonspatial microscopic model

Among the simulation tools used throughout this thesis to validate the proposed macroscopic models there is a probabilistic microscopic model. This microscopic model was first proposed in [48, 63, 64] by Martinoli and coworkers for studying robotic experiments (e.g., stick-pulling and cluster formation).

The central idea of the microscopic probabilistic model is to describe the experiment as a series of stochastic events with probabilities based on simple geometrical considerations and systematic interaction experiments with a single embodied or real agent. The probability for any particular agent to encounter any particular object present in the environment is given by the fraction of the total extended area of the object related to the total area of the arena in which the agent is moving.

The extended area occupied by each object (resp. agent) is computed by considering the detection range of the object (resp. agent) by any active, mobile agent. In the case of a sensing task (i.e., event discovery), the extended area of a given event is calculated in a similar manner, i.e., by the area defined by the detection range of the event by any active agent roaming in the environment.

In this probabilistic model a finite state machine represents all the possible states of the agents, but instead of computing the detailed sensory information and trajectories of the agents (this is a nonspatial model) state transitions are determined by rolling a dice (using the probabilities above). To obtain statistically significant results, several simulation runs of the experiment are carried out and the overall behavior of the system under consideration is computed by averaging the results of those runs.

The microscopic model has been shown to be accurate and computationally efficient [48, 63, 64, 2, 3]. For instance, for the cluster formation experiment with 10 agents, the speed between the microscopic model and Webots is about 70,000 on a PC Pentium III 800 MHz when considering the amount of time it takes to carry out a single simulation run [3]. Finally, the microscopic model is easily applicable to different (sensing and robotic manipulation) experiments and can be used for both embodied and non-embodied simulations. For instance, for non-embodied (point) simulations one just has to set the radii and detection ranges of the different objects present in the arena to 0 and neglect the time spent in obstacle/collision avoidance.

1.3.4 Webots: a sensor-based simulator

Most of the experiments presented in this thesis were implemented in *Webots* [74], a 3-D, kinematic, sensor-based simulator of *Khepera* robots (see Figure 2.18). This simulator computes trajectories and sensory inputs of the mobile robots in an arena corresponding to a given physical setup. The resulting simulation is sufficiently faithful for the controllers to be transferred (i.e., downloaded) to real robots without changes and for the behaviors of the simulated robots to be very similar to those of the real robots, as shown in several previous publications [41, 48, 63, 64]. Webots represents a reliable and computationally efficient tool (as compared with experiments using real robots) for the study of mobile swarm-based systems.

1.4 Contributions

In this thesis, we propose a series of models for describing and studying swarm systems inspired by complex biological colonies. This work covers two specific areas within the (artificial) swarm paradigm: swarm-based mobile sensor networks and swarm robotics.

1.4.1 Modeling information flow in swarm-based mobile sensor networks

As of the time of this thesis, to the best of our knowledge, there are no models that specifically describe swarm-based networks of mobile, artificial sensors. In the first chapter of this thesis, we applied the models used to study population dynamics and insect colonies to artificial swarm sensor systems and show how the systems can be fully characterized by calculating the parameters of the models using solely the characteristics of the environment (e.g., arena radius) and those of the agents (e.g., speed, body size, communication and sensor ranges, movement pattern).

As opposed to usual networks of packet radio terminals, this work elaborates on highly scalable *ad hoc* networks of minimalist mobile sensors acting locally on a single sensing task (e.g., localized odor sensing, mapping at a reduced scale) which may be a subtask in a larger system. Therefore, we do not account for transmission protocols or routing strategies. Our goal is to devise models and techniques for analyzing and designing local communication systems for distributed, mobile sensor systems. For instance, in a robotic system, this type of local communication can be in the form of low range radio, infrared, laser, or sonic signals transmission. Biology offers a variety of local communication mechanisms in social animals colonies, e.g, ants sometimes communicate to each other by antenna contact. With this goal, we look at natural phenomena for inspiration.

1.4.2 Modeling distributed manipulation in swarm robotic systems

The main difference between sensing tasks and manipulation tasks is that the environment in which the manipulation tasks are accomplished, in addition to being passive, is also a shared resource whose modifications are generated by the parallel actions of the robots. For instance, groups of robots that clean up radioactive waste in a nuclear plant would modify the physical configuration of the environment while those that scout (just in search of defects in the construction structure of the nuclear plant) would not modify the environment. The former have a manipulation task while the latter accomplish a detection/sensing task.

The main advantages of the application of the SI approach to the control of multiple robots are fourfold: (i) *scalability*: the control architecture is kept exactly the same from a few units to thousands of units; (ii) *flexibility*: units can be dynamically added or removed, they can be given the ability to reallocate and redistribute themselves in a self-organized way; (iii) *robustness*: the resulting collective system is robust in facing a priori unknown environmental and team changes not only through unit redundancy but also through an adequate balance between exploratory and exploitative behavior; and (iv) *individual simplicity*: an effective way to naturally obtain an adequate exploratory-exploitative balance and at the same time allow for unit miniaturization and overall system cost reduction is to minimize individual complexity.

The central idea of our modeling methodologies is to describe each experiment as a series of

(stochastic) events, then to summarize the task accomplishment using a Finite State Machine (FSM) whose transition probabilities/rates are calculated using the geometrical properties of the environment and systematic experiments results (with one or two embodied agents). In a macroscopic model (like those introduced in Chapters 3 and 4), a single FSM summarizes the whole robotic team, each of its states representing the (average) number of teammates in a particular state during a certain time step. The FSM is coupled with the environment as its parameters change when the configuration of the environment changes due to the actions of the robots.

In Chapters 3 and 4 we present two case studies of non-collaborative and collaborative distributed manipulation, both belonging to the category of experiments in which everything is based on local interaction and local communication strictly following the SI principles as proposed by several other researchers [11, 44, 55]. In this thesis we refer to a *collaborative* task as one that requires a strictly cooperative effort of several agents to be successfully accomplished, and to a *non-collaborative* task as one that can be accomplished without a cooperation between the workers. Note that although the use of multiple agents is not strictly required for the tasks under consideration but we used several agents to speed up the accomplishment of the task and increase the rate of interaction between the robots. In the following, each case study is described in turn.

A. A case study of non-collaborative task: the aggregation experiment

The experiment presented in this case study is the follow-up of experiments performed with real robots presented by Martinoli *et al.* [62, 63, 64]. The task is to collect small objects, referred to as *seeds*, in a square arena and to gather them in a single cluster[3] using Khepera robots [78] equipped with grippers and capable of distinguishing the seeds with their frontal infrared sensors. The robots have local sensing capabilities only such that their encounters with seeds and clusters of seeds happen randomly. Moreover, as the robots do not have a global perception of the environment they do not know when the task is finished, motivating the need for a distributed task allocation mechanism to allow each individual to stop working based on its own estimate of the availability of work.

Division of Labor

We propose a swarm-based, scalable, and fully distributed mechanism that allows an efficient worker allocation for obtaining improved clustering results without the need for a centralized controller. More generally, in the corresponding section of the thesis, we investigate a suite of task allocation and labor division mechanisms exploited in social insect societies that are suitable for artificial embedded systems such as multiple mobile robot platforms. The proposed worker allocation algorithm is based exclusively on the local estimation of the demand by the individuals. The individuals are all characterized by the same threshold but since the agents do not perceive the demand globally but rather estimate it locally, they do not work or rest all at the same time, a behavior that would

[3] Although the experiment is not intended to reproduce a biological system, the experiment presents several similarities with the nest cleaning and dead ants clustering performed by some ant colonies [20].

arise if the demand was broadcast from an external supervisor. We hope that this work not only will contribute to the autonomous robotic literature but will also stimulate the discussion among biologists interested in principles of division of labor in social insect societies.

Modeling

We propose a mathematical model for the aggregation experiment. This is a continuous-time model in which we use a set of differential difference equations (DDEs) to describe the dynamics of the variables of the aggregation experiment and predict their quantitative outcome. This model also accurately captures the behavior of the system when the agents are endowed with the distributed task allocation mechanism. Finally, the proposed methodology has the additional advantage of being easily applicable to other distributed manipulation experiments such as puck clustering [11] and object sorting [44, 72].

B. A case study of collaborative task: the stick-pulling experiment

The probabilistic modeling methodology presented here describes the experiment as a series of stochastic events with the transition between different states of the FSM representing the system happening probabilistically. The resulting model is a discrete-time model in which we use a set of difference equations (DEs) to capture the dynamics of the different constituents of the system and accurately predict some counter-intuitive behaviors resulting from the cooperation between the robots.

C. Modeling assumptions for swarm robotic systems

The modeling methodologies proposed for the two case studies of robotic experiments rely on two main assumptions: the spatial uniformity and the Markov property, each briefly described in turn.

Spatial Uniformity.— The proposed models are nonspatial (following the definition given in this thesis) and the methodologies rely on the assumption that the group of robots covers the arena surface uniformly. The trajectories of the robots therefore are not considered in the model. We also assume that the absolute position of a given object to manipulate in the arena does not play a role, only the quantitative characterization of the group to which it belongs does. For instance, in the aggregation experiment, a seed always belongs to a cluster of a certain size, however, the probability of that seed being manipulated is the same whether the cluster is placed in the center or in the periphery of the arena. Similarly, in the stick-pulling experiment, a stick can be pulled out of ground no matter its position on the arena.

Markov Property.— We assume that each robot's future state depends only on its present state and on how much time it has spent in that state. This assumption is correct for a reactive robot controller comprising a series of actions (or sequence of actions) with fixed durations (e.g., grasping or dropping a seed, holding a stick). The robots and the environment in the aggregation case study

clearly obey this Markov property if we assume that trajectories, i.e., position and heading, can be neglected for computing the desired metric.

Both assumptions are valid in most of the experiments presented. Nevertheless, in Chapter 4 we discuss in detail the limitations of such mathematical models, in particular, the validity of the assumptions of spatial uniformity and Markov property in the case of overcrowded arenas, when the macroscopic model no longer delivers quantitatively correct predictions.

Chapter 2

Modeling Information Flow in Swarm-Based Mobile Sensor Networks

2.1 Introduction

In this chapter, we are interested in modeling event discovery (i.e., sensing) and information exchange and dissemination through local, explicit communication in swarm-based networks of mobile sensors. The proposed models take into account different characteristics of the agents, e.g., embodiment/non-embodiment, sensor and communication ranges, movement pattern, etc. We study the efficiency with which agents in such networks discover a changing stream of events scattered throughout a region and the dissemination of information in the networks. We present microscopic and macroscopic models of how such networks can monitor events occurring throughout a given territory, and maintain an ongoing information flow about the environment through information spread. We present two categories of model: *non-embodied* models (i.e., models in which agents are considered as points and their physical characteristics are ignored) and *embodied* models (i.e., models that take into account the embodiment of the agents and the resulting physical interference between the teammates). Within each category we distinguish *spatial* models (i.e., models that keep track of the agents' trajectories or make use of the spatial distribution resulting from their movement pattern) and *nonspatial* models (i.e., models in which we assume that agents occupy independent, random positions at consecutive time steps). We show that models can not only deliver good quantitative and qualitative predictions, but also serve as a useful optimization tool.

2.2 Problem Description

To study the collectively coordinated behavior that emerges from local interactions between the agents of a swarm-based sensor network requires simple and computationally efficient tools that

avoid the potential burden of the size of the network and deliver accurate predictions.

In this chapter, we refer to *event* as a single information source, to *embodied* agents as those whose physical embodiment is taken into account in the models as opposed to *non-embodied* agents represented as bodiless points. Although self-locomotion is a specific form of mobility, in this chapter, the mobility (of the agents) is achieved solely through self-locomotion, therefore, in the rest of the chapter, we do not make any distinction between *mobile* and *self-locomoted* agents.

The systems under investigation here are fully *ad hoc* networks that consist of a bounded region (or enclosed arena) in which a single or multiple events appear. Each event may present the following characteristics:

Constant duration.— This corresponds to events that last a fixed amount of time before disappearing from the environment. Examples of such events are a flower in bloom, a grenade launch, or an activated bomb with a fixed timeout.

Probability of spontaneous disappearance.— Events with this characteristic can disappear at any time step with a certain probability. An example of such events is the monitoring (by satellite or unmanned airplane) of hit-and-run operations in a guerrilla warfare that happen at undetermined moments and last undetermined amounts of time upon each occurrence.

Disappearance after discovery.— In addition to their original latency, some events sometimes present the characteristic of disappearing after been discovered. An example of such events with a fixed duration is a ticking bomb with a fixed timeout that can be discovered (for example by any agent of a swarm of bomb sniffing sensors) and deactivated (before it explodes) to remove the threat that it represents, thus the activated bomb would disappear as an event before the end of its fixed duration. Another example in the randomly disappearing events category is that of the guerrilla operatives that can be neutralized when discovered.

Multiple mobile sensing agents are used for monitoring, exploring, and discovering the events. The sensors are minimalist agents endowed with basic sensing, self-locomotion, and communication capabilities. The general characteristics of the agents are the following.

Individual simplicity.— Agents in the swarm-based systems under investigation here are not equipped with a camera nor are they provided with any information about their position. Each agent only possesses the sensory and actuator capabilities to perform basic behavioral actions: sensing, collision avoidance, and (wireless) local communication.

Local information exchange capability.— Upon encounter, information exchange between two agents can take place locally through a direct wireless communication mechanism (see Figure 2.1 for example).

Mobility.— The agents are self-locomoted, i.e., capable of moving on their own and capable of moving around the whole event monitoring/exploration region.

The main motivation for using mobile agents in the sensing tasks that we are concerned with here

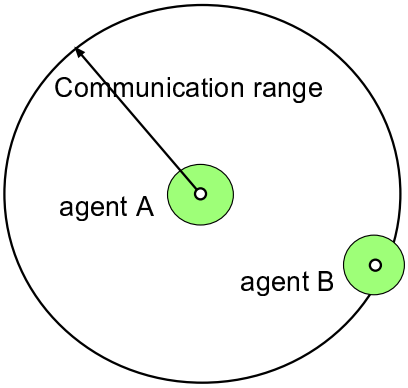

Figure 2.1: Two mobile agents within communication range. Peer-to-peer information exchange happens through local wireless communication.

is that mobility allows the team of agents (with limited sensory capabilities) to explore and cover a much wider area than a network of fixed sensors (with the same sensory capabilities) would. These characteristics of the agents are also motivated by the two pioneering theoretical works of Gupta and Kumar [37] and Grossglauser and Tse [36] on *ad hoc* sensor networks introduced in chapter 1. However, it is worth noting that these works were focusing on the communication aspect of *ad hoc* sensor networks whereas in this chapter we investigate the sensing (i.e., information-mining) aspect in similar networks. Thus, this work addresses another system issue (i.e., the information discovery aspect) of mobile sensor networks, this represents one of its originalities.

2.3 Methodology

We use different approaches to describe the agents and their interactions with each other. First, we do not account for the physical characteristics of the agents, i.e., the embodiment of the agents (this is equivalent to considering the agents as points in a two dimensional space) and, in a second step, we consider models of embodied agents.

The agents may cover the environment with different movement patterns that we describe differently as we distinguish spatial and nonspatial models. Nonspatial models do not take into account the trajectories of the individuals, instead, agents are assumed to occupy independent positions (randomly assigned) during consecutive time steps. This is equivalent to assuming that the agents hop around randomly with equal probability of occupying any position during any time step. With the spatial models, for the embodied models, we make use of the spatial distribution resulting from the movement pattern of the agents and, for the non-embodied models, we consider the movement pattern through a general correlated random walk model defined by two parameters: a step length and a turning angle, each described with a probability distribution function. For instance, movement patterns with sharp turns are easily described with large values of the standard deviation of

the turning angle.

We analyze event discovery by the agents in both the cases where the agents do not share any sensory information and that where they do so upon encounter.

Moreover, as in the real world, agents can only receive communication signals from a finite number of other agents during any time step, we define the *information reception capacity* of the agents, denoted C, as the maximum number of agents that any given agent can receive information from during any given time step. For instance, for an information reception capacity $C = 1$, when during any time step two or more agents transmit to a same receiver, all information messages are simply lost upon reception. The information reception capacity is a simplified model of the hardware and computational limitations of the receiver and its ability to deal with signal interference. For instance, directional antennas may allow signal transmission in/reception from several directions, thus, information exchange with several different agents at a time and a CDMA-like mechanism may also allow for successful signal reception in the presence of a high level of interference.

Figure 2.2 summarizes the areas covered and the different model descriptions presented in this chapter. For every area on the graph that is not specifically covered in this chapter, we explain, in the corresponding section, how the proposed models can be modified or extended to do so. In the

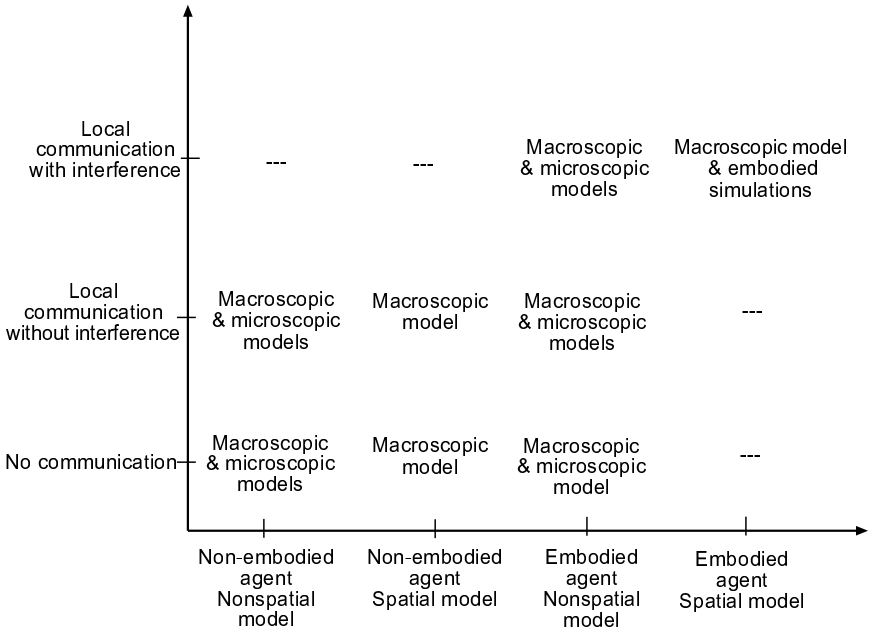

Figure 2.2: Areas covered and models presented in this chapter. (y-axis) Complexity of agent-to-agent interaction vs. (x-axis) complexity of agent and model.

case of event discovery with events that disappear after being discovered, we estimate the expected residency time of the events. In the case of information spread, the metric used for efficiency comparison is the *fraction of informed agents*, which, as a function of time, translates the rate at

which information about a particular event is spread to the agents in the network. Moreover, in the case of optimal design of local communication, the metrics used are the *probability of successful information reception*, used to predict the optimal communication range in the models presented in subsection 2.6.2, and the *number of messages successfully received per agent per time unit*. The last metric is an experimental metric (i.e., obtained from simulation results) that is used to show that the control parameter (i.e., the communication range) that maximizes the probability of successful information reception also maximizes the number of messages that agents successfully received over time. In this chapter, d_0 and N represent the radius of and the total number of agents in the environment, respectively.

Finally, we present a method for calculating the parameters of the different proposed models, using solely the geometrical characteristics of the system, e.g., radius of the arena, the sensory capabilities of the agents, and the event detection range.

The general context of this work is the integration of simple individual behavior in more complex group-level behavior through local information exchange between the constituent agents of a swarm of mobile sensors. This work is inspired by studies in biology and ecology concerning population dynamics (diseases and genes spread in human populations) and insect societies (foraging, scouting, recruitment, and information exchange). In this work we present a set of mathematical tools for describing and analyzing information gathering and information dissemination in large size mobile sensor networks.

2.4 Nonspatial Microscopic Model of the Swarmed-Based Sensor Networks

We applied the microscopic model introduced in Chapter 1 (see subsection 1.3.3) to the distributed sensing and information spread experiment. Figure 2.3 illustrates how a 2-D system consisting of agents and events and characterized by the virtual extended areas occupied by the agents and events is mapped onto a linear (probabilistic) scale. Note that here, the boundary can be taken into account too as agents bounce back into the arena when they encounter the boundary of the exploration/sensing zone.

Following this modeling approach, for instance, the probability P_i that any given agent encounters a round object of type i (of radius R_i) that it can detect from a detection range D_i is given by the following equation

$$P_i = \frac{(R_i + D_i)^2}{d_0^2} \tag{2.1}$$

where d_0 is the radius of the arena assumed to be a circle, without loss of generality.

Figure 2.4 shows the Probabilistic Finite State Machine representation used in the microscopic

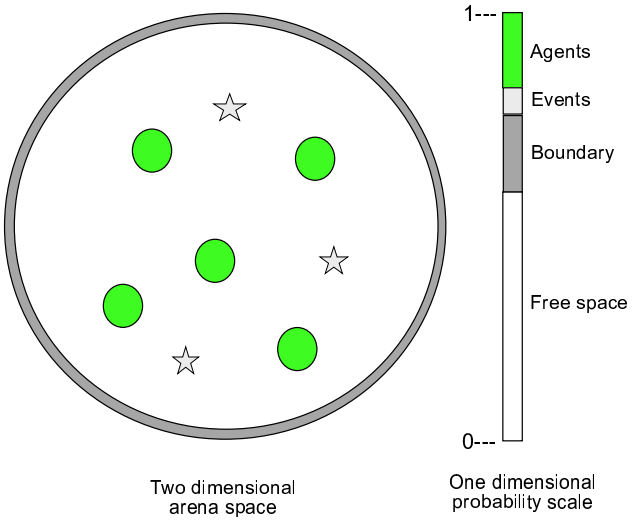

Figure 2.3: Example of transformation of a two dimensional arena space into a one dimensional probability scale. Extended areas of agents and events are determined by their detection ranges.

model. These PFSMs depict only the set of behaviors of the susceptible agents as the informed agents no longer search for the event (or other informed agents) to acquire the information. Each of the PFSMs is characterized by two types of state: several states that the agent can occupy for a limited amount of time, for example, a state of physical interference with other robots for τ_i iterations, a wall avoidance state for τ_w iterations, etc., and a default state, e.g., the event searching state. The transition probabilities between the states are calculated using geometrical considerations as presented in Equation (2.1). Note that in the system described in Figure 2.4 right, we do not include communication interference that may occur, for instance, when several agents transmit to a same receiver agent. In this graph, we assume that agents always successfully exchange information upon encounter. It is also worth noting that communication interference can easily be taken into account by distinguishing between the following two cases for the information exchange state: the case where transmission and reception of messages between the two agents take place without any other agent attempting to transmit to any of them during the process and the case where information exchange fails due to signal interference (the agents will then perform an obstacle avoidance maneuver and proceed with event searching).

In Figure 2.4 right we do not include the scenario where agents can simultaneously search for events and exchange information about their finding, but the simplicity of the model guarantees that this can easily be done. Note that agent-agent interference is different in the two cases represented in Figure 2.4. In fact in Figure 2.4 left, for any given agent, every other agent represents an obstacle to avoid (represented by the probability P_r), whereas in Figure 2.4 right there are only two scenarios of agent-agent physical interference: when two susceptible or two informed agents encounter (we represent these two scenarios by the probability P_s).

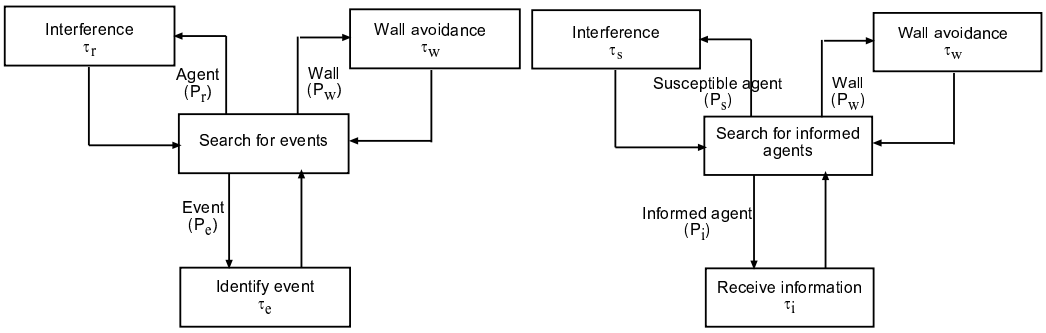

Figure 2.4: Example of Probabilistic Finite State Machine (PFSM) representing the susceptible agent's set of behaviors. Left: PFSM of the event discovery task without communication. Right: PFSM of the sensing task with local information exchange between agents.

In the embodied-models presented in this chapter, in the case of event discovery, we consider only the detection range of the event and in the case of information spread through local communication, we consider only the detection range of the agents (i.e., communication range and agent-agent detection range are assumed equal).

2.5 Macroscopic Models of Non-Embodied Agents

2.5.1 Nonspatial models

Non-spatial models describe the dynamics of the agents without taking into account their trajectories. This is equivalent to assuming that each agent hops around randomly in the region and occupies independent, consecutive positions. This work is inspired by studies of food source discovery and organization through information passing in ant colonies [81, 102, 1, 34].

A. Nonspatial models of event discovery without communication between agents

In this section we consider a system that consists of N autonomous, mobile agents capable of sensing different events present in the territory. Note that here, event discovery corresponds merely to activities such as monitoring and surveillance in the enclosed territory using multiple autonomous agents without any communication taking place between the agents.

Event with probability of spontaneous disappearance.— This is an event that presents the characteristic of random disappearance. Let γ_0 be the probability that an agent discovers a particular event during a particular time step. Let q_N be the probability that at least one of the N agents finds this event on a particular time step, since agents act in parallel during the course of a single time step, we then have:

$$q_N = 1 - (1 - \gamma_0)^N \tag{2.2}$$

Assuming that events disappear spontaneously per time step with probability p, the probability, say Φ_N, that the event is first discovered by one of the N agents is the ratio of the probability that agents find the event on a given time step to the probability that the event is found or disappears on the time step. Thus

$$\Phi_N = \frac{q_N}{1-(1-p)(1-q_N)} = \frac{q_N}{p+q_N-pq_N} \tag{2.3}$$

In the case where $\gamma_0 \ll 1$ and $N\gamma_0$ is small as well, $q_N \approx N\gamma_0$ then we obtain

$$\Phi_N \approx \frac{N\gamma_0}{p+N\gamma_0} = \frac{N\Phi_1}{1-\Phi_1+N\Phi_1} \tag{2.4}$$

where

$$\Phi_1 = \frac{\gamma_0}{p+\gamma_0} \tag{2.5}$$

Note that Φ_1 is the probability that the event is first discovered by a particular agent.

Equations (2.2)-(2.5) are the same as those used in [1] to describe the probabilities that ants in a colony discover a particular event (e.g., a food source) by wandering.

Event with constant duration.— An event with constant duration is one that lasts a fixed duration before disappearing from the environment. Before the event disappears, if at any time step, a single agent has probability ρ_1 of discovering that event, at the same time step the N agents in parallel have probability of success ρ_N given by

$$\rho_N = 1-(1-\rho_1)^N \tag{2.6}$$

This expression was proposed by Oster and Wilson [81] to describe foraging by groups of ants with the assumption that the individuals search and act independently from each other.

Comparing the probability of discovery of events with constant duration and that of events with spontaneous disappearance .— Let us assume that the probability that an event is first discovered by a particular agent during a time step, in the case of an event with spontaneous disappearance, is the same as the probability that the agent discovers a different event with constant duration, i.e., $\rho_1 = \Phi_1$. It appears that the probability of success of the group of N agents, i.e., ρ_N, in the latter scenario is larger than that of the group in search of the event with spontaneous disappearance[1]. Thus, as the group size increases, ρ_N increases to 1 much more rapidly than Φ_N. This difference between the probabilities of successful event discovery obtained for events that disappear spontaneously and those that have a constant duration is due to the different way events interact with the environment, not to a different way that agents interact with the events or each other. In other words, the probability of spontaneous disappearance of the first event reduces its likelihood of being discovered at any time

[1] This can be easily shown by inserting Equation (2.5) into Equation (2.2) and subtracting Equations (2.3) and (2.6) then, noting that $(1-(1-p)\Phi_1) \geq (1-\Phi_1)$, it follows that $\Phi_N \leq \rho_N$.

compared with the probability that the event with constant duration is first discovered by one of the agents.

Events that disappear when discovered.— When the event can disappear randomly, its expected residency time, T_{rand}, is independent of its age and corresponds to the inverse of the probability that at least one agent discovers the event during any time step. This is expressed as follows:

$$T_{rand} = \frac{1}{1-(1-p)(1-q_N)} \tag{2.7}$$

where q_N is the probability that at least one of the N agents finds the event on a particular time step as defined by Equation (2.2). Note that for values of γ_0 such that $q_N \ll 1$, we can approximate the value of the expected residency time as $T_{rand} \approx \frac{1}{p}$, which expresses the fact that the residency time would depend more on the probability of spontaneous disappearance as agents are less likely to discover the event. Equivalently, when $p \ll 1$ we have $T_{rand} \approx \frac{1}{q_N}$, which expresses the fact that the expected residency time depends more on the probability of the event being discovered by one of the agents as the event tends not to disappear on it own. For instance, it can be seen in the example presented in Figure 2.5 left, that the residency of the event converges to $1/q_N$ when the probability of disappearance (i.e., p) decreases.

In the case of an event with constant duration, the expected residency time, T_{const}, is expressed as follows (where T_0 stands for the constant duration of the event):

$$T_{const} = T_0(1-\rho_N)^{T_0-1} + \sum_{i=1}^{T_0-1} i(1-\rho_N)^{i-1}\rho_N = \frac{1-(1-\rho_N)^{T_0}}{\rho_N} \tag{2.8}$$

On the right-hand side of Equation (2.8), the first term represents the case where the event is not discovered before it disappears and the second term corresponds to the different cases where the event is discovered during a time step. Note that $\lim_{\rho_N \to 0} \frac{1-(1-\rho_N)^{T_0}}{\rho_N} = T_0$, in particular, if ρ_N is small such that $N\rho_N \ll 1$, then the expected residency time is approximately equal to the constant duration of the event as the event is unlikely to be noticed by any agent before it disappears, this is illustrated with the example in Figure 2.5 left.

Figure 2.5 presents a comparison of residency time, predicted or measured using the macroscopic model and simulations, respectively, for an event with constant duration and another with spontaneous disappearance. For each value of the constant duration (or the probability of disappearance) 100 simulation runs were conducted and the results are the average of those runs; error bars represent the standard deviations among the runs. For each plot, the x-axis represents the logarithm of the corresponding parameter, i.e., constant duration or probability of spontaneous disappearance.

From Figure 2.5 we can notice that as predicted above, on average, the residency time of the event with constant duration is larger than that of the event with spontaneous disappearance (when

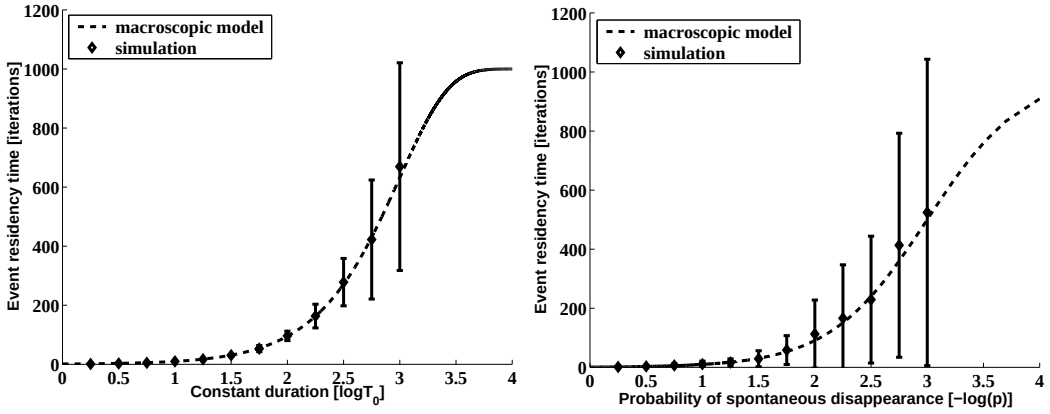

Figure 2.5: Macroscopic model predictions and simulation results of residency time of an event. Left: Plot of $T_{const}(T_0) = \frac{1-(1-\rho_N)^{T_0}}{\rho_N}$ with $\rho_N = 0.001$. Right: Plot of $T_{rand}(p) = \frac{1}{1-(1-p)(1-q_N)}$ with $q_N = 0.001$.

both events have the same probability of been discovered at a particular time step). Finally, Figure 2.5 shows a good agreement between the macroscopic model and simulations, and illustrates the accuracy of the proposed model.

B. Nonspatial models of information spread through local communication

We divide the population into two groups consisting of *Informed* and *Susceptible* agents, represented by $I(t)$ and $S(t)$, the respective numbers of informed and susceptible agents at time t. Moreover, we make the following assumptions: population size is large and constant, $I(t) + S(t) = N$, and no agent removal or new agent introduction.

Information spread through event discovery: stigmergic communication.— The concept of *stigmergic* communication [11, 44] comes from observations of biological systems where it has been found that agents "communicate" by modifying the environment so as to trigger as set of reactions from their teammates in response to those modifications [15, 102]. For instance, the laying of a particular pheromone in the nest can trigger a defensive behavior of individuals in ants colonies, during a nest building, the presence of construction materials at a location can prompt other ants to secrete substances that cement the materials transforming them into concrete blocks, and the accumulation of dead corpses in the nest can trigger a task switching behavior of foraging ants to take on the duty of nest cleaning. Therefore here, we assume that the stigmergic event is created by one of the teammates and used as a communication medium. We can think of this as a way for the mobile agents to communicate by leaving markers, e.g., alarm signs, at different locations. This type of events (e.g., alarm signs) would usually disappear after a certain amount of time. However, when the marker is not created by the teammates, this is simply equivalent to a one-event system. In both cases, clearly, the susceptible agents need to encounter the marker sign before being aware

of the information that it carries.

During the lifespan of the event we have the following process probability:

$$\Pr\{S \to S-1 \quad \text{and} \quad I \to I+1 \quad \text{in} \quad t, t+1\} = \alpha S(t)dt \tag{2.9}$$

where α represents the rate at which an agent encounters the event. If this process is linearized and after replacing $S(t)$ by the deterministic $s(t)$, we obtain the following first order differential equation

$$\frac{ds}{dt} = -\alpha s(t) \tag{2.10}$$

which solution is $s(t) = s_0 e^{-\alpha t}$ where s_0 is the initial value of s. Thus, eventually all the agents discover the event and become aware of the information it contains. Note that here we do not put an upper bound on the number of agents that can discover and/or sense the event during a given time step but in the real world, due to the embodiment of the agents and the position of the event, such upper would exist.

Information spread through local communication.— In this case, a susceptible agent becomes informed only by receiving the information (through local peer-to-peer communication) from an informed agent that it may encounter in the environment. We make the following additional assumptions: homogeneous mixing of informed and susceptible agents and information spread rate is proportional to the number of informed agents. Thus, the process probability is

$$\Pr\{S \to S-1 \quad \text{and} \quad I \to I+1 \quad \text{in} \quad t, t+1\} = \lambda I(t)S(t)dt \tag{2.11}$$

where λ is the encounter rate of two particular agents. From Equation (2.11) we derive a nonlinear first order differential equation that can be expressed in terms of the deterministic fraction of informed agents, denoted by $i(t)$, as follows:

$$\frac{di}{dt} = N\lambda i(t)(1 - i(t)) \tag{2.12}$$

This differential equation is known as the logistic growth equation, proposed by the mathematician *Pierre Verhulst* in 1845 for population growth. The equation above is separable, thus it is solved by dividing and integrating as follows:

$$\int_{i_0}^{i(t)} \frac{du}{u(1-u)} = \int_{i_0}^{i(t)} \left(\frac{1}{u} + \frac{1}{1-u}\right) du = \int_0^t N\lambda dt' \tag{2.13}$$

Some algebra gives

$$i(t) = \frac{1}{1 + \frac{1-i_0}{i_0} e^{-N\lambda t}} \tag{2.14}$$

with the initial condition $i(0) = i_0$. Equation (2.14) describes what is known as the *logistic curve*. It is clear that as $t \to \infty$, $i \to 1$, so eventually, every agent becomes informed. Note that $i(t)$ converges faster to 1 as the total number of agents increases, maybe a small illustration of the fact that "the more people there are in a given area, the more easily a rumor can spread in the area". We illustrate the result above in a later section with an information spread example and the graphs in Figure 2.10.

2.5.2 Spatial models

These models take into account the fact that agents do not hop randomly around the region but tend instead to remain in the same area during consecutive time steps. Multiple search of the same area occurs (more frequently here than in the simplified non-spatial model), as a consequence, efficiency is reduced. A similar behavior is observed in insect colonies and described in [21, 87, 26, 1]. A way of studying the relationship between the areas covered and the path of the mobile agent is to use a correlated random walk model [49, 17, 1].

A. Correlated random walk

A correlated random walk is a sequence of steps with (fixed or variable) step length and direction chosen from a probability distribution wrapped around the region (here assumed to be a circle arena) with mean equal to the direction of the previous step. Note that sharper turns are observed when the standard deviation of the turning angle becomes large. Figure 2.6 presents an example of correlated random walk in a 2-D space.

Biologists and ecologists have shown that animal trajectories are well described using a correlated random walk model [49, 17]. Our intuition is that a similar correlated random walk model can be applied to reactive, mobile robotic systems. For instance, when the movement pattern of the agents is designed to resemble a random walk for particular purposes, e.g., surface coverage, to avoid predators or being tracked, etc. Another example is the case where direction change happens only when the agent avoids an obstacle. In this case, we can look at the distance traveled by the agent between two consecutive encounters with obstacles and the turning angle as the step length and the turning angle of a correlated random walk; in such a system, the average step length is directly related to the density of agents as the agent-agent physical interference rate increases with the total number of agents and agents travel shorter distances between two consecutive encounters.

Agent's dispersal.— One interesting characteristic of a random walk is the dispersal of the agents from a starting region. In other words, when an agent starts out at a given position, its displacement over any observation period depends on the degree to which its path is straight. In fact, the higher the turning angle of the agent is, the more the agent tends to remain in the same region over time. Let $h(l)$ be the probability distribution of the length l of each move step and $g(\theta)$, the

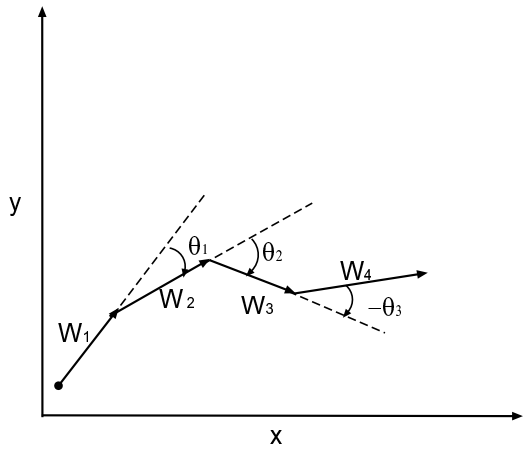

Figure 2.6: Correlated random walk. Path length has distribution $h(l)$ and turning angle has distribution $g(\theta)$ wrapped around previous direction.

probability distribution of the angle between two consecutive steps (measured clockwise). If each kth move is represented by the 2-D vector $W_k = (x_k,\ y_k)$, so that the total displacement after m consecutive moves, D_m, and the expected square displacement are given by Equations (2.15) and (2.16), respectively.

$$D_m = \sum_{k=1}^{m} W_k \tag{2.15}$$

$$E\left(D_m^2\right) = \sum_{k=1}^{m} E\left(W_k \cdot W_k\right) + 2\sum_{k>p} E\left(W_k \cdot W_p\right) \tag{2.16}$$

In Equation (2.16), $E(\cdot)$ represents the expectation of the associated random variable and $W_k \cdot W_p$ is the inner product of the step vectors $W_k = (x_k, y_k)$ and $W_p = (x_p, y_p)$. Since the movement is assumed to be a random walk, the cross-correlation among l_k, l_p, θ_k, and θ_p are taken to be zero, and we have the following equation

$$E\left(D_m^2\right) = mE\left(l^2\right) + 2E(l)^2 \sum_{k>p} E\left(\cos\sum_{j=p}^{k-1}\theta_p\right) \tag{2.17}$$

Using Euler's formulas, it comes that

$$\cos\sum_{i=p}^{k-1}\theta_i = \frac{1}{2}\left(\exp j\sum_{i=p}^{k-1}\theta_i + \exp -j\sum_{i=p}^{k-1}\theta_i\right) \tag{2.18}$$

Since the turning angles $\theta_i (i = 1,2,\ldots)$ are assumed to be independent of each other it follows that

$$E\left(\cos\sum_{i=p}^{k-1}\theta_i\right) = \frac{1}{2}\left[\left(E\left(\cos\theta\right)+jE\left(\sin\theta\right)\right)^{k-p}+\left(E\left(\cos\theta\right)-jE\left(\sin\theta\right)\right)^{k-p}\right] \tag{2.19}$$

With the additional assumption that $g(\theta)$ is symmetric about $\theta = 0°$, $E\left(\sin\theta\right) = 0$ and we obtain

$$E\left(\cos\sum_{i=p}^{k-1}\theta_i\right) = E\left(\cos\theta\right)^{k-p} \tag{2.20}$$

hence

$$E\left(D_m^2\right) = mE\left(l^2\right)+2E(l)^2\frac{E\left(\cos\theta\right)}{1-E\left(\cos\theta\right)}\left(m-\frac{1-E\left(\cos\theta\right)^m}{1-E\left(\cos\theta\right)}\right) \tag{2.21}$$

When the correlated random walk consists of paths of a constant length, say d, then Equation (2.21) becomes

$$E\left(D_m^2\right) = md^2+2d^2\frac{E\left(\cos\theta\right)}{1-E\left(\cos\theta\right)}\left(m-\frac{1-E\left(\cos\theta\right)^m}{1-E\left(\cos\theta\right)}\right) \tag{2.22}$$

It can easily be shown (by applying *L'Hospital*'s rule) that as $E\left(\cos\theta\right) \rightarrow 1$ (i.e., the distribution of θ around the direction of the previous step has zero mean and a standard deviation approaching zero), $E\left(D_m^2\right) \rightarrow m^2d^2$ which corresponds to a straight path movement.

Expected square displacement represents the more convenient parameter with which to quantify dispersal (as opposed to expected displacement) because the former is easier to estimate, moreover, as seen in the literature on random walks (and Wiener processes in general) [90, 22], for large numbers of time steps, the expected square displacement is a linear function of the time spent moving, e.g., the number of steps, whereas the expected displacement is not. We will elaborate more on this idea in section 2.5.3.

B. Spatial models of event discovery without communication between agents

In this section we only elaborate on correlated random walks with a fixed step length d. This model takes into account the fact that the area effectively covered by the agent is reduced due to overlap. In the following we estimate the overlap area covered by an agent and deduce the effective area swept out in search of events.

Correlated random walk with a low turning angle variance.— We first consider the case where the turning angle of the correlated random walk is small such that the agent does not tend to limit it coverage area to a very small region but rather has a large displacement over long observation bouts.

When the variance of the turning angle between two consecutive steps is small, we can approximate the overlapped area during the course of an event as the cumulative overlapped areas between

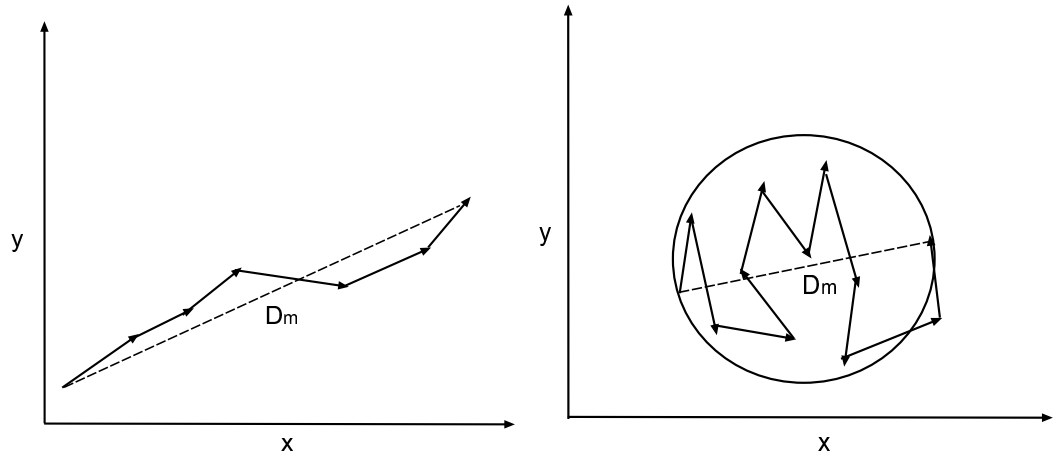

Figure 2.7: Correlated random walk. Turning angle has previous direction as mean. Left: Low variance of the turning angle. Right: Large variance resulting in sharper turns and limited displacement over time.

consecutive steps. This can be intuitively explained by the fact that even though the agent's path is not straight, its displacement over time is large, i.e., the agent tends to move away from its previous location (see Figure 2.7 left). Thus, we approximate the area effectively covered by the agent as the total area covered by a agent moving in a straight line minus the total overlapped area between consecutive steps. Let $a(\theta)$ represent the overlapped area between two consecutive steps, it is easily shown that

$$a(\theta) = \frac{1}{4} w^2 \tan \frac{|\theta|}{2} \tag{2.23}$$

In Equation (2.23), w is the event detection width of the agent (taken to be twice the event detection

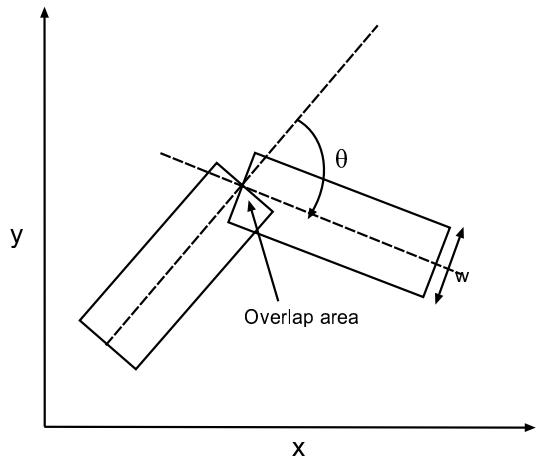

Figure 2.8: Correlated random walk. Overlap of coverage area between two consecutive steps.

range of the agent) and $|\theta|$ is the absolute value of θ. Equation (2.23) is valid for $|\theta| < \arctan(w/2d)$,

i.e., for not too large turning angles. Note that when $w < 2d$, a value of θ that violates the condition above is necessarily above $\frac{\pi}{2}$, which represents quite large a bifurcation. Thus, in the following we implicitly assume that $-\frac{\pi}{2} < \theta < \frac{\pi}{2}$. The expected value of the overlapped area between two consecutive moves, $E(a(\theta))$, is given by

$$E\left(a(\theta)\right) = \frac{1}{2}w^2 \int_0^{\pi} \tan\left(\frac{\theta}{2}\right) g(\theta) d\theta \tag{2.24}$$

Thus the total area effectively explored by the agent after k time steps is obtained by subtracting the total overlapped area from the total area covered by an agent moving straight, i.e., $k(wd) - (k-1)E(a(\theta))$. Hence, over long observation bouts, the fraction of the total area (of the arena) explored per step, γ, can be approximated as follows:

$$\gamma \approx \frac{wd - E\left(a(\theta)\right)}{\pi d_0^2} \tag{2.25}$$

From Equation (2.25) we deduce the probability that the event is first discovered by one of the agents (i.e., Φ_N and ρ_N as defined above for events with spontaneous probability of disappearance and events with constant duration, respectively) as follows:

$$\Phi_N = \frac{N\gamma}{p + N\gamma} \tag{2.26}$$

$$\rho_N = 1 - (1-\gamma)^N \tag{2.27}$$

Note that $\gamma = \gamma_0 - \frac{E(a(\theta))}{\pi d_0^2}$, thus, the performance of the agent in discovering events is degraded by its propensity to remain in the same area.

Example: uniform distribution. Let us consider the uniform turning angle distribution below.

$$g(\theta) = \begin{cases} \frac{1}{\beta}, & \text{if } \frac{-\beta}{2} < \theta < \frac{\beta}{2} \\ 0 & \text{otherwise} \end{cases} \tag{2.28}$$

It comes that

$$E\left(a(\theta)\right) = -\frac{w^2}{\beta} \ln\left(\cos\frac{\beta}{4}\right) \tag{2.29}$$

and finally, the corresponding value of the fraction of area effectively explored per time step, denoted $\gamma_{uniform}$, is given by

$$\gamma_{uniform} = \frac{wd - \frac{w^2}{\beta} \ln\left(\cos\frac{\beta}{4}\right)}{\pi d_0^2} \tag{2.30}$$

Correlated random walk with a large turning angle variance.— In this second scenario of event discovery, we investigate the influence of large turning angles on the total area explored. In fact, due to the high turning angle between consecutive steps, only a small region is explored as repeated

investigation of the same locations is frequent, as a consequence, the agent's efficiency is highly decreased (see Figure 2.7 right). In the following we assume that the event detection width of the agent, i.e., w, is very small related to its displacement over a large number of steps.

We have an estimate of the expected square displacement of the agent over any given number of steps k given by Equation (2.21); taking into account the directionality of the paths, we assume that search remains concentrated in a circle of diameter the expected displacement of the agent. The area of this circle, as a function of the number of steps, is $\frac{\pi}{4}E(D_k^2)$. Thus the fraction of the total area explored by step is given by

$$\gamma(k) = \frac{E\left(D_k^2\right)}{kd_0^2} \tag{2.31}$$

From Equations (2.21) and (2.31) it follows that for large numbers of steps the fraction of the total area explored per step is

$$\gamma = \frac{d^2}{d_0^2}\frac{1 + E(\cos\theta)}{1 - E(\cos\theta)} \tag{2.32}$$

Finally, this value of γ gives the new values of Φ_N and ρ_N, the probabilities that at least one agent discovers the event at any given time step, when this event has a spontaneous probability of disappearance or a fixed duration, respectively (see section above).

Example: uniform distribution. Using the uniform distribution introduced above, we obtain the following value of $\gamma_{uniform}$:

$$\gamma_{uniform} = \frac{d^2}{d_0^2}\frac{1 + \frac{\sin\beta/2}{\beta/2}}{1 - \frac{\sin\beta/2}{\beta/2}} \tag{2.33}$$

To have an idea of the degradation of the agent's efficiency due to repeated coverage of the same area, we compare the estimated coverage area per time step for the low (Equation (2.30)) and large (Equation (2.33)) turning angle cases for the uniform distribution defined by Equation (2.28). For small and large turning angles, we arbitrarily consider $\beta = 20^\circ$ and $\beta = 320^\circ$ respectively and assume that the agent detection range is the same order of magnitude as the fixed path length, i.e., $w \approx d$. We found that $\gamma(\beta = 20^\circ) \approx 2.52\frac{d2}{d_0^2}$ and $\gamma(\beta = 320^\circ) \approx 1.32\frac{d2}{d_0^2}$; in other words, for a uniform distribution of the turning angle around the previous direction and with the assumptions above, the average area effectively explored per step with the narrow turning angle is approximately twice as large as that explored with the large turning angle.

Example: normal distribution. Let us consider a normal formulation of the random walk consisting of a sequence of steps with fixed length d and with direction chosen from a normal distribution with mean the previous direction and standard deviation σ. The corresponding fraction of the area effectively explored per step, denoted γ_{normal}, is given by

$$\gamma_{normal} = \frac{d^2}{d_0^2}\frac{1 + e^{-\sigma^2/2}}{1 - e^{-\sigma^2/2}} \tag{2.34}$$

It is worth noting that high values of σ correspond to high turning angles which are necessary to keep the agents from wandering too far away from a chosen region.

C. Spatial models of event discovery with information spread through local communication between agents

There is an obvious tendency in many multi-agent systems for susceptible agents that are situated a long way from informed agents to have a smaller chance of receiving the information than those in a closer contact with the same informed agents. In the following, we assume that upon encounter agents exchange information locally without any communication interference.

In 1937, *Fisher* [28] introduced a model to describe the invasive spread of a diffusively dispersing organism with logistic local dynamic, in particular an advantageous gene in a population. Fisher's equation is the following

$$\frac{\partial u}{\partial t} = g_r u(1-u) + D\frac{\partial^2 u}{\partial x^2} \tag{2.35}$$

where u is the density (or spatial fraction) of individuals with the gene at location x, g_r is the intrinsic growth rate, and D is the diffusion coefficient. This model in itself does not have such wide application. It is just the prototype equation which admits traveling wave front solutions and is also a convenient equation from which to develop many standard techniques for analyzing single species models with diffusive dispersal [76, 77]. Moreover, Equation (2.35) only applies to a linear population dispersal; in other words, the population is assumed to be distributed along a line! This equation and its traveling waves solutions have been widely studied and the properties of the model are well known [90, 8, 38, 79]. From the introduction of the individuals with the favored gene, the local population of individuals with the gene grows into the surrounding territory. Eventually the invaded and uninvaded regions are separated by a narrow *invasion front* which moves into the uninvaded territory at constant speed. Although the front velocity formally depends on the initial condition, all initial conditions with compact support excite a front traveling at the minimum possible speed $V_{min} = 2\sqrt{g_r D}$.

In the following we apply and generalize Fisher's equation to swarm-based networks of mobile agents and propose two methods for studying the dynamics of the population of informed agents. In the first model we analyze the qualitative dynamics of the population of informed agents as it presents a wave like propagation. In the second method we present a quantitative analysis of the dynamics of the same population of informed agents as it grows into the surrounding regions.

Let s and i represent the spatial densities of susceptible and informed agents, respectively, i.e., the respective fractions of the two types of agents at any given spatial location. These densities are function of both time and local positions, i.e., $s \equiv s(\xi, \eta, t)$ and $i \equiv i(\xi, \eta, t)$, where ξ and η are taken to be the spatial coordinates. Assuming that no new agents are introduced during the experiment,

the generalized Fisher's equation describing the dynamics of $i(\xi, \eta, t)$ can be written as:

$$\frac{\partial i}{\partial t} = \lambda s(\xi, \eta, t) i(\xi, \eta, t) + D\nabla^2 i(\xi, \eta, t) \tag{2.36}$$

where λ represents the encounter rate of two particular agents in the region (same as in previous sections) and D is the diffusion coefficient of the information propagation, this latter parameter also translates the spatial dispersal rate of the agents. An important element of these equations, i.e., Equation (2.35) and Equation (2.36), is the introduction of the Laplace (divergence) operator, i.e., ∇^2, which emerges from the assumption that information spread is local and isotropic. We will now elaborate on Equation (2.36) (which does not have any closed form solutions [8, 79]) to study the qualitative and quantitative dynamics of the population of informed agents.

Qualitative dynamics of the population of informed agents.— We assume that the spatial distribution of the agents is axisymmetric, i.e., it is a function of only the distance between the origin and the location. Thus, Equation (2.36) can be rewritten in axisymmetric form, with r, the distance to the origin.

$$\frac{\partial i}{\partial t} = \lambda i(1 - i) + D\left(\frac{1}{r}\frac{\partial i}{\partial r} + \frac{\partial^2 i}{\partial r^2}\right) \tag{2.37}$$

In the right-hand of Equation (2.37), the first term translates the nonspatial spread of the information through local interaction between the agents and the second term translates the spatial diffusion of the information resulting from the mobility of the agents. If we start for $0 < r < \infty$ and an initial condition that $i(r, 0)$ has a compact support, i.e., $i(r, 0)$ is zero outside a finite domain (see Figure 2.9 left), the solution $i(r, t)$ will evolve as follows. If $i(r, 0) < 1$, the $i(1 - i)$ term is positive thus it causes the solution to grow until $i = 1$; this is illustrated in Figure 2.9 center and right where it can be seen that the local density $i(r, t)$ increases over time to reach a plateau at large t. At the same time diffusion causes wave-like dispersal outwards; this can be seen in Figure 2.9 right where it appears that the local informed agents *invade* the neighboring territories. On the *wave*, $\frac{\partial i}{\partial t} < 0$ so it reduces the term $i(1-i)$. The effect is the decrease of the velocity of the outgoing wave. For large values of r the $\frac{1}{r}\frac{\partial i}{\partial t}$ term becomes negligible so the equation becomes exactly Fisher's equation which solution will tend to a traveling wave-front with velocity $V = 2\sqrt{\lambda D}$. Thus, the total amount of time required for the information to spread to the entire region (of radius d_0) can be approximated as $T = \frac{d_0}{2\sqrt{\lambda D}}$.

Quantitative dynamics of the population of informed agents.— Initially, when the density of susceptible agents can be taken as approximately constant, Equation (2.36) is a standard diffusion equation [8]. This assumption holds more strongly when the number of agents is very large and the density of susceptible agents is much larger than that of informed agents at start. Thus, the

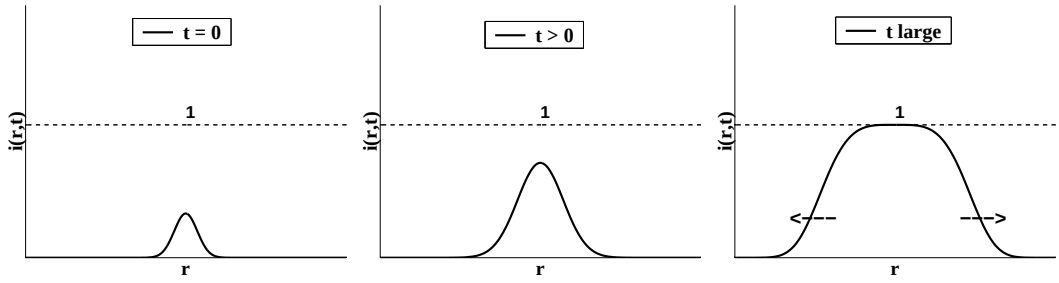

Figure 2.9: Qualitative dynamics of the population of informed agents: example of time evolution of the spatial density of informed agents $i(r,t)$. Left: $i(r,0)$. Center: Increase of $i(r,t)$. Right: Outwards diffusion into surrounding area.

dynamics of the spatial density of informed agents, i.e., $i(\xi,\eta,t)$, can be rewritten as follows:

$$\frac{\partial i}{\partial t} = Ai(\xi,\eta,t) + D\nabla^2 i(\xi,\eta,t) \tag{2.38}$$

where $A = \lambda s$. Equation (2.38) can be solved by means of Fourier transform for spatial coordinates. If $I(p,q,t)$ stands for the Fourier transform for the spatial coordinates of $i(\xi,\eta,t)$, then

$$I(p,q,t) = \iint \exp\left(jp\xi + jq\eta\right) i(\xi,\eta,t) d\xi d\eta \tag{2.39}$$

and $I(p,q,t)$ verifies

$$\frac{\partial I}{\partial t} = AI - D(p^2+q^2)I \tag{2.40}$$

thus,

$$I(p,q,t) = C\exp\left(At - D(p^2+q^2)t\right) \tag{2.41}$$

Finally, by means of inverse Fourier transform we obtain $i(\xi,\eta,t)$

$$i(\xi,\eta,t) = \frac{1}{2\pi}\iint \exp\left(-jp\xi - jq\eta\right) I(p,q,t) dp dq \tag{2.42}$$

which brings

$$i(\xi,\eta,t) = \frac{C}{2Dt}\exp\left(At - \frac{r^2}{4Dt}\right) \tag{2.43}$$

where r represents the distance to the information source, i.e., $r = \sqrt{\xi^2+\eta^2}$, and C is determined by the initial conditions $i = i_0$, at $t = t_0$. Let I_{R_d} be the total number of informed agents outside a circle of radius R_d, we have

$$I_{R_d} = 2\pi\int_{R_d}^{\infty} i(r,t) r dr = 2\pi C\exp\left(At - \frac{R_d^2}{4Dt}\right) \tag{2.44}$$

It follows that

$$R_d = 2\sqrt{AD}t\sqrt{1 - \frac{\ln\left(\frac{I_{R_d}}{2\pi C}\right)}{At}} \tag{2.45}$$

Equation (2.45) expresses a spatial increase of the number of informed agents, i.e., the propagation of a front that grows at an increasing rate R_d/t with limiting velocity $V = 2\sqrt{AD}$. Eventually all the agents become informed.

2.5.3 Calculating model parameters

In this section we present how the parameters of the proposed models of non-embodied agents can be calculated using the characteristics of the system, and we compare information spread results in a same system using stigmergic, i.e. implicit, and explicit communication mechanisms.

A. Nonspatial models

Event discovery without explicit local communication.— The main parameter of this model is α, the rate at which an event is encountered. If the agent travels at an average speed v then during time period dt the agent sweeps out an area of $v(2d_e)dt$ where d_e is the event detection range ($w_e = 2d_e$ is the event detection width); the agent will then encounter the event at a rate approximated as $\alpha = \frac{vw_e}{\pi {d_0}^2}$ where d_0 is the radius of the arena as previously introduced. Thus the fraction of informed agents is

$$i(t) = 1 - x_0 e^{-\frac{vw_e}{\pi {d_0}^2}t} \tag{2.46}$$

Event discovery with local communication.— The main parameter of this model is λ, the rate at which two particular agents encounter that we calculate as follows. Let d_a be the maximum distance from which any given agent can detect another agent and w_a the agent detection width (assumed to be the double of the detection range). During time period dt an agent sweeps out an area of $vw_a dt$, thus, similarly to the calculation of α, we approximate the agent-agent encounter rate as $\lambda = \frac{vw_a}{\pi {d_0}^2}$. Thus the expression of $i(t)$ becomes

$$i(t) = \frac{1}{1 + \frac{1-i_0}{i_0} e^{-N\frac{vw_a}{\pi {d_0}^2}t}} \tag{2.47}$$

Comparing information spread in event discovery with and without local communication between agents.— In Figures 2.10 and 2.11 we present two examples of the speed at which information propagates in a same system, when not using communication (Figure 2.10) and when using explicit local communication (Figure 2.11) for information spread. The graphs in Figures 2.10 and 2.11 were obtained for an arena of radius $d_0 = 1000$ units, a variable event detection width w_e (400 to 1200 units), an agent detection range $d_a = 10$ units, an agent mean speed $v = 10$ distance units per time unit, and an initial fraction of informed agents $i_0 = 0.01$. A team of 100 agents was used.

Note that the timescale of the simulations below is in number of iterations. Therefore, to be able to compare the macroscopic model predictions and the results of the simulations on the same timescale we need to define a conversion coefficient that links any given number of iterations in the simulations to a given number of time units in the macroscopic model. In [63], Martinoli proposed a way of converting iterations to a real timescale and vice versa based on simple geometrical considerations as follows. Without loss of generality, let us consider the case of event discovery without communication. In this scenario, the rate at which a particular agent encounters the event is given by $\alpha = vw_e/\pi d_0^2$ thus, on average, the agent will encounter the event once every $1/\alpha = vw_e/\pi d_0^2$ time units. Similarly, for the simulations, the probability that the agent encounters the event is $P_r = \pi(\frac{w_e}{2})^2/\pi d_0^2$ thus, on average, the agent will encounter the event once every $1/P_r = \pi d_0^2/\pi(\frac{w_e}{2})^2$ iterations. From the two results above we deduce the conversion coefficient denoted τ_{i2t} ("i2t" for "iteration to time") as $\tau_{i2t} = \frac{\pi w_e}{4v}$ (expressed in time units per iteration). A conversion coefficient is calculated for the scenario of information spread through local communication, in a similar way. However, since the value of the conversion coefficient is unique for any experiment, in this calculation we implicitly assume that the agents encounter only one type of element, i.e., an event or another agent for information exchange.

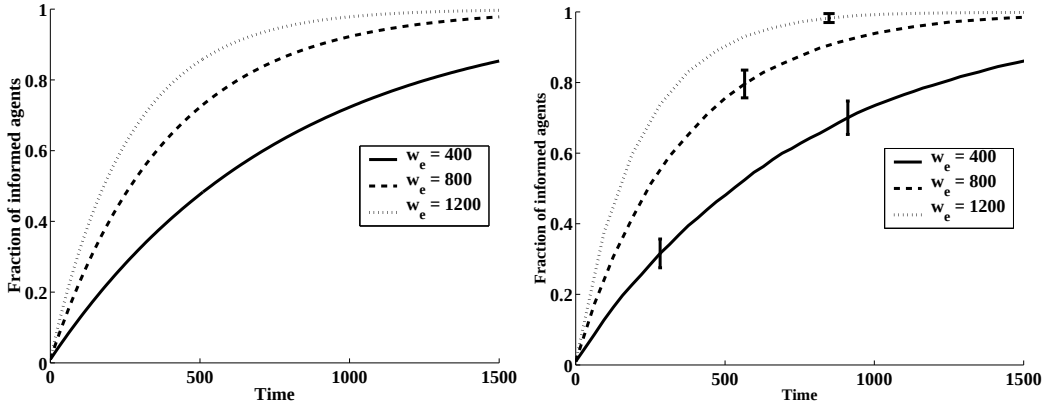

Figure 2.10: Evolution of the fraction of informed agents in an arena of radius $d_0 = 1000$ units in a case of event discovery without communication. Left: Macroscopic model. Right: Non-embodied (point) simulations.

The graphs in Figures 2.10 right and 2.11 right were obtained using the nonspatial, embodied, microscopic model presented in subsection 2.6.1 by neglecting the embodiment of the agents and all agent-agent and agent-wall physical interference. We later analyze the effects of such simplifications. Figures 2.10 and 2.11 illustrate a good overall agreement between the macroscopic model and the microscopic simulations. Figure 2.11 shows that the speed at which the information spreads to all the different agents of the system is an increasing function of the group size when using explicit communication and is accelerated when the number of informed agents grows. This figure also

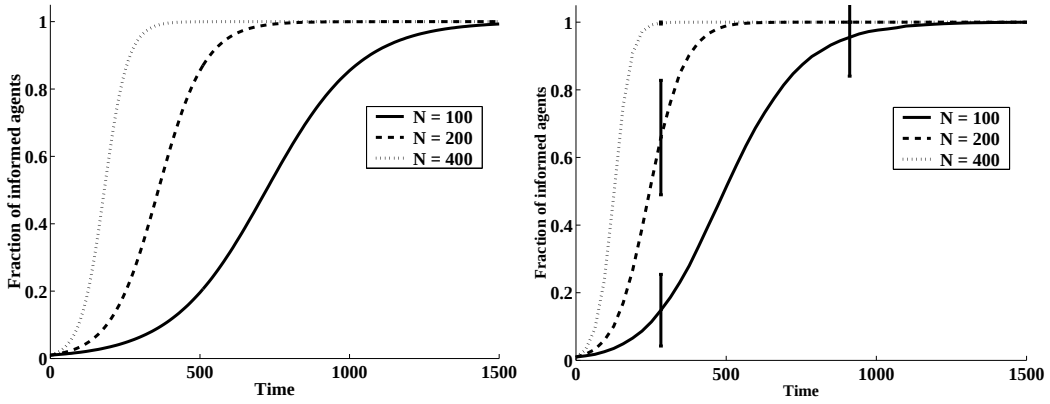

Figure 2.11: Evolution of the fraction of informed agents in an arena of radius 1000 units in a case of information awareness spreading through explicit, local peer-to-peer information exchange mechanism. Left: Macroscopic model. Right: Non-embodied (point) simulations.

clearly shows that for a similar detection range, information spreads faster with explicit, local, peer-to-peer communication between the agents than it does without communication where each agent, in parallel, has to discover a marker (a posted sign) to retrieve the information that it carries. Finally, it is worth noting that to obtain similar performances between the two information propagation schemes, the event detection range of the non communicating agents has to be very large. This may present a design challenge and require each agent to allocate a large fraction of its power resources to a single behavioral maneuver (e.g., event discovery) to the detriments of other maneuvers.

C. Spatial models

The main parameters of the spatial models proposed in section 2.5.2 are λ, the rate at which two particular agents meet and D, the diffusion coefficient of Fisher's equation. While the calculation of λ is known, that of D requires a further analysis. A more practical way of defining D is that it represents the isotropic *emigration* rate of the agents, it can also be easily understood as the isotropic expansion over two consecutive steps in a region where the agents are initially located. Let us consider again the correlated random walk model. According to the square-root law of random walk in two dimensions, the expected square displacement, denoted $< r^2 >$, over a large number of steps, k, is given by

$$< r^2 >= 4Dk \tag{2.48}$$

By comparing Equations (2.22) and (2.48), we derive the value of the diffusion coefficient

$$D = \frac{d^2}{4}\frac{1 + E(\cos\theta)}{1 - E(\cos\theta)} \tag{2.49}$$

Example: Uniform distribution. Let us consider the uniform distribution function $g(\theta)$ of

subsection 2.5.2. The diffusion coefficient, denoted $D_{uniform}$, for this distribution is given by

$$D_{uniform} = \frac{d^2}{4} \frac{1 + \frac{\sin \beta/2}{\beta/2}}{1 - \frac{\sin \beta/2}{\beta/2}} \tag{2.50}$$

Example: Normal distribution. Let $g(\theta)$ be a normal distribution function of zero mean and standard deviation σ as that introduced in subsection 2.5.2. The diffusion coefficient, denoted D_{normal}, for this distribution is given by

$$D_{normal} = \frac{d^2}{4} \frac{1 + e^{-\sigma^2/2}}{1 - e^{-\sigma^2/2}} \tag{2.51}$$

2.6 Macroscopic Embodied Models

In the following, agents are no longer considered as bodiless mobile points, instead, due to their embodiment, they can interfere with each other and trigger obstacle avoidance behavior. Thus, in the following we first extend the nonspatial macroscopic presented in subsection 2.5.1 into new nonspatial models that describe the embodied agents and their interactions with each other. For the study of spatial embodied models, we present a case study concerned with information dissemination in a network of simulated robots that we simulate with a 3-D, sensor-based simulator of real, mobile robots. In the case study, we introduce a mathematical method for finding the communication range that optimizes the local information exchange rate between the agents of the network by maximizing the probability of successful communication transmission in the network.

It is important to note that in the case of information spread through local communication we do not distinguish the agent-agent detection range from the agent-agent communication range. However, this can easily be done by distinguishing agent-agent encounters that result simply in collision avoidance (i.e., informed-informed and susceptible-susceptible encounters) from those that result in information passing (i.e., informed-susceptible encounters) and calculating the values of the corresponding encounter rates appropriately. For instance, in Equation (2.53), there would be two different λ, one (say λ_{com}) that would translate the local communication (for the first term in the right-hand side of the equation) and another (say λ_{col}) that would express the agent-agent collision (for the last two terms in the right-hand side of the equation).

2.6.1 Nonspatial embodied model

We extend the nonspatial models presented in subsection 2.5.1 to take into account the embodiment of the agent and the physical interference that results from the embodiment. In the following, we assume that the agents' movement pattern embeds the avoidance of the surrounding boundary of the arena and we neglect the amount of time the agents spend avoiding that obstacle. This

is similar to the assumption that the agents' movement pattern is increasingly biased toward the inner arena as they move closer to the boundary. Our main motivation for neglecting the influence of the surrounding wall is our particular interest in assessing the influence of embodiment on the agents' performance (e.g., effect of low/high rate of agent-agent interference on event discovery and information spread).

In the mathematical expressions below, we describe the collision avoidance maneuver as a simple delay state (as depicted in Figure 2.3). This results in differential-difference equations, which have the particularity of being defined over an infinite (time) space. Therefore we introduce the step function $\Theta(\cdot)$ which is such that $\Theta(t)$ is zero for $t < 0$ and one otherwise. Moreover, in all the models below, we assume that there are no informed agents until the start (i.e., for $t < 0$).

A. Nonspatial model of event discovery with no communication between agents

Equation (2.10) can be modified to describe the fact that, for any given agent, any other agent represents an obstacle to avoid to prevent (damaging) collision. Thus, the dynamics of the susceptible agents can be formalized as follows:

$$\frac{ds}{dt} = -\alpha s(t) - \lambda s(t) + \lambda s(t - \tau_{obs})\Theta(t - \tau_{obs}) \tag{2.52}$$

In Equation (2.52) α represents the rate at which the agents encounter the event, λ is the encounter rate of two particular agents in the arena, and τ_{obs} is the amount of time required to perform an agent-agent collision avoidance maneuver. In the right-hand side of this equation, while the first term corresponds to the susceptible agents that have just encountered the event, the second term corresponds to those that have just started a collision avoidance maneuver, and the third term translates the return of some susceptible from avoiding a teammate.

Related to Equation (2.10), the last two terms on the right-hand side of Equation (2.52) translate the influence of embodiment. Moreover, a further comparison between the two equations reveals that the right-hand sides of both equations are always negative and that of Equation (2.52) is greater than that of Equation (2.10). As a consequence, the number of susceptible (resp. informed) agents, in both systems, is monotonically decreasing (resp. increasing) and decreases (resp. increases) more slowly (resp. less rapidly) in the embodied system. This can be explained by the fact that embodied agents spend more time in collision avoidance that is ignored in the non-embodied model. This is further illustrated in Figure 2.12 where is appears that the rate at which the agents discover the event is, on average, higher when the agents are modeled as points. In addition, as expected, Figure 2.12 also illustrates that the destructive effect (of embodiment) on the team performance increases with the body size of the agents.

Figure 2.12 was obtained using a group of 100 agents in an arena of radius $d_0 = 1000$ units. For

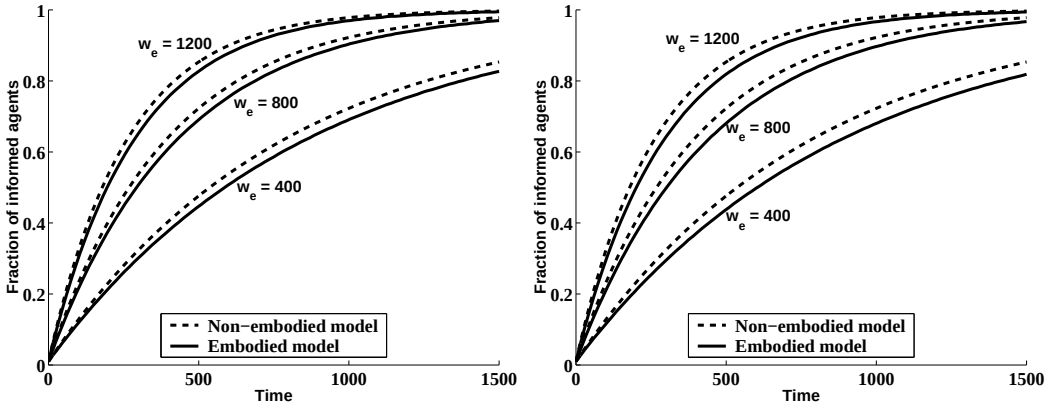

Figure 2.12: Comparing nonspatial models of embodied and non-embodied agents in event discovery without communication. Left: Embodied agents have a radius of 5 units. Right: Embodied agents have a radius of 10 units.

the embodied model, we considered an arbitrary collision avoidance duration of 10 time units and an initial fraction of informed agents $i_0 = 0.01$. $\alpha \approx 31.10^{-6}w_e$, $\lambda \approx 95.10^{-6}N$ per time unit for robot radius of 5 units, and $\lambda \approx 127.10^{-6}N$ per time unit for a radius of 10 units, were calculated using the method proposed in subsection 2.5.3.

B. Nonspatial model of information spread through local communication between agents

Similarly, we modify Equation (2.12) to take into account collision avoidance resulting from the physical interference between the agents. Thus, the dynamics of the fraction of informed agents can be described as follows:

$$\frac{di}{dt} = N\lambda i(t)(1 - i(t)) - N\lambda i(t)(i(t) - \frac{1}{N}) + N\lambda i(t - \tau_{obs})(i(t - \tau_{obs}) - \frac{1}{N})\Theta(t - \tau_{obs}) \qquad (2.53)$$

In Equation (2.53) λ is the same as that of Equation (2.12) and represents the rate at which two particular agent encounter in the arena. The last two terms in the right-hand side of Equation (2.53) translate the dynamics of informed agents engaged in collision avoidance. However, as opposed to the case of Equation (2.52), here not every agent represents an obstacle to avoid systematically. In fact, when an informed agent encounters a susceptible agent, the former transmits the information to the latter. Thus, only informed-informed and susceptible-susceptible encounters represent active interferences. The "$-\frac{1}{N}$" term in Equation (2.53) is used to single out a given informed agent from the rest (as a given agent cannot represent an obstacle to itself).

Figure 2.13 presents a comparison of (nonspatial) embodied and non-embodied models of information spread through local communication between agents. The results in this figure were obtained in an arena of radius $d_0 = 1000$ distance units and using an arbitrary collision avoidance duration $\tau_{obs} = 10$ time units. In both cases, we start with a same initial fraction of informed agents $i_0 = 0.01$.

The left and right graphs of Figure 2.13 were obtained using teams of embodied agents of radius 5 and 10 units, respectively.

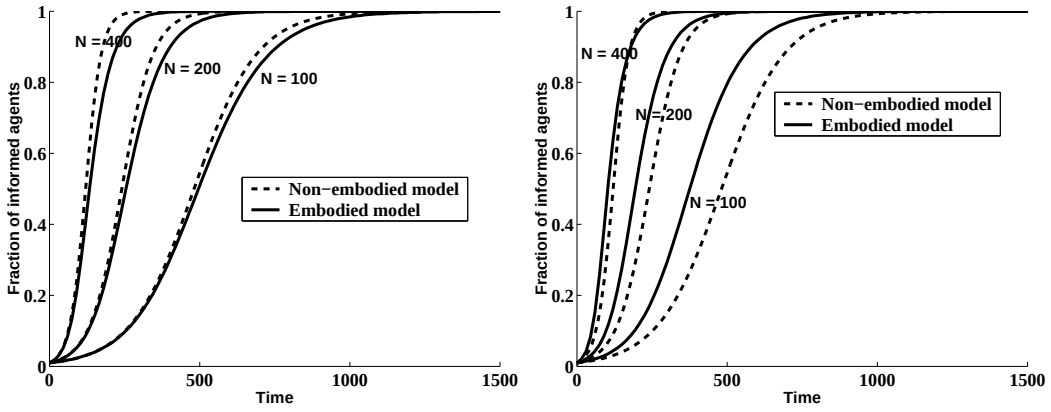

Figure 2.13: Comparing nonspatial models of embodied and non-embodied agents in information spread through local communication. Left: Embodied agents have a radius of 5 units. Right: Embodied agents have a radius of 10 units.

From Figure 2.13 it appears that, as opposed to the event discovery case, the influence of embodiment on the performance of the agents is not always destructive. Here, although frequent collision avoidance slows information spread, the embodiment of the agent, by increasing the likelihood of agent-agent encounter, increases that of information exchange between agents as well. Thus, as a function of the characteristics of the system, one of these two interrelated effects (one positive and the other negative) of embodiment can be dominant. For instance, for small values of the agents's radius (i.e., small body size) and a (relatively) large collision avoidance duration, the destructive effect is dominant and the embodied team performance is worse than that of the non-embodied agents team. This can be seen in Figure 2.13 where (for the same collision avoidance duration) the left graph shows that the rate of information spread in the network of embodied agents is lower than that in the network of non-embodied agents, whereas in the right graph, the larger embodied agents outperform the non-embodied agents.

Although a further mathematical investigation may reveal whether or not there exists a fixed bifurcation point where one of the two effects becomes dominant, the fact that only embodied agents exist in the real world simply suggests that the best approach in designing the system above is to increase the agent body size while minimizing the collision avoidance duration (when all else are equal). Moreover, note the results and analysis above are conditioned conditioned by the assumption that agents communicate only upon detecting each other, i.e., agent-agent encounter results in collision avoidance for the cases informed-informed and susceptible-susceptible encounters and in local information in the case informed-susceptible encounter. As stated previously in section 2.5.1, if we assume different communication and collision detection ranges, Equation (2.53) will change as

well (e.g., there would be two different λ). Nevertheless, based on the results presented above, it is clear that for increasingly larger communication ranges than collision detection ranges the rate of information spread will increase as well (when all else are equal).

Finally, we notice a good agreement between the nonspatial embodied macroscopic models and the microscopic model presented in section 2.4 as illustrated by Figure 2.14. The results in Figure 2.14 were obtained under the same conditions as above.

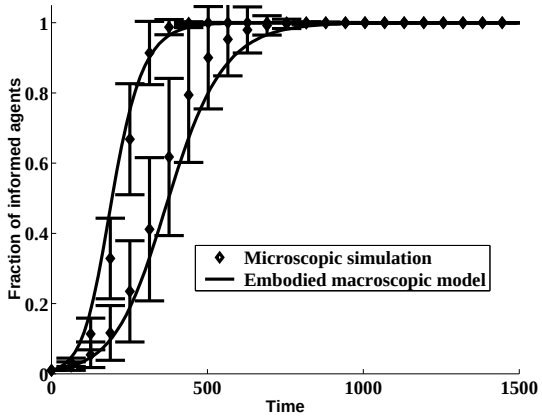

Figure 2.14: Comparison between (nonspatial embodied) macroscopic model predictions and microscopic simulations of information spread through local communication between agents.

2.6.2 Spatial embodied model

The idea that maximizing the rate of information spread in a network of mobile robots can be (fairly accurately) done by maximizing the average probability of local, peer-to-peer, successful information exchange between the robots was first introduced by Yoshida and coworkers [104]. However, their model is nonspatial as it makes no account of the trajectories or the positions of the agents over time. In the following, we extend the model into a spatial model and propose a general method for giving (mathematically) a design parameter (the communication range) that maximizes the average probability of successful information exchange (see also [5]). Although in our model we do not specifically take into account the trajectories of the agents, we consider spatiality through the distribution that results from the movement pattern of the agents. We illustrate that when compared with Yoshida's model, the predictions are more accurate when the model takes into account the distribution of the agents. Finally, we validate our model by comparing the predictions with sensor-based simulations.

A case study of information exchange in a network of simulated robots

Roughly speaking, if we denote by n the average number of sensors within the transmission radius of any given sensor, n is clearly an increasing function of the radius and may also depend on the location of the host, provided that the sensor can receive information from at most C (referred to as the *information reception capacity* of the agent) other sensors at any given time, we see that there must exist an optimal value of the transmission radius, which maximizes the obtainable successful transmission rate and hence the information exchange rate. This case study elaborates on these ideas.

In the following we assume that sensors are capable of transmitting and receiving information simultaneously and we do not account for communication protocols. We also assume that each agent transmits messages continuously by broadcasting within its communication range. Note that in the system considered here, information exchange does not happen systematically when two agents come within communication range, instead only the agents that receive less than C messages at a time are able to process the messages.

A. Mathematical analysis

In the following, we show how a simple macroscopic model can also serve as an optimization tool. We develop a model to accurately predict a control parameter (i.e., the communication range) that maximizes the information exchange rate between the mobile sensors. The proposed model makes use of very little a priori information about the system, as it only requires that the distribution of the agents in the environment be provided. This can be in the form of the probability of finding any given agent at a distance r from the center on the arena at any time. We refer to that distribution of the fraction of agents (expressed in percentage of agents per area unit) as $f(r)$, where r represents the distance between the center of the sensor, i.e., its antenna, and the center of the environment (taken as the spatial origin). Note that $f(r)$ depends on the movement pattern of the agents and also represents both the fraction of the team present at a distance r and the probability of finding any agent at the same distance r from the origin at any given time.

As seen in the literature [73], given P, the probability of successful transmission, the probability that it takes k time steps to succeed in transmission is $P(1-P)^{k-1}$ and the average information transmission time, W, is given by Equation (2.54) as follows:

$$W = \sum_{k=1}^{\infty} kP(1-P)^{k-1} = \frac{1}{P} \tag{2.54}$$

From Equation (2.54) it follows that the communication range, denoted R_{opt}, that maximizes P (when all else are equal), will in turn minimize W, thus, will increase the information exchange rate in the network.

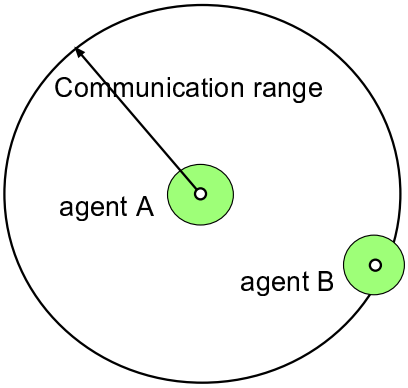

Figure 2.15: Two mobile sensors within communication range. Peer-to-peer information exchange happens through a local, broadcast wireless communication.

For any given agent in the environment, the average probability that the agent successfully receives information sent by another agent (i.e., P) is derived as a function of the communication range R in Equation (2.55), where d_0 is the radius of the arena, $n(R,r)$ is the average number of agents within the transmission area determined by the communication range R at location r. In other words, with the design parameter considered here, i.e., the communication range R, it is clear that $P \equiv P(R)$.

$$P(R) = \frac{1}{d_0} \int_0^{d_0} f(r) \left(\sum_{i=1}^{C} \frac{n(R,r)^i}{i!} \right) e^{-n(R,r)} dr \tag{2.55}$$

We refer to P as an average probability because it represents the average probability of successful transmission over all different distances to the center of the arena (taken as the spatial origin).

Note that in Equation (2.55), while $f(r)$ represents the probability of any agent being at location r, the rest of the integral represent the probability to encounter less than C agents within the communication area defined by R following a Poisson distribution. A Poisson distribution is used because when objects are randomly disposed on a plane the number of objects in a certain area is Poisson-distributed and the same distribution is often assumed for dynamically moving objects in an enclosed region (see for instance [95, 104]).

When $\rho(A,r)$ represents the density of agent at location r for a certain area A, then $n(R,r)$ is estimated as in Equation (2.56) where $A(R)$ represents the communication area defined by the communication range R as shown in Figure 2.16.

$$n(R,r) = \int_{A(R)} \rho(A,r) dr \tag{2.56}$$

The unit area dA as illustrated in Figure 2.16 depends on R (and also a silent variable, e.g., x) as

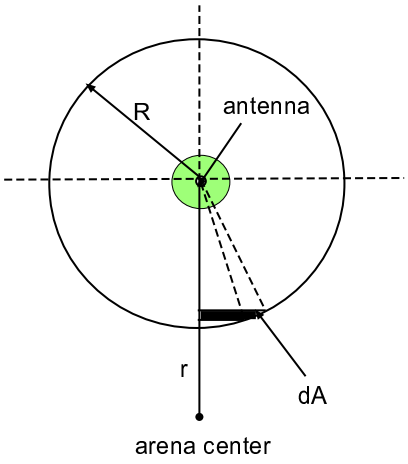

Figure 2.16: Cylindrical coordinates.

follows:

$$dA = x \tan\left(\cos^{-1}(x/R)\right) dx \tag{2.57}$$

We propose a practical method for estimating $\rho(r)$ based on the unique information provided about the system, i.e., $f(r)$, as follows. By definition, the density of mobile sensors in the environment as a function of their distance to the center of the arena is given by the following equation:

$$\rho(r) = \lim_{\delta \to 0} \frac{\int_{r-\delta}^{r+\delta} f(x)dx}{\pi(r+\delta)^2 - \pi(r-\delta)^2} \tag{2.58}$$

However, the mathematical expression of $f(r)$ often does not allow for a closed form in the calculation of Equation (2.58), therefore, an approximation of $\rho(r)$ can be used if necessary, as presented in Equation (2.59) with 2δ, the thickness of an infinitely small ring of radius r surrounding the center of the arena as shown in Figure 2.17.

$$\rho(r) \approx \frac{1}{4\pi\delta r} \int_{r-\delta}^{r+\delta} f(x)dx \tag{2.59}$$

Note that for any distribution function $f(\cdot)$ the condition stated in Equation (2.60) is respected, guaranteeing that all the mobile sensors are accounted for in the approximation.

$$\int_0^{2\pi} \int_0^{d_0} \rho(r) r dr d\theta = 1 \tag{2.60}$$

It is also worth noting that here the influence of embodiment is implicitly contained within the spatial distribution of the agents. For instance, for the same detection range and movement pattern, the minimum distance possible between the centers of two agents would increase as the size of the agents increases. Moreover, when the communication range is much larger than the agent

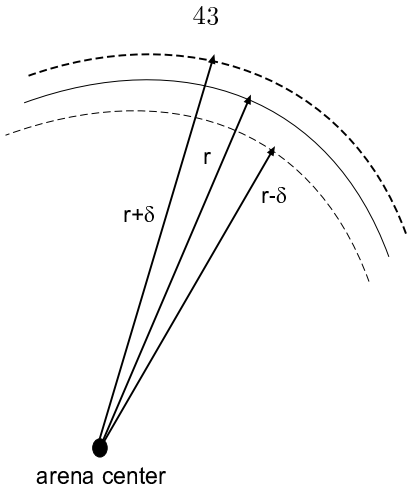

Figure 2.17: Cylindrical coordinates: spatial density location.

detection range, information exchange maneuver will happen much more frequently than collision avoidance resulting from physical interference. Finally, if the controller of the agents is designed such that agents move away from each other after exchanging information (e.g., a U-turn to decrease the likelihood of consecutive information exchange with the same agents), then pure agent-agent collision avoidance will happen rarely.

B. The experiment

Figure 2.18 left shows a simulated, mobile, autonomous agent endowed with local peer-to-peer transmission capabilities (i.e., a broadcast communication antenna) in Webots and Figure 2.18 right presents a network consisting of 10 agents.

Figure 2.15 is a sketching of two agents within communication range. Information exchange occurs between the mobile agents upon encounter. However, due to communication interference, an agent can receive messages only from a maximum number of other agents at a time (referred to as the agent's information reception capacity, C). For instance an information reception capacity 1 corresponds to the case where all transmissions received from more than one agent at a time are always lost (due to communication interference, signal jamming, low computational power of individual nodes, etc.)

C. Results

For this experiment, a team of 10 mobile agents was used in 100 x 100 cm^2 and 142 x 142 cm^2 arenas, and 50 Webots simulations runs were conducted. Error bars represent standard errors among runs. In the methodology detailed in the mathematical analysis above, the first step is to estimate the distribution of the agents in the environment. We obtained that distribution by computing the mean fraction of the team in a set of partitioning regions of the environment over time. These regions

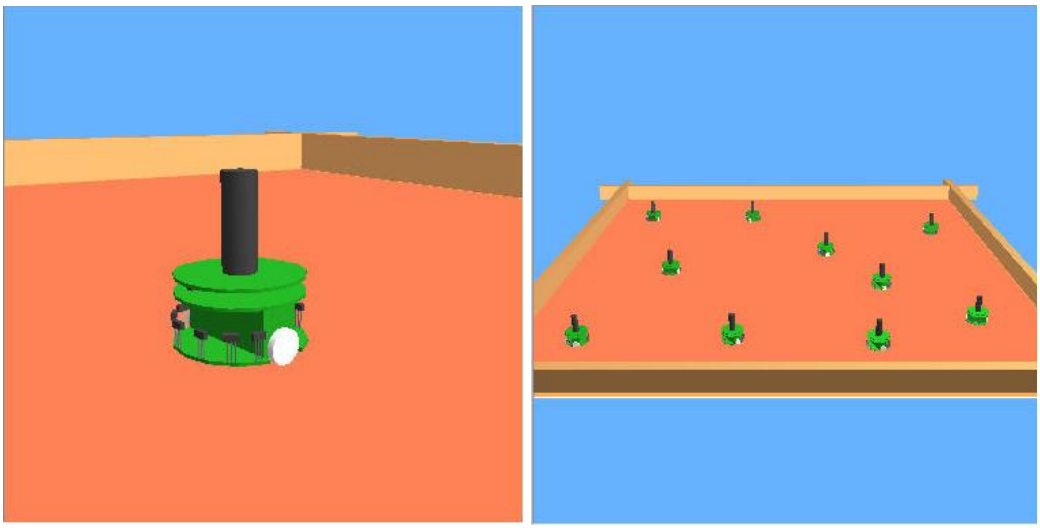

Figure 2.18: Left: Close-up of a simulated mobile sensor equipped with a broadcast communication antenna mounted on top at its center. Right: A network of 10 simulated mobile sensors in an enclosed 142 x 142 cm^2 region.

correspond to a set of concentric rings of width 9 cm centered at the middle of the arena. The width, (i.e., 9 cm) was chosen (arbitrarily) because an agent (i.e., a Khepera robot) is 5.5 cm in diameter and travels at an average speed of 8 cm/s, thus, capable of moving from one region to another, on average, every second. We chose a width of 9 cm (instead 8 cm for instance) to leave a small margin for the 10% noise embedded with the outputs of the actuators (e.g., motors) of each simulated robot (see [74]). Consecutive snapshots of the environment were taken every second for one hour and repeated 50 times.

The histograms in Figures 2.19 and 2.20 show the distribution of agents (in fraction of team size) in the different spatial regions mentioned above for teams of 1 and 10 agents in 100 x 100 and 142 x 142 cm^2 arenas. The similarity between the graphs obtained with 1 agent and that obtained with 10 agents confirms that both graphs also represent the average probability of finding a particular agent (characterized by this specific movement pattern and embodiment) within the spatial regions defined above at any given time. In each of these plots, the solid-line curve represents the analytical function approximation of the measured spatial distribution of the agents in the arena. The distribution of the teammates appears to follow a Gaussian-like pattern (although it seems slightly asymmetric) as illustrated in each plot by the good agreement between the Gaussian approximation and the measured distribution. Note that this Gaussian approximation corresponds to the $f(r)$ mentioned in subsection A and the plots in Figures 2.19 and 2.20 do not represent the probability distribution function itself but the integrated values of the function over consecutive intervals of width 9 cm (i.e., the concentric rings).

The Gaussian approximation function used to fit the measurements in the 100 x 100 cm^2 arena

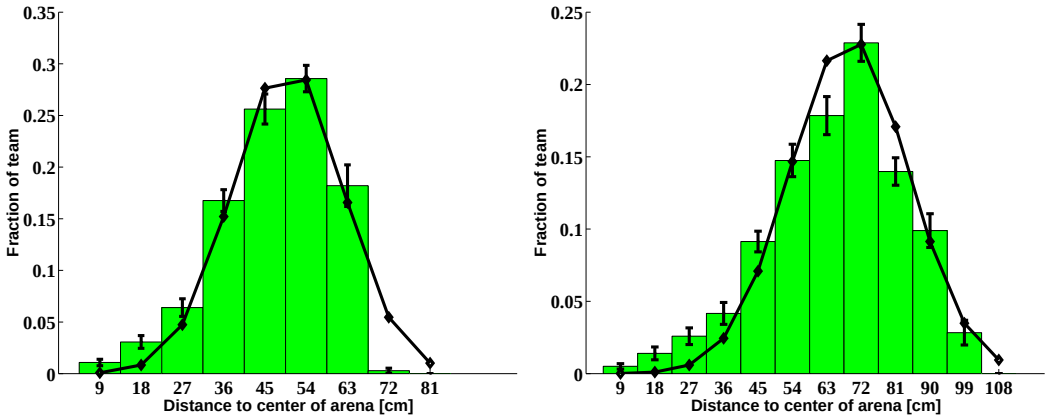

Figure 2.19: Spatial location of 1 agent over time, solid-line curve represents a Gaussian approximation fitting the measurement. Left: 100 x 100 cm^2 arena, Right: 142 x 142 cm^2 arena.

has parameters: $\mu = 50$ cm and $\sigma = 13$ (see Figures 2.19 left and 2.20 left). That used for the 142 x 142 cm^2 arena has parameters: $\mu = 71$ cm and $\sigma = 17$ (see Figures 2.19 right and 2.20 right).

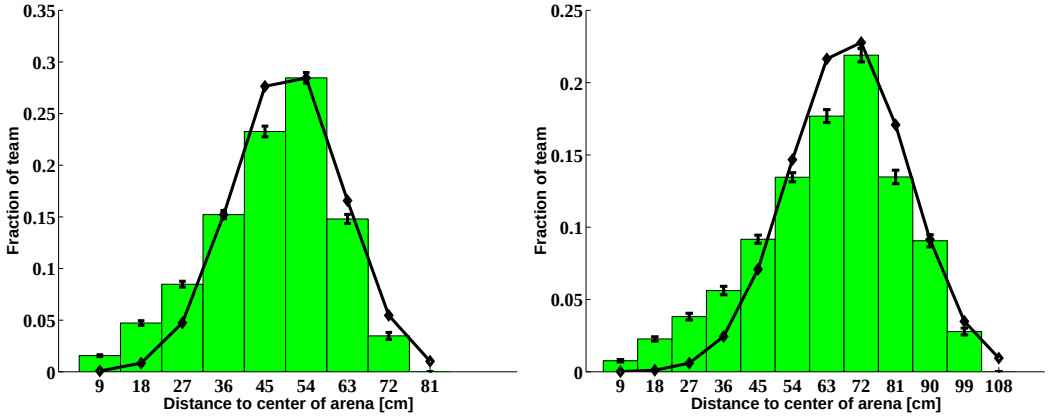

Figure 2.20: Spatial location of 10 agents over time, solid-line curve represents a Gaussian approximation fitting the measurement. Left: 100 x 100 cm^2 arena. Right: 142 x 142 cm^2 arena.

Once the distribution function was estimated, we proceeded using the method detailed in subsection A to find the communication range that maximizes the average probability of successful transmission, then we compared the prediction with the measured communication range that maximizes the number of messages successfully received per agent par time unit for a team of 10 mobile robots and different values of C. Here, the number of messages successfully received per agent per minute provides a simple way of assessing the information exchange rate in the network.

Figures 2.21 left and 2.22 left present the predictions of the average probability of successful information exchange as a function the communication range in 100 x 100 cm^2 and 142 x 142

cm^2 arenas, respectively. Figures 2.21 right and 2.22 right show the measured number of messages successfully received per agent per minute in the same networks, respectively.

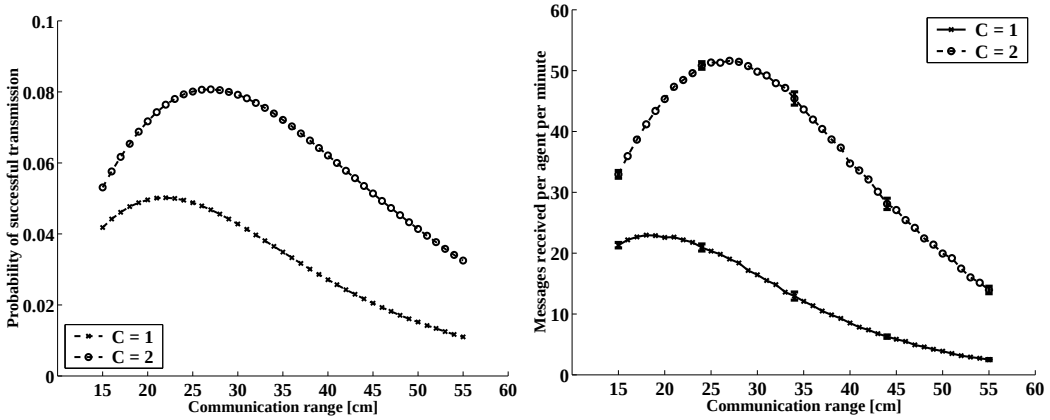

Figure 2.21: Network of 10 agents in a 100 x 100 cm^2 arena. Left: Prediction of the probability of successful transmission obtained using $f(r) = \frac{1}{\sigma\sqrt{2\pi}} \exp(\frac{-(r-\mu)^2}{2\sigma^2})$ with $\mu = 50$ cm and $\sigma = 13$. Right: Simulation results of the number of messages successfully received per agent per minute.

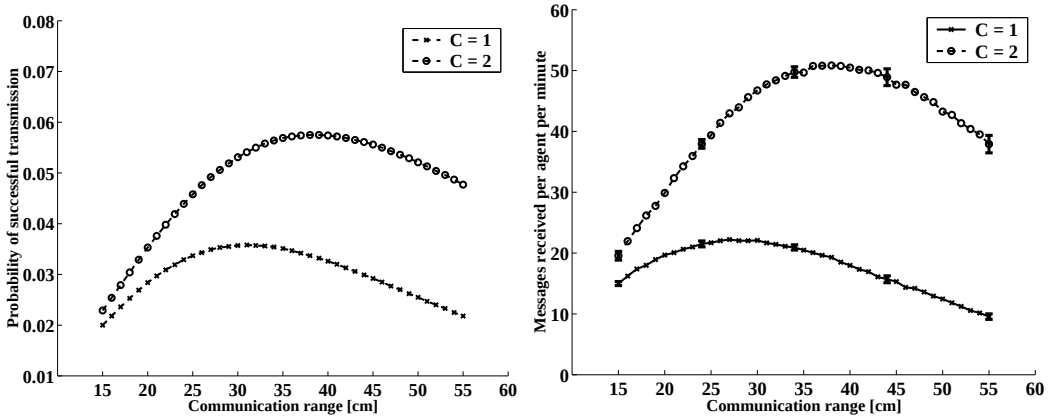

Figure 2.22: Network of 10 agents in a 142 x 142 cm^2 arena. Left: Prediction of successful transmission obtained using $f(r) = \frac{1}{\sigma\sqrt{2\pi}} \exp(\frac{-(r-\mu)^2}{2\sigma^2})$ with $\mu = 71$ cm and $\sigma = 17$. Right: Simulation measurements of the number of messages successfully received per agent per minute.

Figures 2.21 and 2.22 clearly show that the proposed methodology accurately predicts the optimal communication range for maximizing the information exchange rate in the different networks presented above. Nevertheless, to assess the influence of the spatial distribution of the agents on the prediction, we compared the predictions of a uniform spatial distribution assumption to those of the heterogeneous distribution (i.e., the Gaussian approximation) and the measurement results presented above. The uniform distribution assumes that agents have the same probability of being at any location on the plane at any given time (for instance, $f(r) = \frac{1}{d_0}$ for a circular arena of radius

d_0, thus, we approximated the square arena of size L as a circle of radius $\frac{L}{\sqrt{\pi}}$). Figure 2.23 presents the predicted average probability of successful transmission in a network of 10 mobile agents in arenas of sizes 100 x 100 cm^2 and 142 x 142 cm^2, respectively for a 2-D isotropic distribution.

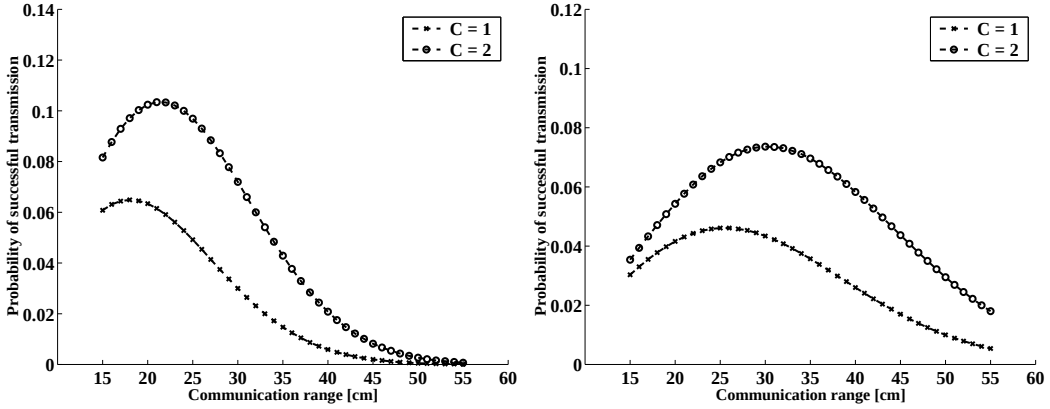

Figure 2.23: Prediction with assumption of uniform distribution, 10 mobile agents. Left: Predicted probability of successful transmission in a 100 x 100 cm^2 arena. Right: Predicted probability of successful transmission in a 142 x 142 cm^2 arena.

First it is worth noting the qualitative similarity between the model prediction and the simulation curves in Figures 2.21 and 2.22 whereas the curves in Figure 2.23 present a different shape. Moreover, there is a better quantitative agreement between the predictions of the Gaussian distribution and the simulation results than between those of the uniform distribution assumption and the same simulation results. For instance in the 142 x 142 cm^2 arena, from Figure 2.23 right it appears that the collaboration rate for a communication range of 55 cm is worse than that of a communication range of 15 cm, whereas the simulation reveals the opposite and rather agrees with the prediction of Figure 2.22 left.

Table 2.1: Performance summary

	R_{opt}		Prediction Error	
Result	Arena100	Arena142	Arena100	Arena142
Gaussian distribution				
$C = 1$	21 cm	30 cm	10.5%	7.1%
$C = 2$	28 cm	38 cm	3.7%	2.7%
Isotropic distribution				
$C = 1$	18 cm	25 cm	5.2%	22.2%
$C = 2$	21 cm	30 cm	10.7%	18.9%
Simulation				
$C = 1$	**19 cm**	**28 cm**	—	—
$C = 2$	**27 cm**	**37 cm**	—	—

In Table 2.1 arenas called Arena100 and Arena142 stand for the 100 x 100 cm^2 and 142 x 142

cm^2 arenas, respectively and R_{opt} is the optimal communication range as predicted using the mathematical model or measured with simulation. Table 2.1 shows that while all predicted optimal communication ranges obtained using the heterogeneous distribution function are within 2 cm of the measurement, the predictions obtained using the arbitrary uniform distribution are off by as much as 22.2% related to the measured optimal communication ranges in both the 100 x 100 cm^2 and 142 x 142 cm^2 arenas for information reception capacities of 1 and 2, respectively.

Finally, it appears that the distribution function resulting from the movement pattern of the agents (and the boundary conditions) plays a major role in predicting the optimal communication range as the more closely it approximates the distribution of the agents in the environment, the more accurate the predictions are.

2.7 Conclusion

We have presented a set of mathematical models, inspired by studies in biology, ecology, and robotics, for describing and analyzing event discovery and information dissemination in swarm-based mobile sensor networks. Although most of the analysis in this chapter was illustrated with experiments using networks of relatively small sizes, the scalability of the proposed models guarantees that they can be applied to networks consisting of hundreds or thousands of mobile sensing agents. Event discovery or information spread in such networks have similar characteristics (e.g., occurrence, dissemination) as biological phenomena of geographical spread (e.g., infectious diseases or gene spread in human populations) and foraging and scouting in insect colonies. We have shown that both can be described using similar mathematical equations.

We have proposed two main categories of models: non-embodied and embodied models. Each category of models consists of two subsets: spatial and nonspatial models. In the first category of models we describe the agents as mobile bodiless points, while in the second category we take into account the physical interference between agents resulting from embodiment. In the nonspatial model we assume that agents hop around randomly and occupy independent positions over consecutive time steps. The spatial models often keep track of the trajectory of each agent, the correlation between the positions occupied by the agents over consecutive time steps or make uses of the spatial distribution resulting from the movement pattern of the agents.

We have proposed simple methods for calculating the parameters of the proposed models using the (geometrical) characteristics of the environment (e.g., radius of the arena, and those of the agents (e.g., body size, communication range, sensing range.)

While the analysis of non-embodied models was mostly theoretical, in the analysis of the embodied models we have presented a case study in which we have introduced a mathematical design method for maximizing the information exchange rate in a fully *ad hoc* network of mobile sen-

sors. This analytical design method makes use solely of the distribution of the agents in the region over time to accurately predict the communication range that maximizes the successful information exchange rate between the teammates. We have validated the proposed models by comparing the theoretical predictions with results obtained using a 3-D sensor-based simulator of real mobile robots and have illustrated the influence of the spatial distribution function by showing that the better the function agrees with the agent spatial distribution over time, the more accurate the predictions are.

Chapter 3

Modeling Non-Collaborative Distributed Manipulation in Swarm Robotic Systems

3.1 Introduction

In this chapter, we present a macroscopic model that captures the dynamics of a non-collaborative, distributed manipulation experiment and a suite of threshold-based, scalable, and fully distributed algorithms for allocating the workers to the task whose demand evolves dynamically over time. In the allocation algorithm, individual robots estimate the availability of work based solely on their local perceptions. We compare the predictions of the model with those delivered by Webots (see 1.3.4 for a description of Webots). Instead of computing trajectories and sensory information like an embodied simulator would, the model represents the collective dynamics of the system as a set of differential-difference equations whose parameters are obtained using simple geometrical considerations and systematic tests with one or two embodied agents. The proposed model has been built up incrementally, the matching between model and embodied simulations verified at each step. This in turn allows us to clearly define how to calculate the different parameters of the model and avoid the introduction of any free parameter into the model. We support the discussion with a concrete case study concerned with the gathering and clustering of small objects initially scattered in an enclosed arena. Experiments were carried out with teams consisting of one to ten individuals in either or both the macroscopic model and embodied simulations. Results show that teams using a number of active workers dynamically controlled by the allocation algorithm achieve similar or better performances in aggregation than those characterized by a constant team size while using, on average, fewer robots over the whole aggregation process. Moreover, like the models presented in the previous chapter, results show that the macroscopic model proposed here represents a useful tool for generalizing the dynamics of highly stochastic and nonlinear robotic systems, delivering both qualitatively and quantitatively correct predictions in time lapses that are

at least four orders of magnitude smaller than those required by embodied simulations. Finally, we conclude the chapter by presenting a study of four threshold-based allocation algorithms applied to the aggregation experiment, comparing their efficiency and robustness in the presence of static and dynamic perturbations and by proposing a few suggestions for future work.

The Case Study

The case study presented in this chapter belongs to the category of experiments in which everything is based on local interaction and local communication strictly following the swarm intelligence principles [11, 44, 55]. The experiment presented here is the follow-up of experiments conducted with real robots presented by Martinoli *et al.* [62, 63, 64] and represent a combination of sensing (object search and discovery) and more importantly manipulation (objects are displaced and gathered in clusters). The task is to collect small objects referred to as *seeds* in a square arena and to gather them in a single cluster[1] using *Khepera* robots [78] equipped with grippers and capable of distinguishing the seeds with their frontal sensors. The robots have local sensing capabilities only, so their encounters with seeds and clusters of seeds happen randomly. Moreover, as the robots do not have a global perception of the environment, they do not know when the task is finished, motivating the need for a distributed task allocation mechanism (because of the swarm-based approach) to allow each individual to stop working based on its own estimate of the availability of work.

Division of Labor

We propose a swarm-based, scalable, and fully distributed mechanism that allows an efficient worker allocation for obtaining improved clustering results without the need for a centralized controller. More generally, we are interested in understanding task allocation and labor division mechanisms exploited in social insect societies that are suitable for artificial embedded systems such as multiple mobile robot platforms. Recently, several hypotheses have been proposed to explain these mechanisms in natural colonies: some of them based on threshold responses [15, 96], others focus only on task-switching probabilities [82]. However, to the best of our knowledge, no attempt has been made so far to link both the macroscopic (i.e., group behavior) and microscopic (i.e., individual behavior) descriptions of the system. For instance, how workers in such systems gather the information necessary to decide whether or not to perform a particular task or to switch task, has not been investigated so far.

In this chapter, we propose a mathematical model that describes accurately the dynamics and the quantitative characterization of the system and we present a series of threshold-based, scalable, distributed worker allocation algorithms applied to an aggregation experiment conducted with a

[1] Although the experiment is not intended to reproduce a biological system, it presents several similarities with the nest cleaning and dead ants clustering performed by some ant colonies [20].

homogeneous group of autonomous robots. In these allocation algorithms, whether the individuals are all characterized by the same threshold or not, since the agents do not perceive the demand globally but rather estimate it locally, they do not work or rest all at the same time, a behavior that would arise if the demand was broadcast from an external supervisor.

In section 3.4 we first introduce and analyze two simple systems in which the robots wander, looking for seeds, and avoid obstacles. Section 3.5 presents the full system without task allocation and its mathematical description, and in section 3.6, we illustrate the efficiency and benefits of task allocation with both mathematical predictions and experimental results. In section 3.7 we discuss the strengths and limitations of the worker allocation algorithm. Concluding remarks are presented in section 3.8.

3.2 The Aggregation Experiment

The experiment described in this chapter and used for assessing the efficiency of the worker allocation algorithm is concerned with the gathering and clustering of small objects scattered around an enclosed arena. We refer to these small objects as seeds. In most of the work done so far, e.g., [44], and more specifically in [63, 64], the size of the working team was kept constant during the whole aggregation process. Those experiments define the baseline for efficiency comparison with and without the worker allocation algorithm. In the following, we use three primary team performance measurements: the *average cluster size*, the *average number of clusters*, and the *average number of active workers* in the environment.

3.2.1 The sensor-based simulator and the experimental setup

In [63], the experiment was carried out in a square arena (of size 80 x 80 cm^2) delimited by a wall. 20 seeds (cylindrical wooden objects of diameter 1.7 cm and height 2.5 cm) were randomly distributed around the working zone within the arena. Groups of one to five Khepera robots equipped with gripper turrets were used to collect and cluster the seeds. Because of their thinness, the seeds can be distinguished from the wall and other robots using the Khepera's six frontal infrared proximity sensors. In the experiments presented here, the arena consists of an inner area (or working zone) where the cluster formation takes place and a surrounding area (or parking zone) where the robots that decide to stop working following the distributed worker allocation mechanism go and stay idle to save energy. Finally, we implemented the experiment in Webots. The mean speed ratio for this experiment with five robots between Webots and real time is about 10 on a Pentium III, 933 MHz machine.

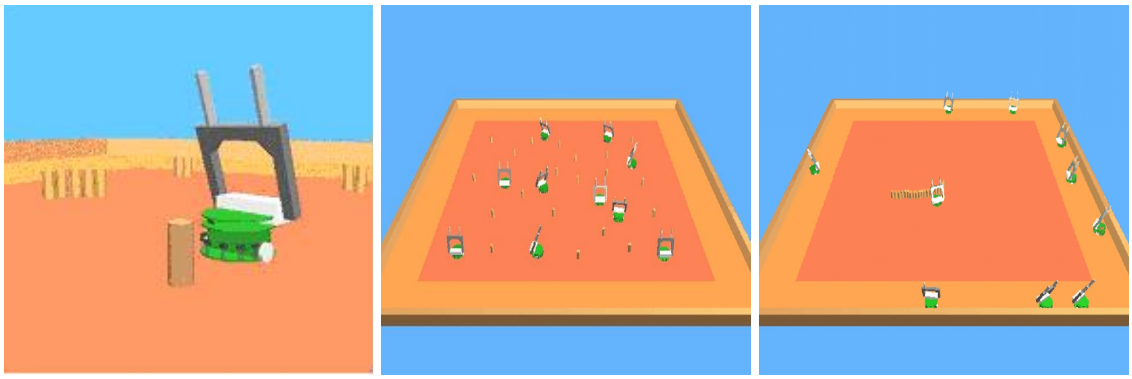

Figure 3.1: Left: close-up of a simulated robot equipped with a gripper and standing in front of a seed. Center: setup in the embodied simulator (10 robots, 20 seeds, 178 x 178 cm^2 arena). Inner area represents the working zone, surrounding area is the parking/resting zone where robots that decide to stop working stay idle. Right: typical end of aggregation experiment, e.g., 5 hours of simulated time, in a 178 x 178 cm^2 arena.

3.2.2 The robot's controller

The behavior of a robot is determined by a simple hand-coded program that can be represented with a standard flow chart or a Finite State Machine (FSM), as depicted in Figure 3.2. Note that the FSM representation in Figure 3.2 does not include the distributed worker allocation mechanism introduced and detailed in section 3.6. The behavioral granularity shown in Figure 3.2 is chosen by the experimenter such that the FSM captures all the details of interest.

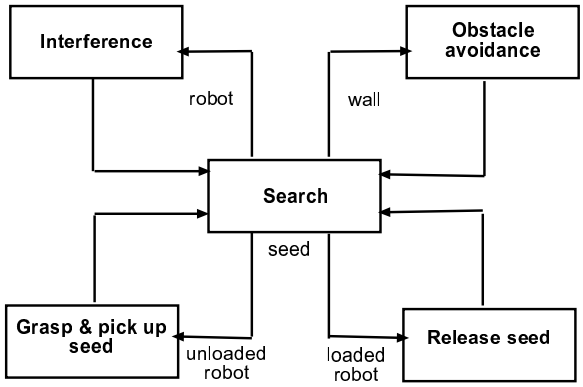

Figure 3.2: FSM representing the robot's controller. Transitions between states are triggered by sensory measurements.

Without considering the mode-switching behavior (i.e., the worker allocation mechanism explained in detail in section 3.6), we can summarize each agent's behavior with the following simple rules. In its default behavior the agent moves straight forwards within the working zone looking for seeds. When at least one of its six frontal proximity sensors is activated, the agent starts a

distinguishing procedure. Two cases can occur: if the agent is in front of a large object (a wall, another agent, or the body side of a cluster of seeds), the object is considered to be an obstacle and the agent avoids it. In the second case, the small object is considered to be a seed. If the agent is not already carrying a seed, it grasps this one with the gripper, otherwise it drops the seed it is carrying close to the one it has found; then in both cases, the agent resumes searching for seeds. With this simple individual behavior, the team is able to gather the objects in clusters of increasing size. A cluster is defined as a group of seeds whose neighboring elements are separated by at most one seed diameter. Note that agents can identify only the two ends of a cluster as seeds (as opposed to obstacles), as a consequence, clusters are built in line.

3.3 Mathematical Modeling

The mean speed ratio for this experiment with five robots between the macroscopic model (implemented in C language) and Webots simulations is about 5000 on a Pentium III, 933 MHz machine. This comparison is based on the time it takes to obtain a macroscopic model prediction of an experiment and to conduct a single embodied simulation, although several runs of the latter are needed to obtain statistically significant information. For instance, for each experimental result presented in this paper, 30 Webots simulation runs were carried out, corresponding to a speed ratio of 150000 between the macroscopic model and the embodied simulator.

Calculating model parameters

In each of the mathematical models presented below, we use a set of parameters to link the models to the embodied simulations. The main parameters used are γ_w, γ_r, γ_k^{dec}, and γ_k^{inc}, the rate at which a robot encounters the wall, another robot, and a cluster of size k and decreases or increases the size of that cluster, respectively. Consistent with previous publications [2, 48, 63, 64, 66, 67], we compute these transition rates (and the corresponding probabilities) from one state to another using simple geometrical considerations and the robots' sensing and interaction capabilities (e.g., detection areas, approaching perimeters) as follows. We used the method proposed in subsection 2.5.3 to calculate the rate γ_i at which the robot will detect any object of type i present in the arena. Note that the robot's motion can be considered as a correlated random walk whose direction change occurs only when the robot encounters an obstacle. Therefore, in reality the rate at which the robot encounters the surrounding wall depends mostly on its speed and sensory capability (wall detection range). The robot moves straight forwards in the arena until it encounters an obstacle, therefore our simple way of computing γ_w is by assuming that, on average, the robot encounters the wall every time after traveling a certain distance. The average distance we consider is the side of the square arena, S, minus twice the wall detection range, R_w, mathematically, $\gamma_w = v/(S - 2R_w)$.

The numerical values used for all the geometrical parameters defined above are measured in systematic tests with one or two embodied agents. Furthermore, the current implementation of the model does not allow for overlapping areas of detection between objects of different types (i.e., robot, wall, and clusters of seeds). Equation (3.1) expresses this limit mathematically, where M is the maximum number of seeds in the arena.

$$\gamma_w + \gamma_i(t) + \sum_{k=1}^{M} \max(\gamma_k^{dec}(t), \gamma_k^{inc}(t)) \leq 1, \quad \text{at any time step} \quad t \tag{3.1}$$

This in turn sets an upper bound for the total extended area of the group of robots as the modeling methodology reaches its limitations in overcrowded arenas scenarios, i.e., scenarios where the sum of all extended areas exceeds the total area of the arena.

The numerical values used for the robot speed, the approaching angle for decreasing or increasing the size of a cluster, and the different detection radii of the objects are summarized in Table 3.1 and Table 3.2. In Table 3.1 the construction (resp. destruction) angle corresponds to the angle of the access perimeter for incrementing (resp. decrementing) a cluster of seeds. Table 3.2 summarizes the durations of the robot's different maneuvers. All the parameters presented in these tables were obtained using systematic simulation experiments using one or two embodied agents .

Table 3.1: Parameters used in the models of the aggregation experiment. Detection distances are measured from the center of the robot to the center of the object detected.

Mean robot speed v[cm/s]	Cluster construction angle $\alpha_{inc}(k)$ [degrees]	Cluster destruction angle $\alpha_{dec}(k)$ [degrees]	Wall detection distance R_w[cm]	Seed radius r_s[cm]	Seed detection distance R_s[cm]	Robot detection distance R_r[cm]	Arena size [cm^2]
8	65 if $k > 1$ 180 if $k = 1$	60 if $k > 1$ 180 if $k = 1$	6	0.85	6.4	10	80 x 80 114 x 114

Table 3.2: Durations of the different robot maneuvers described in subsection 3.2.2.

Seed dropping (or grasping) τ_d (or τ_g) [seconds]	Obstacle avoidance τ_w [seconds]	Interference τ_i [seconds]
10	1	2

3.4 Examples of Simplified Models

Before describing the implementation details and results of the full aggregation system, we introduce two examples of simplified scenarios of Probabilistic Finite State Machines (PFSMs). Both examples

represent key sub-chains of the full aggregation system's Markov chain that we will describe in sections 3.5 and 3.6. It is worth noting that the sets of differential-difference equations ($DDEs$) describing these two simplified systems at the macroscopic level are linear and contain time-delays, but the model of the full aggregation system at the macroscopic level will be a set of nonlinear, coupled DDEs affected by time-delays due to the (time-dependent) modification of the environment, e.g., seeds are moved and placed in different clusters whose sizes therefore vary over time.

Unless otherwise stated, simulations using Webots have been repeated 50 times and error bars represent the standard deviations among the runs. At the macroscopic level, of course, one run suffices, since only the mean swarm performance is predicted.

3.4.1 Search and obstacle avoidance

In general, mobile robotics experiments have one characteristic in common which is the existence of delay states. A delay state simply represents a behavior which the robot may perform for a certain duration τ. In this first simplified scenario, the robot's default state is the searching (or wandering) state, while in this state it may encounter the wall and another robot at the rates γ_w and γ_i, respectively. Upon encounter with the wall (resp. another robot) the robot enters the wall avoidance state (resp. interference state), staying in that state for a duration τ_w (resp. τ_i). Here, the rates at which the robot enters or leaves those states are independent of the robot's interaction with the environment, i.e., they are fixed. As described in subsection 3.2.2, since the default behavioral state in the aggregation experiment is the search state, a first key sub-chain of the system is represented by the search state coupled with an obstacle avoidance state and an interference state. Figure 3.3 represents graphically the state diagram of this simple PFSM. Note that the numerical values used in this example (and presented in Figure 3.3) were derived from the values of Table 3.1 and Table 3.2 using parameters for a generic obstacle: $\tau_w = 1$ second, $\tau_i = 2$ seconds, and γ_w and γ_i are functions of the setup (for example, $\gamma_w = 0.158$ and $\gamma_i = 0.157$ for the wall and robot encounter rates in an arena of 80 x 80 cm^2 with a group of 10 robots). w_s, w_i, and w_w represent the numbers of workers in the search, interference, and obstacle (i.e., wall) avoidance states, respectively.

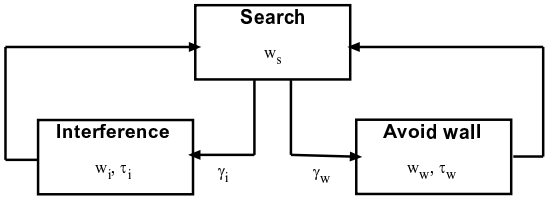

Figure 3.3: Simple sub-chain consisting of a search, a robot-robot interference state, and an obstacle avoidance state.

The state variables of the macroscopic model are represented by non-discrete quantities (i.e., real numbers) thus, the PFSM of Figure 3.3 representing the whole swarm can be described by the following system of DDEs:

$$\frac{dw_s}{dt} = -(\gamma_w + \gamma_i)w_s(t) + \gamma_w w_s(t-\tau_w)\Theta(t-\tau_w) + \gamma_i w_s(t-\tau_i)\Theta(t-\tau_i) \tag{3.2}$$

$$\frac{dw_w}{dt} = \gamma_w w_s(t) - \gamma_w w_s(t-\tau_w)\Theta(t-\tau_w) \tag{3.3}$$

$$w_i(t) = W_0 - w_s(t) - w_w(t) \tag{3.4}$$

In Equations (3.2)-(3.4), $w_s(t)$, $w_i(t)$, and $w_w(t)$ represent the number of robots in search, interference, and wall avoidance states at time t, respectively. W_0 is the total number of worker robots in the arena. Equation (3.2) states that the average number of robots in the search state during time step t decreases (resp. increases) when some robots encounter a wall or other robots (resp. when some robots finish interfering with each other or finish avoiding the wall). The step function $\Theta(t-\tau)$ is the same as that introduced in subsection 2.6.1 (i.e., equal to 0 when $t < \tau$ and 1 otherwise). Equation (3.4) simply exploits the conservation of the total number of robots for calculating the average number of robots in the interference state. It is worth mentioning that, unless otherwise stated, we assume that no robots exist when $t < 0$. Mathematically speaking, we have

$$w_s(t) = w_w(t) = w_i(t) = 0 \quad \text{if} \quad t < 0 \tag{3.5}$$

A. Steady-state analysis

The DDE system above is linear therefore we can analyze its steady-state using the Laplace transform. Equation (3.2) can be transformed using the right shift and left shift theorems as follows ($W_s(s)$ representing the Laplace transform of $w_s(t)$):

$$sW_s(s) - W_s(0) = -(\gamma_w + \gamma_i)W_s(s) + \gamma_w W_s(s)e^{-\tau_w s} + \gamma_i W_s(s)e^{-\tau_i s} \tag{3.6}$$

Solving for $W_s(s)$ and applying the limit theorem brings

$$w_s^* = \lim_{s\to 0} sW_s(s) = \frac{W_0}{1 + \gamma_w\tau_w + \gamma_i\tau_i} \tag{3.7}$$

Similarly, Equations (3.8) and (3.9) follow as

$$w_w^* = \lim_{s\to 0} sW_w(s) = \frac{\gamma_w\tau_w W_0}{1 + \gamma_w\tau_w + \gamma_i\tau_i} \tag{3.8}$$

$$w_i^* = \lim_{s\to 0} sW_i(s) = \frac{\gamma_i\tau_i W_0}{1 + \gamma_w\tau_w + \gamma_i\tau_i} \tag{3.9}$$

B. Model and simulation results

Figure 3.4 shows a comparison between the mathematical model predictions and embodied simulations in the steady-state for different team sizes. In Figure 3.4 right, the height of the embodied simulations histogram represents the mean value of w_s^*/W_0 over a 30-minute time window and over 50 runs. The error bars represent a mean standard deviation calculated as an average of the standard deviations measured in each run over the 30-minute time window.

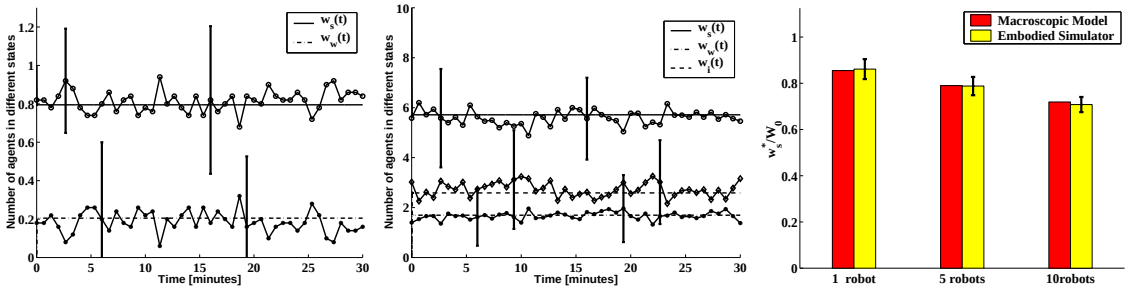

Figure 3.4: Comparing macroscopic model predictions and embodied simulations results for the simple PFSM depicted in Figure 3.3. Left: $w_s(t)$ and $w_w(t)$ (1 robot, 80 x 80 cm^2 arena). Center: $w_s(t)$, $w_i(t)$, and $w_w(t)$ (10 robots, 80 x 80 cm^2 arena). Right: Steady-state value of w_s^*/W_0.

Figure 3.5 shows graphically how the average number of searching robots in the steady-state is influenced by the time delays τ_i and τ_w.

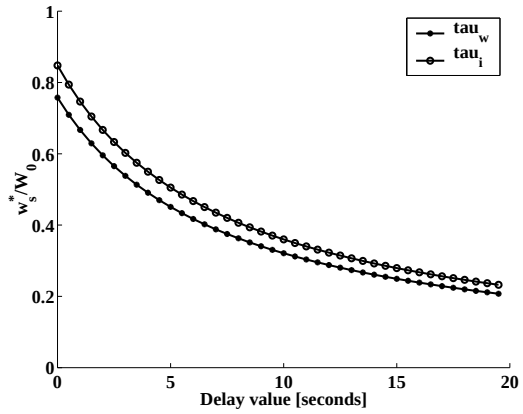

Figure 3.5: Graphical representation of Equation (3.7) for different τ_i and τ_w delay durations in an 80 x 80 cm arena.

3.4.2 Search, seeds detection, and obstacle avoidance

In this section we analyze at first hand the robot-seed interaction and measure the agreement between the mathematical model predictions and the embodied simulations, in particular for the seed detection and identification by the team of robots. In this second simplified scenario, the robots

search for seeds, avoid the wall (and each other) upon encounter, and when they detect a seed, they identify it following the distinguishing procedure detailed in section 3.2.2 but do not proceed with picking up the seed.

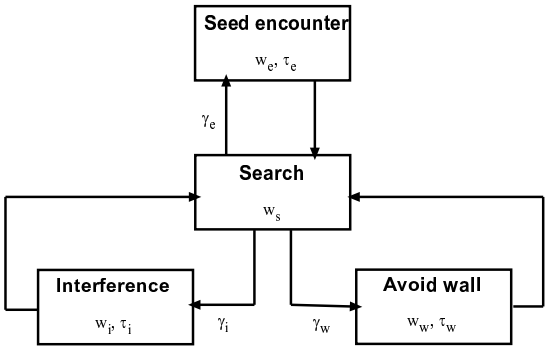

Figure 3.6: Sub-chain representing a simple detection and identification interaction of the robots with the seeds in the aggregation experiment.

The following system of DDEs represents the macroscopic model for this sub-chain (here we present only the dynamics of the numbers of agents in search and seed encounter states, those of the agents in wall avoidance and interference states are obtained much like in the previous example):

$$\begin{aligned}\frac{dw_s}{dt} &= -(\gamma_w + \gamma_i + \gamma_e)w_s(t) + \gamma_w w_s(t-\tau_w)\Theta(t-\tau_w) \\ &\quad + \gamma_i w_s(t-\tau_i)\Theta(t-\tau_i) + \gamma_e w_s(t-\tau_e)\Theta(t-\tau_e)\end{aligned} \tag{3.10}$$

$$\frac{dw_e}{dt} = \gamma_e w_s(t) - \gamma_e w_s(t-\tau_e)\Theta(t-\tau_e) \tag{3.11}$$

In Equations (3.10) and (3.11), γ_e and τ_e represent the rate at which a robot encounters a seed and the amount of time the robot needs to fully identify the object as a seed following the distinguishing procedure of section 3.2.2, respectively. The measured value of τ_e is half that of τ_g given in Table 3.2 (i.e., $\tau_e = 5$ seconds) because here the robot only rotates in front of the small object to fully identify it as a seed and does not proceed with lowering its gripper and grasping the seed, a maneuver that lasts additional 5 seconds. The additional term in Equation (3.10), related to Equation (3.2), translates the fraction of robots leaving the search state to examine a small object they have just found and those returning to the search state after fully identifying a small object as a seed.

A. Steady-state analysis

We use the same methodology as in subsection 3.4.1, i.e., Laplace transform and limit theorem, for finding the steady-state values of the dynamic variables of the simplified model scenario represented

by Figure 3.6.

$$sW_s(s) - W_s(0) = -(\gamma_w + \gamma_i + \gamma_e)W_s(s) + \gamma_w W_s(s)e^{-\tau_w s} + \gamma_i W_s(s)e^{-\tau_i s} + \gamma_e W_s(s)e^{-\tau_e s} \quad (3.12)$$

It follows that

$$w_s^* = \lim_{s\to 0} sW_s(s) = \frac{W_0}{1 + \gamma_w\tau_w + \gamma_i\tau_i + \gamma_e\tau_e} \quad (3.13)$$

$$w_w^* = \lim_{s\to 0} sW_w(s) = \frac{\gamma_w\tau_w W_0}{1 + \gamma_w\tau_w + \gamma_i\tau_i + \gamma_e\tau_e} \quad (3.14)$$

$$w_i^* = \lim_{s\to 0} sW_i(s) = \frac{\gamma_i\tau_i W_0}{1 + \gamma_w\tau_w + \gamma_i\tau_i + \gamma_e\tau_e} \quad (3.15)$$

$$w_e^* = \lim_{s\to 0} sW_e(s) = \frac{\gamma_e\tau_e W_0}{1 + \gamma_w\tau_w + \gamma_i\tau_i + \gamma_e\tau_e} \quad (3.16)$$

B. Model and simulation results

Figure 3.7 shows a comparison between macroscopic model predictions and the embodied simulations of the robots' interaction with the seeds for different arena and swarm sizes. We notice a good

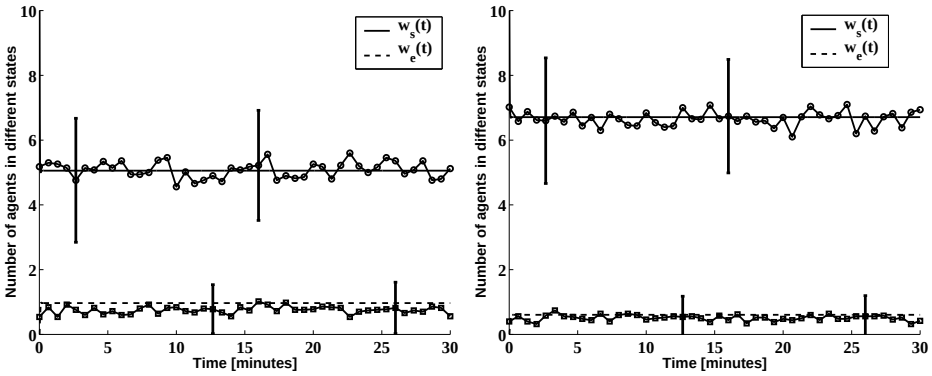

Figure 3.7: Comparing macroscopic model predictions and embodied simulation results. Number of robots in different states (10 robots and 2 seeds). Left: 80 x 80 cm^2 arena. Right: 114x114 cm arena (i.e., twice as large as the previous arena).

agreement between the mathematical model predictions and the embodied simulations of the robots' interaction rate with the seeds in different arenas and with different team sizes.

3.5 Aggregation Experiment without Worker Allocation

In most of the previous aggregation experiments [44, 63, 64], the size of the working team was kept constant during the whole process. Furthermore, the experiments were monitored and terminated when a single cluster arose in the arena. Thus the (destructive) effect of the lack of a mechanism to allow the robots to stop working was never clearly studied. Here we do not end the experiment

when the robots create a single cluster of seeds, instead we let the experiment run for 10 hours. The outcome presented in subsection 3.5.3 justifies the need for a distributed mechanism that enables each robot to switch off when the aggregation task is finished. In concordance with previous works on cluster formation and aggregation [2, 44, 50, 63, 64], we are using three primary team performance measurements: the average cluster size, the average number of clusters, and the average number of active workers in the environment.

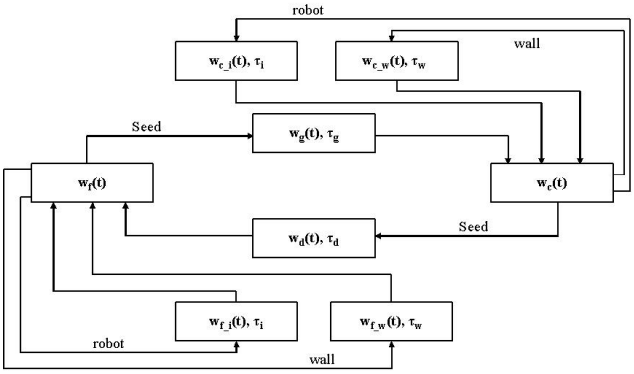

Figure 3.8: PFSM representing aggregation task performance by the robotic team in the macroscopic model (distributed worker allocation mechanism not included).

3.5.1 Mathematical model

After introducing two key sub-chains characterizing the dynamics of the aggregation experiment, we are now ready to analyze the full system model mathematically. As mentioned above, Figure 3.2 shows the FSM of the full system. In this full aggregation system, the number of robots in the search state breaks in two categories: the unloaded and the loaded robots. While the robots in the former category look for a seed to grasp and pick up, robots in the latter look for clusters of seeds to enlarge, releasing the seeds they are carrying next to the clusters they have found. In the DDE system below, we refer to the unloaded robots as the *free* workers and their number is represented by w_f, and to the loaded robots as the *charged* workers whose number is represented by w_c. As a consequence, and also for the sake of clarity, we will distinguish the unloaded robots in obstacle avoidance or interference states (i.e., w_{f_w} and w_{f_i}, respectively) from the loaded robots in the same obstacle avoidance or interference states (i.e., w_{c_w} and w_{c_i}, respectively). Moreover, the side of a large cluster also constitutes an obstacle that the robots will avoid in the same manner as they avoid the surrounding wall.

In the equations below, w_g and w_d represent the number of robots in the seed-grasping and the seed-drop-off states, respectively, Equations (3.19) and (3.20) describe their dynamics and express that w_g increases (resp. decreases) when unloaded robots find isolated seeds and start to pick them

up (resp. when newly loaded robots finish grasping the seeds they have found). Similarly, w_d is increased by the loaded robots that find a cluster and start to drop off the seeds they carry, and decreased when the newly free robots resume looking for seeds. Equations (3.21)-(3.23) translate the dynamics of the number of unloaded workers in wall avoidance and interference states and the number of loaded workers in wall avoidance state, respectively, and express that these numbers change simply as robots enter or leave the states that they represent after the appropriate time-delays. In the last equation we use the conservation of the total number of robots in the environment to derive the number of loaded robots in interference state. M and n_k represent the total number of seeds in the arena and the number of clusters of size k, respectively. The dynamics of n_k are described in detail in a later paragraph of this section.

Thus, the macroscopic model of the full system can be formalized as follows[2]:

$$\begin{aligned}\frac{dw_f}{dt} &= -(\gamma_{f_w}(t)+\gamma_i)w_f(t) - \sum_{k=1}^{M}\gamma_k^{dec}n_k(t)w_f(t) + \gamma_{f_w}(t-\tau_w)w_f(t-\tau_w)\Theta(t-\tau_w) \\ &\quad + \gamma_i w_f(t-\tau_i)\Theta(t-\tau_i) + \sum_{k=1}^{M}\gamma_k^{inc}n_k(t-\tau_d)w_c(t-\tau_d)\Theta(t-\tau_d)\end{aligned} \tag{3.17}$$

$$\begin{aligned}\frac{dw_c}{dt} &= -(\gamma_{c_w}(t)+\gamma_i)w_c(t) - \sum_{k=1}^{M}\gamma_k^{inc}n_k(t)w_c(t) + \gamma_{c_w}(t-\tau_w)w_c(t-\tau_w)\Theta(t-\tau_w) \\ &\quad + \gamma_i w_c(t-\tau_i)\Theta(t-\tau_i) + \sum_{k=1}^{M}\gamma_k^{dec}n_k(t-\tau_g)w_f(t-\tau_g)\Theta(t-\tau_g)\end{aligned} \tag{3.18}$$

$$\frac{dw_g}{dt} = \sum_{k=1}^{M}\gamma_k^{dec}n_k(t)w_f(t) - \sum_{k=1}^{M}\gamma_k^{dec}n_k(t-\tau_g)w_f(t-\tau_g)\Theta(t-\tau_g) \tag{3.19}$$

$$\frac{dw_d}{dt} = \sum_{k=1}^{M}\gamma_k^{inc}n_k(t)w_c(t) - \sum_{k=1}^{M}\gamma_k^{inc}n_k(t-\tau_d)w_c(t-\tau_d)\Theta(t-\tau_d) \tag{3.20}$$

$$\frac{dw_{f_w}}{dt} = \gamma_{f_w}(t)w_f(t) - \gamma_{f_w}(t-\tau_w)w_f(t-\tau_w)\Theta(t-\tau_w) \tag{3.21}$$

$$\frac{dw_{f_i}}{dt} = \gamma_i(t)w_f(t) - \gamma_i w_f(t-\tau_i)\Theta(t-\tau_i) \tag{3.22}$$

$$\frac{dw_{c_w}}{dt} = \gamma_{c_w}(t)w_c(t) - \gamma_{c_w}(t-\tau_w)w_c(t-\tau_w)\Theta(t-\tau_w) \tag{3.23}$$

$$w_{c_i}(t) = W_0 - w_f(t) - w_c(t) - w_g(t) - w_d(t) - w_{f_w}(t) - w_{f_i}(t) - w_{c_w}(t) \tag{3.24}$$

Note that in the DDE system above (i.e., Equations (3.17)-(3.24)) γ_w is a function of time whereas γ_i is not. This is due to the fact that the team size does not vary in the experiment therefore γ_i is a constant as well. Meanwhile γ_w depends both on the surrounding border between the working and the parking zones and the cluster sides that constitute new obstacles that the robots avoid, therefore it depends on the number of different cluster types present in the arena which in turn varies over

[2] In a previous publication [4] we neglected the amount of time spent in obstacle avoidance and interference states. That assumption holds only when that duration is short and the density of robots is low.

time as the seeds are put in clusters of increasing size.

Since the loaded and unloaded workers (equivalently w_c and w_f) do not perceive exactly the same approach angle to manipulate the clusters (i.e., the construction and destruction angles are different) they dot not perceive the same obstacle upon encounter with the side of a cluster. Therefore, γ_w has to be separated in two parameters corresponding to each of these two cases (i.e., γ_{c_w} and γ_{f_w}, respectively). This is clearly linked to the fact that the aggregation experiment is a distributed manipulation process in which robots modify a shared environment. For instance, γ_{f_w} is calculated using geometrical considerations as described in Equation (3.25).

$$\gamma_{f_w}(t) = \frac{v}{S - 2R_w} + \sum_{k=1}^{M} \frac{2R_s v n_k(t)}{S^2} \frac{2\pi R_s + 4(k-1)r_s - 2\alpha_{dec}(k)R_s}{2\pi R_s + 4(k-1)r_s} \tag{3.25}$$

In the right hand side of Equation (3.25), the first term represents the rate at which a robot encounters the surrounding border separating the working and the parking zones, the second term corresponds to the rate at which an unloaded robot encounters the sides of clusters containing more than a single seed. In the last term, the first expression refers to the rate of encountering a cluster of side k and the other expression corresponds to the probability of an unsuccessful approach.

Equation (3.17) tells us that the number of unloaded robots in search state at any time is decreased by the robots which transition to a seed grasping state and those which start avoiding a wall or a teammate; w_f is increased by the robots which come back from releasing a seed and the unloaded robots which finish a wall or robot avoidance maneuver. Similarly, Equation (3.18) describes the dynamics of the number of loaded robots in search state.

The last variables of the aggregation system are the numbers of seed clusters of different sizes that arise and/or disappear over time during the aggregation process. We describe the quantitative dynamics of the seed clusters in the following rate equation:

$$\begin{aligned} \frac{dn_k}{dt} = &(\gamma_{k+1}^{dec} n_{k+1}(t-\tau_g) - \gamma_k^{dec} n_k(t-\tau_g)) w_f(t-\tau_g)\Theta(t-\tau_g) \\ &+ (\gamma_{k-1}^{inc} n_{k-1}(t-\tau_d) - \gamma_k^{inc} n_k(t-\tau_d)) w_c(t-\tau_d)\Theta(t-\tau_d) \end{aligned} \tag{3.26}$$

Equation (3.26) represents the general rate equation for the number of clusters of any size k. On the right hand side of this equation, the first term corresponds to the contribution of the agents that are searching for seeds to pick up. These agents can increase (resp. decrease) the number of clusters of size k by removing a seed from one of size $k+1$ (resp. k). Similarly, the second term translates the contribution of loaded agents which can increase (resp. decrease) the number of clusters of size k by adding a seed to a cluster of size $k-1$ (resp. k).

3.5.2 Steady-state analysis

As shown by Martinoli *et al.* in [65], if the agents do not withdraw (i.e., the team size is constant), do not drop a seed unless it is next to another seed, and do not pick up an internal seed from a cluster, the number of clusters monotonically decreases and eventually a single cluster always arises. A more practical way of defining the steady-state value is that it represents the average size of the unique cluster remaining in the environment as the robots continue picking up and dropping seeds at its two end tips. Based on this assumption, intuitively, the size of the remaining cluster corresponds solely to the total number of seeds in the arena minus the number of those being held by the loaded workers.

$$L^* = M - (w_c^* + w_{c_w}^* + w_{c_i}^* + w_d^*) \tag{3.27}$$

In Equation (3.27), w_x^* is the steady-state value of w_x and L^*, the average size of the single cluster remaining in the arena in the steady-state. L^* is obviously a function of several (steady-state) variables whose values cannot be derived in closed forms based solely on the equations above (i.e., Equations (3.17)-(3.24)). Therefore, we proceed with an approximation of L^* by neglecting the number of robots in obstacle avoidance, interference, seed grasping, and seed drop off states at steady-state (i.e., $w_f^* + w_c^* \approx W_0$) —[3]. Thus by equating (3.27) to zero we obtain Equation (3.28), where $\psi = \frac{\gamma_k^{inc}}{\gamma_k^{dec}} = \frac{\alpha_{inc}}{\alpha_{dec}}$ is a constant for all k such that $1 < k < M$, and Equation (3.29):

$$\psi = \frac{w_f^*}{w_c^*} \tag{3.28}$$

$$L^* \approx M - \frac{W_0}{1+\psi} \tag{3.29}$$

Note that L^* is a decreasing function of the team size W_0. This approximation of the average size of the unique cluster remaining in the arena agrees with the embodied simulation results presented in the next section.

3.5.3 Results

Figures 3.9 and 3.10 present the model predictions and the embodied simulation results of the aggregation experiment without the use of any worker allocation algorithm for groups of 1, 5, and 10 robots in an 80 x 80 cm^2 arena. Figures 3.9 and 3.10, left, present the increasing average size of the clusters over time and Figures 3.9 and 3.10, right, show the decreasing average number of clusters over time. The good agreement between the results collected at both implementation levels illustrates how reliable the macroscopic model is at predicting the outcome of the aggregation experiment.

[3] This assumption holds as long as the density of robots in the arena is low and the different maneuvers (e.g., obstacle avoidance, interference, and seed grasping/drop off) durations are short.

When the team consists of 1 or 5 robots, Figure 3.9 left, presents an increasing average cluster size and, in concordance, Figure 3.9 right shows that the number of clusters monotonically decreases over time. However, with a team of 10 robots, Figure 3.10 left consists of a first phase when the average cluster size increases steadily from 1 seed to about 15 seeds and a second phase when the average cluster size remains, on average, constant around 15 seeds. Similarly, during the first phase Figure 3.9 and 3.10, right, show that the average number of clusters decreases asymptotically from 20 to about 1 then remains close to 1 during the second phase of the aggregation process. This is clearly a destructive effect of the lack of a mechanism to allow the workers to stop performing the aggregation task once a single cluster arises, which is confirmed by the fact that after 10 hours into the aggregation process the average cluster size built by the team of 5 robots is larger than that obtained with the team of 10.

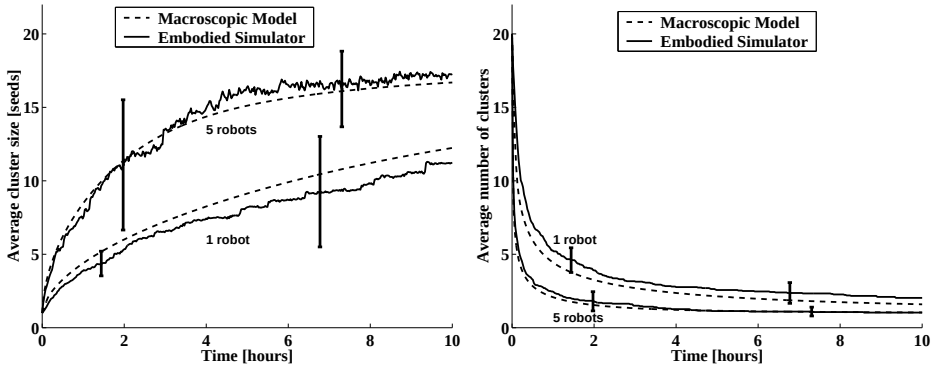

Figure 3.9: Results of aggregation experiment with groups of 1 and 5 robots and 20 seeds in an 80 x 80 cm^2 arena. Group sizes are fixed. Left: Average cluster size over time. Right: Average number of clusters over time.

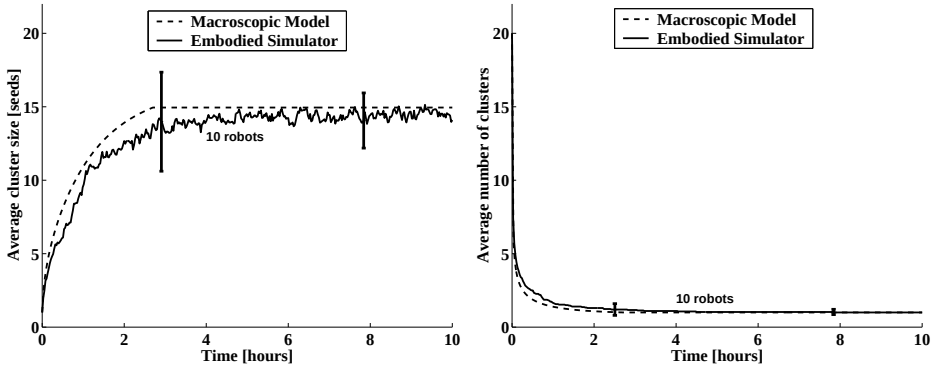

Figure 3.10: Results of aggregation experiment with a group of 10 robots and 20 seeds in an 80 x 80 cm^2 arena. Group size is fixed. Left: Average cluster size over time. Right: Average number of clusters over time.

3.6 Aggregation Experiment with Worker Allocation

In addition to the mathematical model developed for its analysis, the other main objective of this case study is to show that the introduction of worker allocation mechanisms allows the team of robots to increase its efficiency as a whole by allocating the right number of workers as a function of the demand intrinsically defined by the aggregation process. Intuitively, we can imagine that at the beginning of the aggregation there are many possible manipulation sites (i.e., several scattered seeds) that allow for a parallel work of several agents. As the aggregation process goes on, the number of these sites is reduced and having more agents competing for the same manipulation sites decreases their efficiency. Therefore reducing the number of active agents during the last phase of aggregation should consequently increase the size of the single cluster.

However, the main challenge here is to find a simple mechanism that allows each minimalist agent to autonomously evaluate the progress in task accomplishment (i.e., how far along the aggregation process is) so as to be able to decide whether to continue performing the task or not. Swarm intelligence offers a variety of biologically inspired, distributed solutions to this type of problems. In the swarm robotic literature, we find already algorithms inspired by threshold-based models discovered in natural systems [54]. In Krieger and Billeter's experiment the demand related to the nest energy is assessed by an external supervisor and globally transmitted to all the robots. Using this method, the team of robots must be heterogeneous and a different threshold must characterize each agent in order to regulate the team activity. This in turn results in a different exploitation of the teammates, the one endowed with the lowest threshold systematically being more active than that with the highest threshold. We propose here a swarm intelligence-based solution in the form of a threshold-based, distributed worker allocation mechanism based on local demand estimation.

3.6.1 A threshold-based, distributed worker allocation mechanism

In threshold-based systems the propensity of any agent to act is governed by a *response threshold*. If the demand is above the agent's threshold then that agent continues to perform the task, conversely if the demand is below its threshold then the agent stops performing that particular task. In the algorithm presented here the time an agent spends before finding some work to accomplish (i.e., to pick and drop a seed) represents the agent's estimation of the demand stimulus associated with the aggregation task.

Our current worker allocation mechanism is as follows. When an agent has not been able to work (i.e., to pick up and drop a seed) for a reasonable amount of time, its propensity to accomplish the task is decreased. If the stimulus goes below a certain threshold (i.e., if the amount of time spent in the search for work to do is above a given τ_{search} time-out) a deterministic switching mechanism prompts the agent to leave the working zone for resting in the surrounding parking zone. A loaded

agent that decides to become inactive cannot do so until it finds an appropriate spot (i.e., one tip of a cluster) to drop the seed. Thus, with this simple algorithm characterized by a single threshold, each agent is able to estimate the aggregation demand locally and to decide whether to work or rest with no need for a central controller. This task allocation mechanism is similar to that observed in some ant colonies [34, 102], for which it has been shown that an individual performs a task as long as the demand stimulus associated with the task (e.g., a pheromone concentration) exceeds the individual's threshold for that task [15, 96].

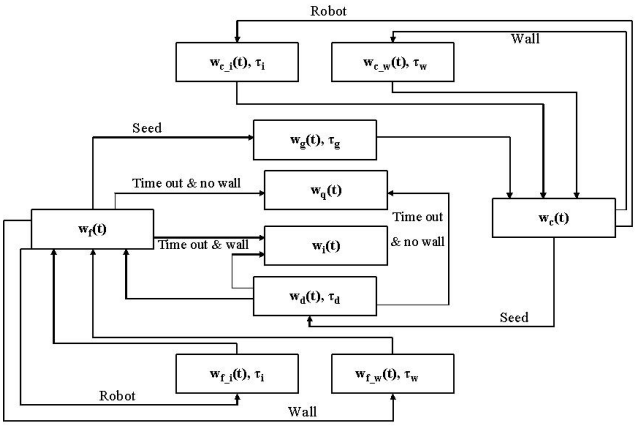

Figure 3.11: PFSM representation of the complete aggregation experiment with distributed worker allocation. The numerical values used in this PFSM have been derived from Table 3.1 and Table 3.2.

Following the worker allocation mechanism above, we assign the same response threshold to all the agents. The team of robots is therefore homogeneous from control point of view. The resulting agents' behavior (i.e., rhythm of activity) is not identical for every agent because it is based on the local, private assessment of the current status of the shared resource (i.e., the environment). In other words, diversity in activity is created by exploiting the intrinsic noise of the system as well as local perceptions and interactions.

3.6.2 Mathematical model

The PFSM of the complete aggregation experiment with distributed worker allocation depicted in Figure 3.11 differs from that in Figure 3.8 by two new states, the *quit* and the *idle* states, represented by w_q and w_i, respectively. The quit state corresponds to the state of the robots that have autonomously decided to stop working and have started traveling from the inner area of the arena (i.e., the working zone) to the surrounding parking zone (see Figure 3.1). The idle state corresponds to the state of the robots that have reached the parking zone after stopping performing the aggregation task. For the sake of brevity, we introduce here only the equations that describe the dynamics of the new variables and those of the old variables (i.e., those from section 3.5) whose

dynamics have changed due to the task allocation mechanism.

$$\begin{aligned}\frac{dw_f}{dt} = &-\left(\gamma_{f_w}(t)+\gamma_i(t)\right)w_f(t) - \sum_{k=1}^{M}\gamma_k^{dec}n_k(t)w_f(t) + \gamma_{f_w}(t-\tau_w)w_f(t-\tau_w)\Theta(t-\tau_w) \\ &+ \gamma_i(t-\tau_i)w_f(t-\tau_i)\Theta(t-\tau_i) - \gamma_{sw}w_f\Gamma_f(t;\tau_{search}) \\ &+ \sum_{k=1}^{M}\gamma_k^{inc}n_k(t-\tau_d)w_c(t-\tau_d)\left(1-\Gamma_c(t;\tau_{search})\right)\Theta(t-\tau_d)\end{aligned} \tag{3.30}$$

$$\frac{dw_q}{dt} = \sum_{k=1}^{M}\gamma_k^{inc}n_k(t-\tau_d)w_c(t-\tau_d)\Gamma_c(t;\tau_{search})(1-P_{sw})\Theta(t-\tau_d) - \gamma_{sw}w_q(t) \tag{3.31}$$

$$\begin{aligned}\frac{dw_i}{dt} = &\sum_{k=1}^{M}\gamma_k^{inc}n_k(t-\tau_d)w_c(t-\tau_d)\Gamma_c(t;\tau_{search})P_{sw}\Theta(t-\tau_d) + \gamma_{sw}w_q(t) \\ &+ \gamma_{sw}w_f(t)\Gamma_f(t;\tau_{search})\end{aligned} \tag{3.32}$$

Note that γ_i is now a time-varying parameter as the number of robots present in the working area changes over time due to the worker allocation algorithm. In Equations (3.30)-(3.32), P_{sw} and γ_{sw} represent the probability that a robot encounters the surrounding wall at any time step and the rate at which a robot encounters the same border separating the inner area (i.e., the working zone) from the outer area of the arena (i.e., the parking zone). These two parameters are fixed and computed using the method based on geometrical considerations explained in subsection 2.6.1. γ_{sw} differs from γ_w, as the latter, in addition to the surrounding border detection, also includes the detection of the sides of the clusters present in the arena, which constitute large obstacles that the robots avoid. P_{sw}, when used in the equations above, translates a successful attempt by the robots to leave the working zone, for resting in the parking zone, and $1-P_{sw}$, a failure to do so.

$\Gamma_f(t;\tau_{search})$ and $\Gamma_c(t;\tau_{search})$, the fractions of *disappointed* unloaded and loaded workers, are the probabilities that an unloaded (resp. loaded) robot has been unsuccessful in finding a single seed to pick up (resp. in dropping a seed) over the past time period $[t-\tau_{search}, t]$ (resp. $[t-\tau_d-\tau_{search}, t-\tau_d]$). Note that $\Gamma_c(t;\tau_{search})$ is defined over the time period $[t-\tau_d-\tau_{search}, t-\tau_d]$ because robots are not allowed to leave the arena with a seed, therefore, a loaded robot that decides to stop working has to find an appropriate spot to drop the seed it is carrying before leaving. Hence, Γ_c only applies to the loaded robots that have finished a seed drop-off started τ_d time steps ago.

The last two terms in the right hand side of Equation (3.30) represent a change compared with Equation (3.17). The first of these terms translates the dynamics of the unloaded robots that decide to stop working and successfully cross the border separating the working and the parking zones; the other term translates the dynamics of the robots that enter the search state after releasing a seed without the intention to stop working.

Equation (3.31) expresses the dynamics of w_q the number of robots that have decided to stop working but have not succeeded to leave the working zone yet; this corresponds to the quit state

depicted in Figure 3.12, a robot in this state leaves the working zone at the rate γ_{sw}.

Finally, Equation (3.32) expresses the dynamics of the robots in the idle state as the number of robots in that state depends on the robots that successfully cross the border on their way to the parking zone whether their decision to stop working was made while they were loaded or unloaded.

To calculate $\Gamma_f(t;\tau_{search})$ (and similarly, $\Gamma_c(t;\tau_{search})$), we use a method similar to that proposed in [57] for calculating the fraction of robots abandoning a stick pulling action. We divide the time interval $[t-\tau_{search}, t]$ into K small intervals of length $\delta t = \tau_{search}/K$. Hence the probability for an unloaded robot not finding a seed to pick up during the time interval $[t-\tau_{search}, t-\tau_{search}-\delta t]$ is $1-\delta t R_{dec}(t-\tau_{search})$, where $R_{dec}(t)$ represents the rate at which an unloaded robot manipulates (at least) one seed during time step t. Similarly, $R_{inc}(t)$ is the rate at which a loaded robot would find at least one cluster tip at which to drop the seed it is carrying during the same time step. It follows that

$$\Gamma_f(t;\tau_{search}) = \prod_{i=1}^{K} \left[1-\delta t R_{dec}(t-\tau_{search}+i\delta t)\right] \Theta(t-\tau_{search}) \tag{3.33}$$

$$\Gamma_f(t;\tau_{search}) \equiv \exp \sum_{i=1}^{K} \ln\left(1-\delta t R_{dec}(t-\tau_{search}+i\delta t)\right) \Theta(t-\tau_{search}) \tag{3.34}$$

Finally, expanding the logarithm in equation (3.34) and taking the limit $\delta t \to 0$ brings

$$\Gamma_f(t;\tau_{search}) = \exp\left(-\int_{t-\tau_{search}}^{t} R_{dec}(t')dt'\right)\Theta(t-\tau_{search}) \tag{3.35}$$

$\Gamma_c(t;\tau_{search})$ is calculated in a similar way. The distributed allocation mechanism is essentially determined by $\Gamma_f(t;\tau_{search})$ and $\Gamma_c(t;\tau_{search})$ which in turn strongly depend on $R_{dec}(t)$ and $R_{inc}(t)$, respectively. $R_{dec}(t)$ is calculated as in Equation (3.36), $R_{inc}(t)$ is calculated in a similar manner.

$$R_{dec}(t) = \sum_{k=1}^{M} \gamma_k^{dec} n_k(t) \tag{3.36}$$

3.6.3 Steady-state analysis

A thorough analysis of the steady-state here is better conducted by equating Equations (3.30)-(3.32) to zero, thus we obtain the following results:

$$w_q^* = 0 \tag{3.37}$$

$$w_f^* \Gamma_f^*(\tau_{search}) = 0 \tag{3.38}$$

$$w_c^* \Gamma_c^*(\tau_{search}) \sum_{k=1}^{M} \gamma_k^{inc} n_k^* = 0 \tag{3.39}$$

$$w_f^* \sum_{k=1}^{M} \gamma_k^{dec} n_k^* - w_c^* \sum_{k=1}^{M} \gamma_k^{inc} n_k^* = 0 \tag{3.40}$$

In fact, if the arena is not overcrowded and the total number of agents is less than that of the seeds (this latter condition prevents the scenario where all the seeds are picked up and the robots cannot drop them next to any other seeds), there will always be at least one cluster of one or more seeds present in the arena. Thus, there exists at least two rate values (e.g., R_{min} and R_{max} positive and non null) such that $R_{min} < R_{dec}(t) < R_{max}$ for all time step t. Hence:

$$e^{-\tau_{search} R_{max}} < \Gamma_f(t; \tau_{search}) < e^{-\tau_{search} R_{min}} \tag{3.41}$$

Equation (3.41) implies two consequences: the first is that $\lim_{\tau_{search} \to \infty} \Gamma_f(t; \tau_{search}) = 0$ and the other is that $\Gamma_f^*(\tau_{search}) = \lim_{t \to \infty} \Gamma_f(t; \tau_{search})$ is never null for a fixed τ_{search}. The latter result and Equations (3.37)-(3.40) in turn imply that there exists a unique steady-state which is the following: $w_f^* = w_c^* = w_{f_w}^* = w_{f_i}^* = w_{c_w}^* = w_{c_i}^* = 0$ and $w_i^* = W_0$. In other words, in a system where the proposed distributed task allocation mechanism is implemented, all the agents will stop working eventually. Note that in reality it may take a long observation bout before obtaining this stable steady-state depending on the value of the fixed activity threshold.

Nevertheless, when the aggregation experiment with worker allocation is observed over a relatively short period of time (which, in fact, depends on the number of seeds to cluster and the total number of robots performing the task), it has different outcomes depending on both the initial setup and, more strongly, on the fixed time-out τ_{search}. On the one hand, if τ_{search} is too small, the propensity of the robots to work will be reduced substantially and this will result in the whole team stopping working too early, preventing the seeds from being gathered in a single cluster. On the other hand, if τ_{search} is too large, the outcome of the experiment will be exactly similar to that of the scenario without worker allocation, if observed over a relatively short period of time.

3.6.4 Results

Figures 3.12 and 3.13 present the outcome of the aggregation experiment using the proposed distributed worker allocation algorithm with teams of 1, 5, and 10 robots in an 80 x 80 cm^2 arena. Each aggregation run lasted 10 hours. We carried out 30 simulation runs using the embodied simulator. All error bars represent the standard deviations among the runs. For all the results presented in Figure 3.12, we hand-coded $\tau_{search} = 25$ minutes.

Figures 3.12 and 3.13 show that here, in contrast to the case without worker allocation, the average cluster size remains an increasing function of time for all three group sizes. When using a single worker, the average cluster size over time is the same as that obtained in the scenario of section 3.5 and the robot does not leave the task; with a group of 5 workers, although the average cluster size over time is similar to that obtained in the scenario without worker allocation, here the average number of workers necessary to obtain the same results is substantially lower, as on average, only about half of the agents remain in the arena after 10 hours. The most important improvement comes with the large team of 10 robots. In fact, with 10 workers, during the second phase of the aggregation, the average cluster size remains an increasing function of time, eventually reaching 20 seeds, the optimal largest value possible, while the number of active workers in the environment decreases such that after 10 hours only about 2 workers remain active.

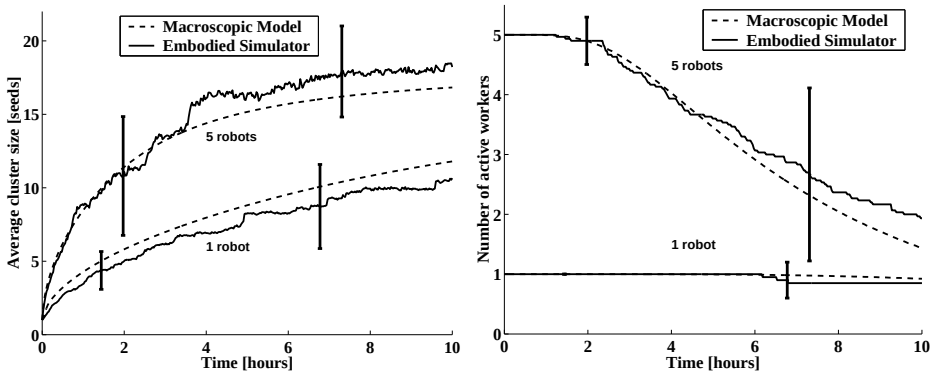

Figure 3.12: Results of aggregation experiment with worker allocation and groups of 1 and 5 robots with 20 seeds in an 80 x 80 cm^2 arena. Left: Average cluster size over time. Right: Average number of active workers over time.

3.7 Discussion

In sections 3.5 and 3.6, we have presented results of the aggregation experiment obtained at two levels of implementation: the macroscopic model and embodied simulations with and without the implementation of the worker allocation mechanism. In this section, we discuss the efficiency of the worker allocation algorithm and compare its performance with that of the case study without worker

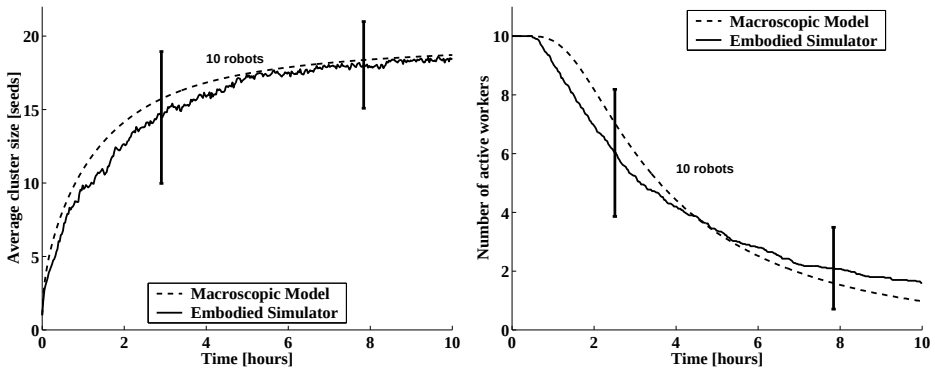

Figure 3.13: Results of aggregation experiment with worker allocation (10 robots, 20 seeds, 80 x 80 cm^2 arena). Left: Average cluster size over time. Right: Average number of active workers over time.

allocation.

3.7.1 Aggregation without worker allocation

Without worker allocation, the aggregation experiments are characterized by a constant team size from the beginning till the end of the experiment. Figures 3.9 and 3.10 present the average size of the clusters over time. The result in Figures 3.9 and 3.10 can be explained by the fact that, once a single cluster arises only two manipulation sites remain in the environment (i.e., the two end tips of that cluster); the seed pick-up and drop-off rates (i.e., γ_k^{dec} and γ_k^{inc}) are empirically very close, therefore at any given time during the last phase of the aggregation process, on average, half of the active workers will be carrying a seed and the other half will not. This explains why the average cluster size at steady-state (i.e., L^*), as expressed in Equation (3.28), is a decreasing function of the team size. For instance, $L^* \approx 15.2$ seeds for the team of 10 robots and Figure 3.10 shows that after 10 hours into the aggregation process the average cluster size built by that team is close to this value for both the macroscopic model and the embodied simulator.

3.7.2 Aggregation with threshold-based distributed worker allocation

When using the worker allocation algorithm, in contrast to the case without worker allocation, the average cluster size remains an increasing function of time. Intuitively, this can be explained by the fact that, with only two manipulation sites remaining in the arena, on average, half of the active agents are carrying a seed while the other half are not, so reducing the number of active agents consequently increases the size of the single cluster.

A variety of threshold-based allocation algorithms may be developed and applied to the aggregation experiment with different successes. In [3] we introduced a suite of threshold-based worker allocation algorithms.

A. Diverse threshold-based allocation algorithms

In [3] we introduced a series of threshold-based allocation algorithms and studied their performance in different situations. In the following, *public* refers to the existence of an explicit, collaborative information flow between the workers on the demand, and *private*, to no information sharing at all. Thus, in what follows we present three different worker allocation algorithms: a private fixed-threshold algorithm, a private variable-threshold algorithm, and a public fixed-threshold algorithm.

The Private Fixed-Threshold worker allocation algorithm (PrFT).— The PrFT algorithm is the same as that detailed in section 3.6, i.e., characterized by a fixed threshold value identical to all the robots, thus the team is homogeneous from the control point of view.

The Private Variable-Threshold worker allocation algorithm (PrVT).— We proposed a variable-threshold algorithm based on a threshold self-calibration rule that works in two steps: a threshold estimation phase followed by a worker allocation phase. During the estimation phase, each autonomous agent evaluates the spatial density of the demand and then sets its response threshold based on that individual estimation. During the allocation phase, the algorithm works (i.e., seed clustering) as explained above. Two parameters govern the self-calibration mechanism: the *Estimation Steps* (ES) and the *Estimation Factor* (EF). Each robot estimates the availability of work in the environment by averaging the amount of time it spends to find some work to accomplish over its first ES successful attempts. The agent then computes its response threshold by multiplying that average amount of time by EF.

Note that although all the robots are characterized by the same values of ES and EF (and therefore the same calibration rule, i.e., this is a homogeneous team), due to their partial perceptions in time and space, the robots will end up with different thresholds in the allocation phase. In addition, since PrVT is a private algorithm, the transition time from one phase to the other is also determined by the agents individually and asynchronously. Equation (3.42) summarizes how each worker computes its response threshold. T_k represents the amount of time an agent spent to find work to do at its k^{th} successful attempt.

$$\tau_{search} = EF\left(\frac{1}{ES}\sum_{k=1}^{ES} T_k\right) \tag{3.42}$$

The Public Fixed-Threshold worker allocation algorithm (PuFT).— In order to allow the multi-agent system to react to dynamic external perturbations (e.g., sudden introduction of additional seeds) more quickly than relying on the slow implicit communication through environmental modifications (referred to as *stigmergy*), we endow all the (simulated) robots with peer-to-peer communication capabilities. The advantage of explicit communication is that each agent is able to gather information about the work demand both from its individual experience and the experience of other teammates that it may encounter. Our current collaborative scheme is as follows. When two agents

are within a certain *Communication Range* (CR), they exchange their estimations of the demand; their individual estimations are then set to the average of their original values. In the PuFT algorithm presented here, the distribution of CR and τ_{search} is homogeneous and fixed *a priori* among the teammates.

Note that the threshold-based algorithms presented above differ from each other by only one characteristic. This was done in the intention of understanding the influence of each of those characteristics (i.e., private vs. public demand estimation and fixed vs. variable threshold mechanisms) on the performance of the algorithms. We show in the rest of the work how this systematic approach can help to forecast the performance of a Public Variable-Threshold algorithm (PuVT).

Having defined a series of threshold-based algorithms, a common metric is necessary to evaluate their performance when applied to the same aggregation experiment.

B. Efficiency and robustness of threshold-based allocation algorithms

In [3] we presented an analysis of the efficiency and robustness of the three allocation algorithms based on a general efficiency metric.

B.1 A metric for efficiency comparison and optimization of algorithmic parameters

To be able to optimize the different algorithmic parameters and to assess the cost effectiveness of these different algorithms, we introduced a cost function whose integrated value over the total observation time, named *Integrated Cost* (IC), corresponds to the total cost of the experiment. The IC represents an efficient combined metric for comparing the influence of each parameter of any given worker allocation algorithm as well as the performances of different algorithms. The cost function is defined as follows:

$$F_{cost}(x_t, y_t, z_t) = \xi_x(X_{opt} - x_t)^2 + \xi_y(Y_{opt} - y_t)^2 + \xi_z(Z_{opt} - z_t)^2 \tag{3.43}$$

Where:

- x_t, y_t and z_t represent the average cluster size, the number of clusters, and the number of active workers at time step t, respectively.
- X_{opt}, Y_{opt} and Z_{opt} are the optimal values of the variables above, e.g., 20, 1, and 0, respectively, when clustering 20 seeds.
- ξ_x, ξ_y and ξ_z are coefficients selected to weight the contribution of each of the variables accordingly.

In the right-hand side of Equation (3.43), the first two terms can be considered as the *penalty cost* (i.e., the cost for work not finished) and the third term can be seen as the *worker cost* (i.e., the cost for hiring workers). The Integrated Cost is computed by using a discrete *Riemann* integration over

any observation time when using a discrete-time model and an integral when using a continuous-time model.

We have not extended the macroscopic model yet to take into account simultaneously the concurrent local communication between the agents and the aggregation task accomplishment nor have we yet mathematically described the dynamics of the variable threshold mechanism. Therefore results presented in Table 3.3 were obtained for an observation time of 10 hours using the probabilistic microscopic model introduced in section 1.3.3, which, in this case study, in addition to being approximately as fast as the macroscopic model, also provides useful statistical information about the aggregation process but requires several runs of the same experiment to deliver statistically significant predictions. We chose $\xi_x = 0.01$ and $\xi_x = \xi_y = 4\xi_z$ in order to obtain a higher penalty cost for unsuccessful aggregation results (cluster size and number of clusters) than worker cost.

We tested the efficiency and robustness of the three algorithms in three different scenarios of aggregation concerned with the presence of static and dynamic perturbations. We started by optimizing their respective parameters using a systematic search for an experimental setup that we refer to as Arena1: an 80 x 80 cm^2 arena, 20 seeds originally scattered in the arena, and 10 active agents at the beginning of the experiment. We then moved the teams of robots running the different algorithms into a 178 x 178 cm^2 arena (same number of seeds, i.e., 20 seeds), which has 5 times the area of the original arena, to study the influence of a static perturbation on the performances of the three different algorithms. We refer to the new scenario as that of Arena2. Finally, we proceeded by studying the influence of dynamic perturbations on the cluster formation for the different algorithms using an 80 x 80 cm^2 arena, initially with 20 seeds, and introducing 5 additional seeds 2 hours after the start.

B.2 Performance summary of the threshold-based allocation algorithms

Table 3.3 summarizes the performance of the allocation algorithms in the different scenarios detailed above. Values with an asterisk sign correspond to the best performances (i.e., lowest cost) calculated based on the following criterion. Assume that (E_i, ϵ_i) represents the pair of mean and standard error of the IC associated with the set of simulation runs i. We consider the IC achieved by a set of simulations i to be better (more efficient) than that of a set j if $E_i + (1+\eta)\epsilon_i < E_j - (1+\eta)\epsilon_j$, where η is a criterion parameter. For the results reported here we used $\eta = 0$. The last row in Table 3.3 represents the results of the aggregation experiment without the use of any worker allocation mechanism.

From Table 3.3 and Figure 3.14 it appears that the PuFT and PrFT algorithms are the most efficient in Arena1 (as their performances are equivalent following the criterion), PrVT in Arena2, and PuFT in Arena3.

PrFT.— Due to its a priori fixed response threshold value the agents behave sub-optimally in an environment different from that for which the algorithm was optimized. For instance, when

Table 3.3: Performance summary: Integrated Cost.

Algorithm	Arena1	**Arena2**	Arena3
PrFT	$138.9 \pm 7.0^*$	324.9 ± 10.8	154.5 ± 7.9
PrVT	155.1 ± 8.0	$231.9 \pm 10.7^*$	152.2 ± 8.7
PuFT	$138.2 \pm 6.9^*$	337.6 ± 10.7	$122.4 \pm 6.4^*$
W/o WA	227.4 ± 4.8	310.8 ± 8.8	197.2 ± 5.9

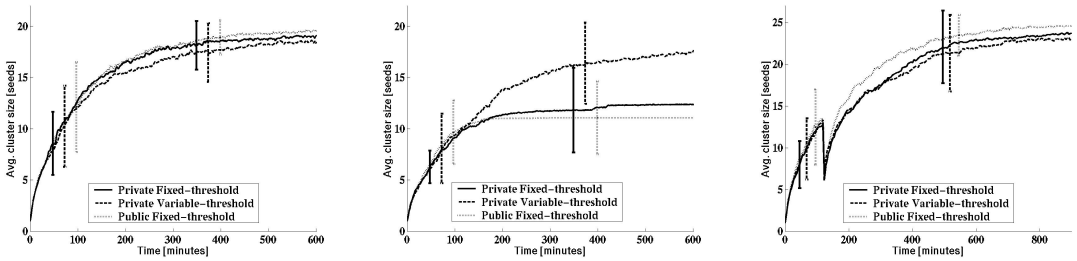

Figure 3.14: Mean cluster size (20 seeds, initial group size of 10 robots). Left: Results from Arena1 (80 x 80 cm^2 arena). Center: Results from Arena2 (178 x 178 cm^2 arena). Right: Results from Arena3 (80 x 80 cm^2 arena, 5 additional seeds introduced 2 hours after the start of the aggregation process).

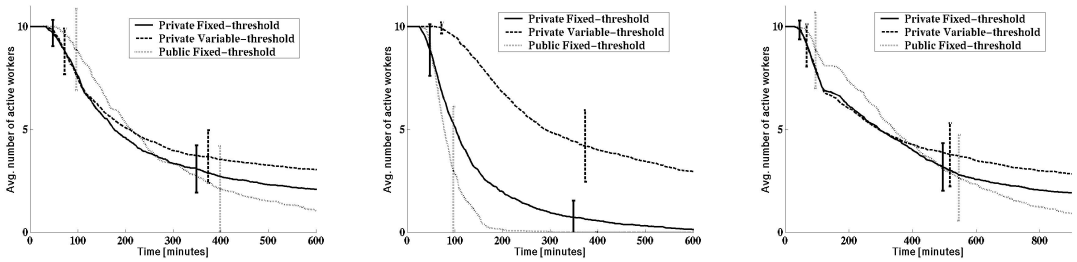

Figure 3.15: Mean number of active workers (20 seeds, initial group size of 10 robots). Left: Results from Arena1 (80 x 80 cm^2 arena). Center: Results from Arena2 (178 x 178 cm^2 arena). Right: Results from Arena3 (80 x 80 cm^2 arena, 5 additional seeds introduced 2 hours after the start of the aggregation process).

performing the same aggregation task in a 178 x 178 cm^2 arena, the average size of the clusters they create are smaller than the average size of those created by the team using the PrVT algorithm (with similar standard deviations) because the agents withdraw too soon.

PrVT.— The density of manipulation sites (seeds that can be manipulated) is higher in the smaller arena and the robots are more likely to encounter them than in the larger arena. In response to this difference in density of manipulation sites, variable-threshold workers autonomously set their response thresholds higher in Arena2. Therefore, they stay active longer in the larger arena than in the smaller one and this in turn allows them to continue performing the task, as most seeds are not gathered yet into a single cluster. This is illustrated by Figure 3.14 center, where it clearly appears that PrFT and PuFT underperform due to a relatively too low homogeneous threshold value and their inability to adapt to a new environment. However, the PrVT algorithm is not appropriate for an optimal response of the agents to a dynamic change in the number of objects to manipulate. For instance, results in Table 3.3 show that when additional seeds are dropped in the arena 2 h into the aggregation process, the efficiency of the PrVT algorithm deteriorates. This results from the absence of a continuous, adaptive activity threshold mechanism that allows the agents to upgrade their activity thresholds when facing a sudden increase in the availability of work.

PuFT.— This algorithm can efficiently deal with a dynamic change in the number of objects to manipulate in a distributed way. However, because the threshold is fixed, the PuFT algorithm does not allow the agents to respond efficiently to all modifications of the arena surface as shown in Figure 3.14 and Table 3.3. Nevertheless, this algorithm has the advantage of providing the team of autonomous agents with the ability to quickly access the information about dynamic changes brought to the working environment through its peer-to-peer communication scheme. As shown in Figure 3.6, the sudden increase in the availability of work due to the additional seeds dropped in the environment was quickly sensed by the teammates. This results in an appropriate number of agents staying active to accomplish the task throughout the aggregation process.

B.3 Performance analysis using the mathematical model

Arena2.— This arena is five times as large as Arena1 therefore for a same cluster distribution, it takes, on average, more time to find a seed to manipulate in this arena than in Arena1. According to the proposed model, whose parameter calculation is based on geometrical considerations (see section 3.3), it results that the rate at which the robots encounter a cluster of any given size in Arena2 is one fifth the rate at which they encounter that cluster in Arena1. Hence, the values of $R_{dec}(t)$ and $R_{inc}(t)$ in Arena2 are one fifth of those in Arena1. Hence, according to Equation (3.35), the value of $\Gamma_f(t;\tau_{search})$ in both arenas (let us name them $\Gamma_{f_1}(t;\tau_{search})$ and $\Gamma_{f_2}(t;\tau_{search})$ in Arena1 and Arena2, respectively) are now linked by the following relationship (a similar relationship links

$\Gamma_{c_1}(t;\tau_{search})$ and $\Gamma_{c_2}(t;\tau_{search})$):

$$\Gamma_{f_2}(t;\tau_{search}) = (\Gamma_{f_1}(t;\tau_{search}))^{\frac{1}{5}} \tag{3.44}$$

This corresponds, in particular, to a considerable increase in the fractions of unloaded and loaded robots characterized by a fixed activity threshold (i.e., PrFT and PuFT) abandoning the aggregation task in Arena2 as compared with Arena1. This outcome is not observed with team characterized by PrVT because, intuitively, we can foresee that the resulting threshold distribution in this team is heterogeneous with some agents characterized by large threshold values.

Arena3.— Introducing 5 additional seeds at a (perturbation) time $\tau_p = 2$ hours suddenly increases the availability of work (i.e., the number of seeds to manipulate) at time τ_p. However, the information about this increase is more rapidly disseminated into the network of agents using the public demand estimation mechanism (i.e., PuFT) than in the networks of agents using the private demand estimation one (i.e., PrFT and PrVT). This can be explained by the fact that information dissemination is faster in the robot network with an explicit communication scheme than in the stigmergic, implicit communication scheme (see also results in Chapter 2).

Mathematically, at time τ_p the rates (and similarly the probabilities) of seed manipulation are suddenly increased. This results in a sudden decrease of the fractions of unloaded and loaded robots abandoning the aggregation task. For instance, $R_{dec}(t = \tau_p)$ in Arena3 is equal to that in Arena1 plus $5\gamma_1^{dec}$, where γ_1^{dec} is the rate at which an isolated seed (i.e., a cluster of size 1) is found and picked up when present in the arena as previously defined. Equation (3.45) expresses that decrease for the fraction of unloaded agents with a reduction factor of $e^{-5\tau_{search}\gamma_1^{dec}}$, which corresponds to the sudden availability of 5 newly introduced seeds. This effect will remain until the newly introduced seeds are all picked up and gathered with the original seeds into a single cluster.

$$\Gamma_{f_2}(t = \tau_p;\tau_{search}) = \Gamma_{f_2}(t = \tau_p;\tau_{search})e^{-5\tau_{search}\gamma_1^{dec}} \tag{3.45}$$

In summary, the fixed-threshold mechanism does not perform as well as the variable-threshold mechanism in the presence of a static, environmental change and the private demand estimation mechanism does not perform as well as the public mechanism in the presence of external perturbations similar to the introduction of additional seeds. Thus, intuitively, by combining the public demand estimation and the variable-threshold mechanisms, we should obtain an algorithm, i.e., the Public, Variable-Threshold worker allocation algorithm (PuVT), whose performance is not undermined when facing static environmental or external perturbations. This is, in fact, confirmed by the results summarized in Table 3.4.

Table 3.4: Performance summary: Integrated Cost.

Algorithm	Arena1	Arena2	Arena3
PrFT	$138.9 \pm 7.0^{*}$	324.9 ± 10.8	154.5 ± 7.9
PrVT	155.1 ± 8.0	$231.9 \pm 10.7^{*}$	152.2 ± 8.7
PuFT	$138.2 \pm 6.9^{*}$	337.6 ± 10.7	$122.4 \pm 6.4^{*}$
PuVT	$141.3 \pm 5.2^{*}$	$227.2 \pm 9.4^{*}$	$134.2 \pm 9.1^{*}$

3.7.3 The mathematical model as an optimization tool

The modeling methodology we present in this paper can be useful also from optimization point of view. First, the proposed model delivers results in time lapses that are at least four orders of magnitude shorter than a corresponding embodied simulation. Second, abstraction in general allows researchers to understand the role of key system parameters, to generalize, and to analyze underlying principles. In particular, we think that the mathematical model may be improved further so as to predict the optimal threshold distribution among the teammates based solely on the experimental setup.

Moreover, we can use the predictions of the macroscopic model to find the best fixed threshold value for the aggregation experiment described in section 3.6. For instance, Figure 3.16 presents the predicted average cluster size obtained after 10 hours of aggregation for different values of the activity threshold τ_{search}. This figure illustrates that for a group of 10 robots, an activity threshold value of about 25 minutes provides the optimal cluster size at time $t = 10$ hours into the aggregation. On the one hand, when the threshold value is lower than 25 minutes, the workers stop performing the task too early without all the seeds already in a single cluster. On the other hand, when the threshold value is larger than 25 minutes, the agents keeps working even after a single cluster has been created, as a consequence, 10 hours into the aggregation some seeds are still been carried by some active agents decreasing the efficiency of the team. This behavior happens more often (i.e., more robots overstaying in the arena) as the threshold value becomes much larger than 25 minutes.

3.8 Conclusion

In this chapter, we have presented a mathematical methodology for studying aggregation experiments (and robotic manipulation in general) using embodied agents in groups of fixed or variable sizes. The proposed macroscopic model uses no free parameters. We have defined a clear mapping between the parameters of the model and the geometrical characteristics of the environment. We have validated the predictions of the macroscopic model with a realistic, sensor-based, embodied simulator. Results show that the proposed approach delivers quantitatively accurate predictions and constitutes a

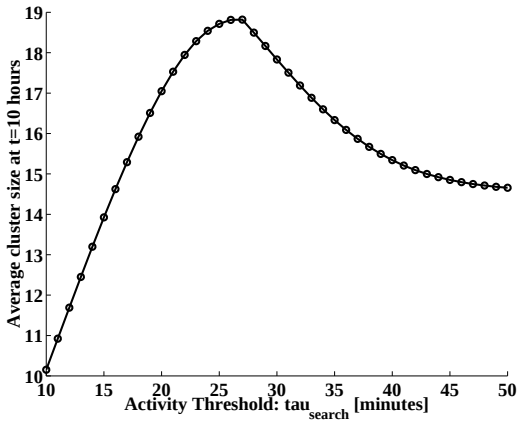

Figure 3.16: Mathematical model prediction of the average cluster size at $t = 10$ hours as a function of the activity threshold τ_{search} (80 x 80 cm^2 arena, 10 robots, 1-minute step from 10 to 50 minutes).

computationally efficient tool.

Moreover, we have presented first a scalable, distributed, threshold-based worker allocation algorithm that allow a team of autonomous, embodied agents to dynamically allocate an appropriate number of workers to a given task based solely on their individual estimations of the progress in the execution of the task. We have shown that teams of active workers dynamically controlled by the allocation algorithm achieve similar or better performances in aggregation than those characterized by a constant team size, while using, on average, a considerably reduced number of agents over the whole aggregation process.

Second, we have presented a comparative study of four scalable, distributed, threshold-based worker allocation algorithms, comparing their efficiency and robustness in a collective manipulation case study concerned with gathering and clustering of seeds. After a systematic optimization process in a given environment common to all the algorithms, it appeared that quantitative differences in efficiency among these allocation algorithms are less apparent. Although PuFT and PuVT appear to be the two most efficient algorithms in the smaller arena, the energy cost due to communication was not included in the experiments. In addition, even if the demand can be estimated more accurately by sharing information and the collective reaction to external perturbations is faster, the PuFT algorithm still suffers from the same drawback as the PrFT algorithm in facing environmental changes for which its threshold was not *a priori* optimized.

After understanding the influence of the type of demand estimation (private vs. public) and of the activity threshold (fixed vs. variable), we combined the characteristics of threshold variability and information sharing to obtain a PuVT algorithm that performs optimally in all the different scenarios presented in this chapter.

It is worth noting that an agent that withdraws from a task can allocate itself to a different task following a similar response-to-stimulus mechanism, thus making these task allocation algorithms

easily applicable to multi-task problems and more complex labor division schemes.

The proposed model has helped us shed light on the decrease in performance observed when using the threshold-based algorithms in certain scenarios. The simplicity of the model suggests that it is easily applicable to other aggregation/segregation or sorting experiments characterized by different agent capabilities and individual control algorithms. Future work can include an effort to investigate different methods of analyzing, at the macroscopic level, the noise due to agents' partial perception. A continuously adapting algorithm as suggested in [96] also seems to be a promising solution to allow the agents to still perform efficiently in the presence of static and dynamic external perturbations brought into the environment. Furthermore, the performance of threshold-based allocation algorithm still needs to be compared with that of other algorithms such as market-based and publish/subscribe messaging algorithms [32].

Finally, we admit that our task is facilitated by the fact that our individuals, instead of being natural creatures, are artificial ones in which we can precisely measure sensors and actuators data or the activation of a given behavior. Nevertheless, we hope that this work will not only contribute to the autonomous robotic literature with the mathematical tool it proposed but also stimulate the discussion among biologists interested in principles of division of labor in social insect societies.

Chapter 4

Modeling Collaborative Distributed Manipulation in Swarm Robotic Systems

4.1 Introduction

In this chapter[1], we present a new discrete-time, incremental methodology for microscopic and macroscopic modeling of the dynamics of distributed manipulation experiments using swarms of autonomous robots. As in Chapter 3, the robots are endowed with reactive controllers. Here, we focus on a concrete case study concerned with using a team of robots to pull sticks out of the ground, an action that requires the collaboration of two robots to be successful. Experiments were carried out with teams consisting of two to 600 individuals at different levels of abstraction (real robots, embodied simulations, microscopic and macroscopic models). Results show that models can deliver both qualitatively and quantitatively correct predictions in time lapses that are at least four orders of magnitude smaller than those required by embodied simulations and that they represent a useful tool for generalizing the dynamics of these highly stochastic and nonlinear systems, often outperforming intuitive reasoning. Finally, in addition to discussing subtle numerical effects, small prediction discrepancies, and difficulties in generating a simple mapping between different abstractions levels, we conclude the paper by reviewing the intrinsic limitations of the current modeling methodology.

4.2 The Stick-Pulling Experiment

In the case study described in this chapter, robots must pull sticks out of the ground, an action that, due to the length of the sticks, requires the collaboration of two robots to be successful. The metric used to quantitatively investigate and model the effects of variation of the parameters of the

[1] The results presented in this chapter derived from a collaborative work with Alcherio Martinoli and Kjerstin Easton and will appear in [68].

system is the collaboration rate among robots, i.e., the number of sticks successfully pulled out of the ground over time.

4.2.1 Experimental setup

The experiment is carried out in a circular arena (40 cm of radius) delimited by a white wall. Four holes situated at the corners of a square with 30 cm edges hold white sticks (15 cm long, 1.6 cm of diameter) that, in their lowest position, protrude 5 cm above the ground (see Figure 4.1 left). Groups of two to six Khepera robots, equipped with gripper turrets, are used to pull the sticks out of the ground. Much like the seeds in the aggregation experiment, the Khepera can distinguish the sticks from the wall and from other robots using its six frontal infrared proximity sensors. The sticks are too long to be pulled from the ground by a single robot grasping and lifting it, therefore, collaboration between two (and only two) robots is required. After a successful collaboration, the stick pulled out of a hole is released by the robot, and replaced in the hole by the experimenter.

Figure 4.1: Left: Physical setup for the stick-pulling experiments. Center: Setup in the embodied simulator (4 sticks, arena of radius 40 cm). Right: Simulated Setup with 16 sticks and arena of radius 80 cm.

4.2.2 Embodied simulations

In order to more systematically investigate the collaboration dynamics, we also implemented the experiment in Webots (see section 1.3.4 and [74]). The mean speed ratio for this experiment with five robots between Webots and real time is about 18 on a Pentium III, 900 MHz machine.

4.2.3 The robot's controller

In a similar manner as presented in the previous chapter, we represent the behavior of the robot with a standard flow chart or a Finite State Machine (FSM), as depicted in Figure 4.2 left.

In addition to a default behavior (i.e., search by moving in a straight line) and an obstacle avoidance behavior, the robot is endowed with a stick-gripping and -pulling procedure. The robot can determine from its arm elevation speed (while lifting a stick from a hole) whether another robot

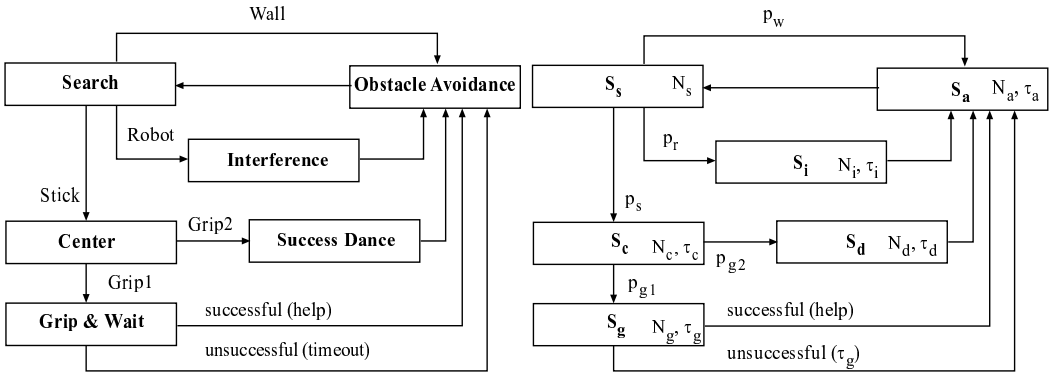

Figure 4.2: Left: FSM representing the robot's controller. Transition between states are deterministically triggered by sensory measurements. Right: PFSM representing an agent or the whole robotic team in the macroscopic model.

is already gripping the same stick. When the robot has pulled a stick half way out of the hole, it holds on and waits for another robot to come and finish pulling the stick out entirely. While waiting for collaboration, the robot can detect another robot's attempt to lift the stick through the arm elevation sensor. If no other robot is holding the stick, we call such a grip a *grip1*. If another robot is already holding the stick, such a grip is called *grip2*. When a robot makes a grip1, it holds the stick raised half-way out of the ground and releases it when either the duration of the grip exceeds the *gripping time parameter* τ_g (i.e., a failed collaboration), or another robot comes to make a grip2 (i.e., a successful collaboration). Once the stick is released, the robot turns away, performs obstacle avoidance for a few seconds, then resumes searching for sticks. When a robot makes a grip2, the robot making the grip1 will release the stick, allowing its teammate to raise the stick completely. The robot making grip2 performs a short "success dance" (moving the arm up and down) to mark the successful collaboration, then releases the stick (which has to be replaced in the hole by the experimenter), performs obstacle avoidance for a few seconds, and resumes searching for sticks too.

Because of the way sticks are recognized (only by their thinness), a stick that is held by one robot can only be recognized when approached from the opposite side within a certain angle (approximately 125 degrees in the physical set-up). For the other angles of approach, both the stick and the robot are detected and the whole is taken for an obstacle. More details are reported in [48] but we note that the acceptable approach angle, expressed as a ratio R_g over the whole approaching perimeter, is an important parameter in the collaboration dynamics of the system.

4.3 Calibrating Model Parameters

4.3.1 Transition probabilities

Consistent with the approach used throughout this thesis, we compute the transition probabilities from a state to another based on simple geometrical considerations about the interaction (e.g., detection areas, approaching perimeters). The numerical values used for these geometrical parameters are measured in systematic tests with one or two real robots as we did in the previous chapter. At each iteration, the probabilities that a robot in the search mode encounters a wall, a stick, or another robot is determined by their corresponding detection areas divided by the total arena area A_a. For instance, the probability of finding a stick is computed as $p_s = A_s/A_a$, where A_s is the detection area of a stick. Similarly, $p_w = A_w/A_a$ and $p_r = A_r/A_a$ represent the probability of encountering a wall and another robot respectively. The total probability of encountering any other robot on the arena can be computed as $p_R = (W_0 - 1)p_r$, with W_0, the total number of robots on the arena. Additionally, since robots can perform a grip1 from any angle of approach, $p_{g_1} = p_s$. On the other hand, since a stick available for grip2 can only be approached from a certain angle as mentioned above, the probability of a grip2 event is $p_{g_2} = R_g p_{g_1}$. The total number of sticks in the arena is M_0.

As explained in the previous chapter, this proposed methodology for calculating the different probabilities does not allow for overlapped areas of detection between objects of different types (e.g., robot, wall). Equation (4.1) expresses this limit mathematically.

$$p_w + p_s M_0 + p_R \leq 1 \tag{4.1}$$

The boundary for the maximal detection areas occupied by robots follows as $p_R = 1 - p_w - p_s M_0$. As we will see in subsection 4.6.2, this remains an approximation as for example additional robots can still "squeeze" into the available free space. In particular, in overcrowded scenarios, the modeling methodology reaches its limitations.

The numerical values used for the robot speed, the approaching angle for grip2, and the different detection radii of the objects are summarized in Table 4.1. These values are exactly the same as those reported in [48], although a single arena of 40 cm in radius was used in that experiment.

4.3.2 Time discretization

The probabilistic methodology proposed in this chapter generates discrete-time models.

Time discretization interval.—Consistent with previous works [2, 48, 63, 64, 66] each iteration of our models corresponds to a time step of a finite duration in real time. The duration of a time step is equivalent to the amount of time a robot (moving at an average speed v, with a certain detection

Table 4.1: Parameters used in the models of the stick-pulling experiment. Detection distances are meant from the center of the robot to the center of the object detected.

Mean robot speed v[cm/s]	Approaching angle ratio R_g	Wall detection distance R_w[cm]	Stick detection distance R_s[cm]	Robot detection distance R_r[cm]	Arena radius R_a[cm]
8	0.35	6	6.4	10	40 to 5000

width w_i for the smallest object i in the arena, i.e., a stick) needs to cover the detection area of the smallest object. Equation (4.2) shows how to compute the real duration, T, of a time step in the modeling methodology. Note that T here is equivalent to the conversion coefficient (τ_{i2t}) introduced in subsection 2.5.3 of Chapter 2 as both represent the duration of an iteration in time units.

$$T = \frac{A_s}{vw_s} = \frac{\pi R_s^2}{v(2R_s)} = \frac{\pi R_s}{2v} \tag{4.2}$$

Choosing the smallest object is a way to ensure that the time granularity is small enough to capture every occurrence of the system's fastest detection event (but this approach does account for the wall properly). This is an approximation which links time partitioning with probability space partitioning and is consistent with the fact that the models are nonspatial. Although Ijspeert *et al.* [48] used a slightly different robot detection width (a value that is a weighted sum of all the different detection widths specific to different objects in the arena), we believe that the approach proposed here helps to avoid the introduction of what could be thought of as a free parameter. Although this way of choosing the time step is completely heuristic, it has provided, combined with the method of calculating the transition probabilities explained in subsection 4.3.1, the best results so far not only in the stick-pulling experiment but also in aggregation experiments performed with different robotic platforms. Nevertheless, this approach has not yet being applied to other experiments using completely different robotic platforms. We will discuss the difficulties inherent to parameter calibration and time discretization more extensively in subsection 4.6.2.

Continuous and discrete models parameter values.— Table 4.2 summarizes the values of delays and time steps used in all the models presented in this chapter. The continuous-time parameters are exactly the same as those used in [48]. As the time step is different from the original publication, the corresponding discrete-time parameters may be slightly different after round up into integer values. Of course, a given gripping time parameter τ_g will be also discretized into T_g iterations in the models.

Table 4.2: Duration of the different robot maneuvers described in subsection 4.2.3.

Delay	Centering	Success dance	Obstacle avoidance	Interference	Time step
Continuous [s]	$\tau_c = 10$	$\tau_d = 6$	$\tau_a = 1$	$\tau_i = 2$	$T = 1.26$
Discrete [iter.]	$T_c = 8$	$T_d = 5$	$T_a = 1$	$T_i = 2$	$T = 1$

4.4 Examples of Simplified 2-State Models

Before describing implementation details and results of the full system model we use the same incremental approach as in the previous chapter and start by introducing two examples of simple 2-state PFSMs representing key sub-chains of the full stick-pulling experiment (which itself will be described in section 4.5).

For all the results presented in this section, experiments using the microscopic model were repeated 100 times and error bars represent standard deviations among the simulation runs.

4.4.1 Search and obstacle avoidance

The default behavioral state in the stick-pulling experiment is the search state thus, a first key sub-chain of the system is represented by the search state coupled with an obstacle avoidance state. In reality, depending on the type of obstacle, a collision avoidance maneuver may be characterized by different durations and different probabilities of occurrence. For the sake of simplicity, in this section we will use mean values for probabilities and durations calculated from Tables 4.1 and 4.2. Figure 4.3 represents graphically the state diagram of this simple PFSM.

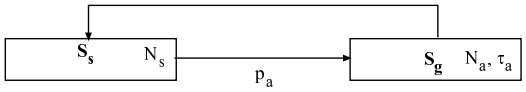

Figure 4.3: A simple sub-chain consisting of a search and an obstacle avoidance state.

The numerical values used in the example of Figure 4.3 were derived from the values of Tables 4.1 and 4.2 using parameters for a generic obstacle: $T_a = 2$ iterations and p_a is a function of the setup (for example, $p_a = 0.65$ for the mean probability of encountering an obstacle in an arena of 40 cm, 4 robots, and 4 sticks; in this case teammates, walls, and sticks are considered all to be obstacles). W_s and W_a represent the numbers of robots in the search state and obstacle avoidance states respectively. In the microscopic model, these are integer values and correspond to the actual

numbers of robots in each state, while in the macroscopic model, these are the average numbers of robots.

The PFSM of Figure 4.3 representing the whole swarm can be described by the following DE system:

$$W_s(k+1) = W_s(k) - p_a W_s(k) + p_a W_s(k - T_a) \tag{4.3}$$

$$W_a(k+1) = W_0 - W_s(k+1) \tag{4.4}$$

where $k = 0, 1, 2, \ldots$, represents the current iteration and $W_x(k)$, the value of the state variable W_x at time kT. The notation $W_x(k)$ instead of $W_x(kT)$ is standard in the automatic control literature. Equation (4.3) states that the average number of robots in the search state at iteration $k+1$ is equal to the average number of robots in the same state at iteration k minus those that left for an obstacle avoidance maneuver plus those that terminated such a maneuver and returned to the search state. Equation (4.4) simply exploits the conservation of the total number of robots for calculating the average number of robots in obstacle avoidance. The system is causal thus, in the proposed model we assume that no robots exist before $k = 0$ (a standard assumption for this type of time-delayed DE-based model, see [6, 67]). This condition is translated mathematically as follows:

$$W_a(k) = W_s(k) = 0 \quad \text{if} \quad k < 0 \tag{4.5}$$

A. Steady-state analysis

Since the DE system (4.3)-(4.4) is linear, we can analyze the steady-state of the system either using a z-transform (frequency domain) or in time domain. Both analysis bring to the same result.

Frequency domain analysis.— Equation (4.3) can be transformed using the right shift and left shift theorems in the z-space as follows:

$$zW_s(z) - zW_0 = W_s(z) - p_a W_s(z) + p_a W_s(z) z^{-T_a} \tag{4.6}$$

Solving for $W_s(z)$ and applying the limit theorem we obtain

$$W_s^* = \lim_{k \to \infty} W_s(k) = \lim_{z \to 1} (z-1) W_s(z) = \frac{W_0}{1 + p_a T_a} \tag{4.7}$$

Therefore, through the conservation of the total number of robots we obtain

$$W_a^* = W_0 - W_s^* = \frac{W_0 p_a T_a}{1 + p_a T_a} \tag{4.8}$$

Let $W^* = [W_s^*, W_a^*]$ represents the state vector of the DE system in the steady-state regime.

Time domain analysis.— Equation (4.4) can been written also in the same form as Equation (4.3), i.e., as a DE, as follows:

$$W_a(k+1) = W_a(k) + p_a W_s(k) - p_a W_s(k - T_a) \tag{4.9}$$

Equation (4.9) represents a delay state. A first intuitive step for calculating the steady-state of this equation could be to set $W_a(k+1) = W_a(k) = W_a^*$ and $W_s(k) = W_a(k-T_a) = W_s^*$. However, we can immediately see that in this case Equation (4.9) will be undetermined, i.e., the trivial equation $0 = 0$. A work-around step for this problem is as follows. We can say that after T_a iterations, every robot that enters the delay state must leave this state with probability one. Since no robots exist for negative iterations, the only time during which W_a can be increased is during the first T_a iterations (full inflow with probability p_a and zero outflow). Although $W_s(k)$ may vary during the first T_a iterations, we can approximate it as a constant if T_a is small.

$$W_a^* = \int_0^{T_a} p_a W_s^* dk = p_a W_s^* T_a \tag{4.10}$$

By combining Equation (4.10) with the conservation of the total number of robots, it becomes easy to demonstrate that W_s^* and W_a^* can be expressed as we calculated via the z-transformation (i.e., Equations (4.7) and (4.8)). This second method of calculating the steady-state vector for delay states is extremely useful if the whole system involving this type of state is nonlinear. This will be the case in the full system model describing the dynamics of the stick-pulling experiment. Figure 4.4 shows graphically how the average number of searching robots at the steady-state is influenced by the time delay T_a. Figure 4.4 shows that the greater the swarm size is, the smaller p_a is, since in

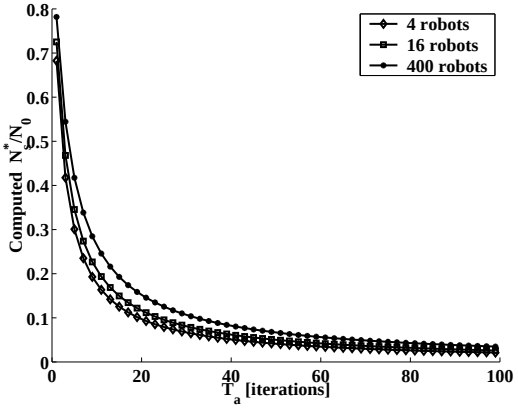

Figure 4.4: Graphical representation of Equation (4.7) for different delay durations and different probabilities of encountering an obstacle.

our case, the wall detection surface becomes proportionally smaller in bigger arenas, decreasing like

$1/R_a$.

B. Model and simulation results

Figure 4.5 shows a comparison between microscopic and macroscopic predictions for different system regimes (transient, steady-state) and different team sizes. In this figure the height of the microscopic column represents the mean value of W_s over a 15-minute time window for 100 simulation runs. The error bars represent a mean standard deviation calculated as an average of the standard deviations measured in each run over a 15-minute time period. It is worth noting that, although the standard deviation among runs and over time of the microscopic model is about inversely proportional to the swarm size, both predictions perfectly agree on the stationary mean number of robots in the different states for different team sizes.

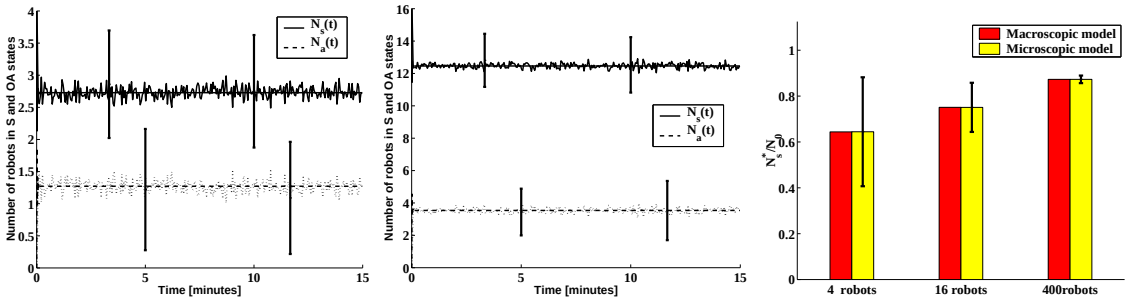

Figure 4.5: Comparing microscopic and macroscopic models results. Left: time evolution of W_s and W_a (4 robots, 40 cm arena). Center: time evolution of W_s and W_a (16 robots, 80 cm arena). Right: Steady-state value of $W_s(W_s^*)$.

4.4.2 Search and stick-gripping states

A second key element of the full system model is specific to the strictly collaborative nature of the stick-pulling experiment and the way robots indirectly communicate through stick manipulation. If we look carefully at the PFSM of Figure 4.2, we can notice that several states can be thought of as simple delay states. The durations of the delays do not exceed a few seconds and are therefore much shorter than most of the values of the gripping time parameter considered in this paper (up to 600 s). As a consequence, we can still capture the essential sub-chain responsible for the collaboration dynamics in the stick-pulling experiment by neglecting these delays and reducing the model to two states: search and grip. The numerical values used in this simplified model have been derived using Table 4.1 and the time discretization step is that mentioned in Table 4.2.

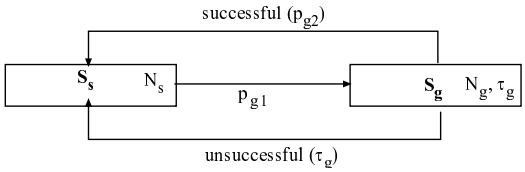

Figure 4.6: The key sub-chain representing the dynamics of collaboration in the stick-pulling experiment.

The following system of DEs represents the macroscopic model for this search-grip sub-chain:

$$\begin{aligned} W_s(k+1) = W_s(k) - p_{g_1}[M_0 - W_g(k)]W_s(k) + p_{g_2}W_g(k)W_s(k) \\ + p_{g_1}[M_0 - W_g(k-T_g)]W_s(k-T_g)\Gamma(k;0) \end{aligned} \tag{4.11}$$

$$W_g(k+1) = W_0 - W_s(k+1) \tag{4.12}$$

In other words, Equation (4.11) tells us that the average number of robots in search state at any time is decreased by the robots transitioning to a gripping state and is increased by the robots coming back from a (failed or successful) collaboration. The number of sticks free (i.e., available) for gripping is equal to the total number of sticks M_0 minus those that are already being manipulated. Note that, in contrast to the previous system (i.e., the simplified search-avoidance case), the manipulation of environmental elements immediately introduces nonlinear coupling between the DEs.

Γ represents the fraction of robots that abandon the grip state after the time spent in the state exceeds their gripping time parameter τ_g and can be expressed as follows

$$\Gamma(k;T_{s_l}) = \prod_{j=k-T_g-T_{s_l}}^{k-T_{s_l}} [1 - p_{g_2}W_s(j)] \tag{4.13}$$

As explained more extensively in [68, 57], this is equivalent to calculating the probability that no other robot came "to help" during the time interval $[k-T_g-T_{s_l}, k-T_{s_l}] = [k-T_g, k]$ with $T_{s_l} = 0$.

Finally, our team metric (i.e., the average collaboration rate $\overline{C_t}$) can be computed from the number of successful collaborations C over the maximal number of iterations T_e:

$$C(k) = p_{g_2}W_s(k)W_g(k) \tag{4.14}$$

$$\overline{C_t} = \frac{1}{T_e}\sum_{k=0}^{T_e} C(k) \tag{4.15}$$

A. Steady-state analysis

Since the system of DEs (4.11) and (4.12) is nonlinear, it cannot be analyzed using z-tranforms, nevertheless, we use a method proposed by Lerman *et al.* in [57] and perform a steady-state analysis

in time domain.

By setting $W_s(i) = W_s^*$ and $W_g(i) = W_g^*$ for all i between $k - T_g$ and $k + 1$ in Equations (4.11)-(4.14), we obtain

$$0 = -p_{g_1}(M_0 - W_g^*)W_s^* + p_{g_2}W_g^*W_s^* + p_{g_1}(M_0 - W_g^*)W_s^*\Gamma^* \tag{4.16}$$

$$W_g^* = W_0 - W_s^* \tag{4.17}$$

$$\Gamma^* = (1 - p_{g_2}W_s^*)^{T_g} \tag{4.18}$$

$$C^* = p_{g_2}W_s^*W_g^* \tag{4.19}$$

First of all, we would like to know when the number of collaborations is maximized as a function of the robots in the search state (or equivalently, in the gripping state). In order to answer this question, we insert Equation (4.17) into Equation (4.19), it easily follows that C^* is maximal when $W_s^* = W_0/2$. By inserting this result, $p_{g_2} = R_g p_{g_1}$, and Equations (4.17) and (4.18) into Equation (4.16), we obtain the following equation:

$$0 = -\left(M_0 - \frac{W_0}{2}\right) + R_g\frac{W_0}{2} + \left(M_0 - \frac{W_0}{2}\right)\left(1 - p_{g_1}R_g\frac{W_0}{2}\right)^{T_g^{opt}} \tag{4.20}$$

Introducing $\beta = W_0/M_0$ and solving the equation for T_g^{opt} brings:

$$T_g^{opt} = \frac{1}{\ln\left(1 - p_{g_1}R_g\frac{W_0}{2}\right)}\ln\frac{1 - \frac{\beta}{2}(1 + R_g)}{1 - \frac{\beta}{2}} \tag{4.21}$$

Equation (4.21) represents a convex curve (i.e., an optimal value of T_g exists) as long as all the arguments of the logarithms are greater than zero. While this condition is always verified for the first logarithm in all our scenarios, the argument of the second logarithm depends on β and R_g. It can easily be shown [2] that an optimum exists if and only if

$$\beta < \beta_c = \frac{2}{1 + R_g} \tag{4.22}$$

For all other cases, the collaboration rate is a monotonically increasing and eventually saturating function of T_g. Figure 4.7 shows graphically the meaning of Equation (4.22). Equation (4.22) tells us that the bifurcation of the system (from the existence of an optimal T_g to no optimum at all) is a function of the collaboration parameter R_g. Note that, for instance, if the collaboration is very difficult (e.g., when R_g is very small), there could be situations where, although we have a larger number of robots than sticks (see Figure 4.7 left with 20 robots), the optimal collaboration rate

[2] While for $2/(1+R_g) \leq \beta \leq 2$ the second logarithm of Equation (4.22) does not exist, more generally, if $2/(1+R_g) \leq \beta$, W_s^* in Equation (4.16) will be always greater than $W_0/2$, i.e., a non optimal value. This is easily demonstrated by introducing $\beta = W_0/M_0$ in Equation (4.16), solving for β and comparing with Equation (4.22).

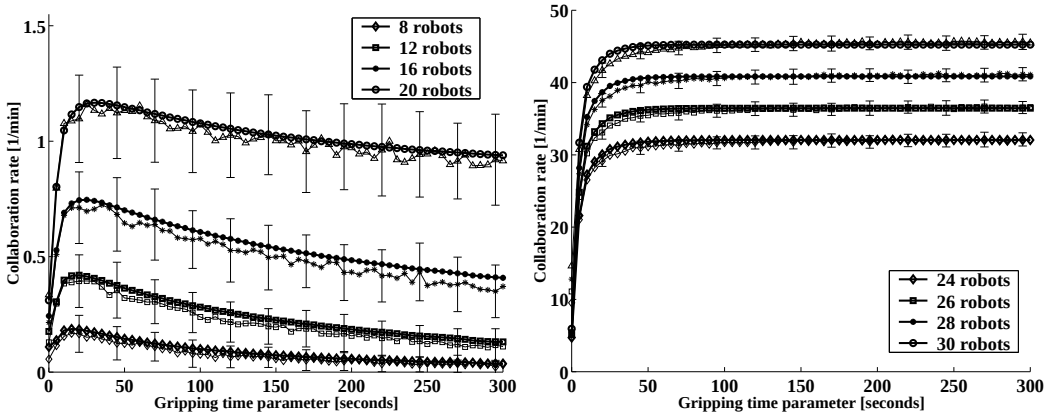

Figure 4.7: Graphical illustration of Equation (4.22) for an arena of 80 cm and 16 sticks (microscopic and macroscopic predictions overlapped). Left: $\beta < 2/(1 + R_g)$, with $R_g = 0.035$. Right: $\beta > 2/(1 + R_g)$ with $R_g = 0.7$.

may still be achieved only with a specific T_g. In other words, when it is difficult to collaborate, in order to enhance the number of collaborations, it is worth abandoning the sticks after a while and (probabilistically) increasing the critical mass working in another area of the arena. Although the precise team size at which the bifurcation happens in the real system cannot be correctly computed with Equation (4.22), this equation allows us to better situate intuitive considerations such as those presented in subsection 4.4.2.

B. Model and simulation results

Figure 4.8 shows a comparison between microscopic and macroscopic predictions of the collaboration rate for different arena and swarm sizes. Note that in this figure, the density of robots and sticks in each arena are invariant. A first striking result is the existence of two different system dynamics

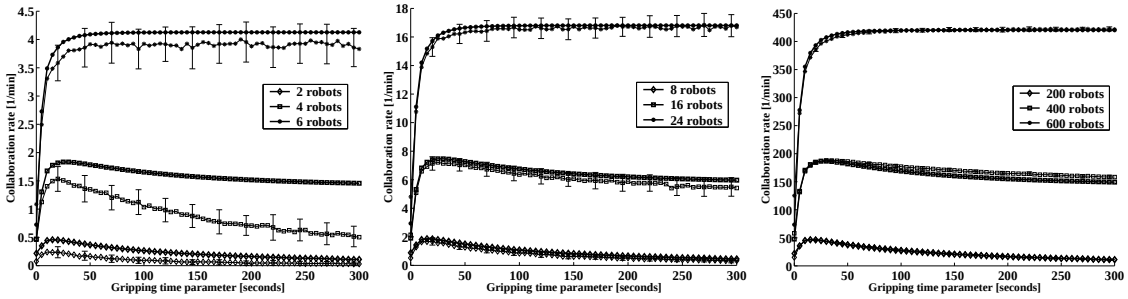

Figure 4.8: Comparing microscopic and macroscopic predictions of the collaboration rate. Left: Arena of 40 cm in radius. Center: Arena of 80 cm in radius. Right: Arena of 400 cm in radius.

as a function of the ratio between number of robots and number of sticks. When there are more robots than sticks, the collaboration rate increases monotonically with the gripping time parameter

and eventually saturates in a plateau corresponding to the optimal collaboration rate. In other words, under these conditions, it is a good strategy for a robot gripping a stick to wait a very long time for another robot to help, because there will always be at least one "free" robot available. By contrast, when there are fewer robots than sticks, waiting a very long time becomes a bad strategy, because all the robots waste time holding different sticks while no other robots are available to collaborate. For instance (this is an extreme case), an infinite gripping time parameter would lead to a null collaboration rate with all robots eventually holding a different stick permanently. Although this intuitive explanation is correct for the physical setup under consideration here, it is not generally correct since the sticks-to-robots ratio at which the system bifurcates is dependent on the R_g parameter, as we will show in the next subsection.

A second observation about Figure 4.8 is that, much like the 2-state example presented in subsection 4.4.1, the smaller the swarm size is, the larger the standard deviation is among the microscopic model simulation runs. However, here we notice that for very small swarm sizes and a smaller or equal number of robots than sticks (see Figure 4.8 left) macroscopic and microscopic models predictions do not agree qualitatively. Another, much smaller discrepancy can be observed between microscopic and macroscopic predictions with the swarm of 400 robots (see Figure 4.8 right). This is counter-intuitive since one would think that the problem shown in Figure 4.8 left would only have to do with the small number of robots and sticks used in the scenario, quantities not large enough to satisfy the law of the large numbers on which macroscopic models are based. A more careful inspection of the state variables in steady-state, much like what was done in subsection 4.4.1, tells us that the collaboration rate is not correctly predicted because, at the source, the steady-state variables used for calculating it are different (see Figure 4.9).

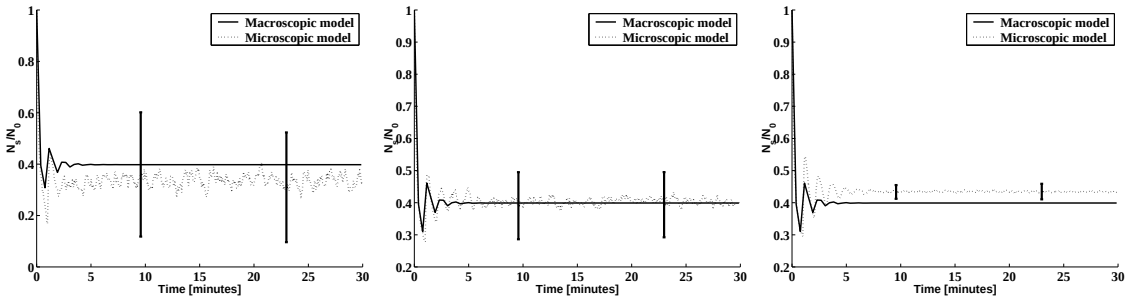

Figure 4.9: Comparing microscopic and macroscopic predictions of the number of robots in search state over time ($\tau_g = 60s$). Left: 4 robots, 40 cm arena. Center: 16 robots, 80 cm arena. Right: 400 robots, 400 cm arena.

In subsection 4.4.1 we showed that albeit with high standard deviation among runs and over time, the mean value of state variables of microscopic and macroscopic predictions coincide also for small teams, therefore something else appears to generate discrepancies here. As we will discuss in subsection 4.6.3, this has to do with the collaborative manipulation aspect of the stick-pulling

experiment, in particular in the way robots can compete for the same sticks in the two types of model.

4.5 Modeling the Full Stick-Pulling System

After introducing the two key sub-chains characterizing the dynamics of the stick-pulling experiment, we are now ready to analyze the full system model. As mentioned above, Figure 4.2 right shows the PSFM of the full system. The macroscopic model of the full system can be formalized as follows[3]:

$$\begin{aligned} W_s(k+1) = W_s(k) &- [\Delta_{g_1}(k) + \Delta_{g_2}(k) + p_w + p_R]W_s(k) + p_w W_s(k - T_a) + p_R W_s(k - T_{ia}) \\ &+ \Delta_{g_2}(k - T_{ca})W_s(k - T_{ca}) + \Delta_{g_2}(k - T_{cda})W_s(k - T_{cda}) \\ &+ \Delta_{g_1}(k - T_{cga})\Gamma(k; T_a)W_s(k - T_{cga}) \end{aligned} \tag{4.23}$$

$$\begin{aligned} W_a(k+1) = W_a(k) &+ \Delta_{g_1}(k - T_{cg})\Gamma(k; 0)W_s(k - T_{cg}) + \Delta_{g_2}(k - T_c)W_s(k - T_c) \\ &+ \Delta_{g_2}(k - T_{cd})W_s(k - T_{cd}) + p_w W_s(k) + p_R(k - T_i) \\ &- \Delta_{g_1}(k - T_{cga})\Gamma(k; T_a)W_s(k - T_{cga}) - \Delta_{g_2}(k - T_{ca}) \\ &- \Delta_{g_2}(k - T_{cda})W_s(k - T_{cda}) - p_w W_s(k - T_a) - p_R W_s(k - T_{ia}) \end{aligned} \tag{4.24}$$

$$W_i(k+1) = W_i(k) + p_R W_s(k) - p_R W_s(k - T_i) \tag{4.25}$$

$$\begin{aligned} W_c(k+1) = W_c(k) &+ \Delta_{g_1}(k)W_s(k) + \Delta_{g_2}(k)W_s(k) - \Delta_{g_1}(k - T_c)W_s(k - T_c) \\ &- \Delta_{g_2}(k - T_c)W_s(k - T_c) \end{aligned} \tag{4.26}$$

$$W_d(k+1) = W_d(k) + \Delta_{g_2}(k - T_c)W_s(k - T_c) - \Delta_{g_2}(k - T_{cd})W_s(k - T_{cd}) \tag{4.27}$$

$$W_g(k+1) = W_0 - W_s(k+1) - W_a(k+1) - W_i(k+1) - W_c(k+1) - W_d(k+1) \tag{4.28}$$

where $T_{xyz} = T_x + T_y + T_z$, W_s, W_a, W_i, W_c, W_d, and W_g represent the mean numbers of robots in the search, obstacle avoidance, interference, stick-centering, success dance, and gripping states respectively. The Γ-functions can be calculated with Equation (4.13) while the Δ-functions are defined as follows

$$\Delta_{g_1}(k) = p_{g_1}[M_0 - W_g(k) - W_d(k)] \tag{4.29}$$

$$\Delta_{g_2}(k) = p_{g_2} W_g(k) \tag{4.30}$$

[3] The system of DEs presented in this chapter is consistent with the description presented in previous publications [66] with one exception: here, the number of sticks available for grip1 has been decreased by a factor W_d since actually, as explained in subsection 4.2.3, the experimenter replaces the stick in the hole only when the success dance of the robot has terminated. Although the change in the collaboration rate with this correction is not noticeable the current model is more faithful to the real experiment.

Equations (4.23)-(4.28) can be interpreted in a similar way as Equations (4.11) and (4.12). For instance, Equation (4.23) tells us that the average number of robots in search state at any time is decreased by the robots which transition to a gripping state (grip1 and grip2) and those which start avoiding a wall or a teammate; W_s is increased by the robots that come back from a successful collaboration either as the "starter" or the "finisher" (the "finisher" of the successful collaboration will dance), those that come back from an unsuccessful collaboration, and those that finish a collision avoidance maneuver. It is worth noting that all the states other than the search and the gripping states are characterized by a W_s factor (either at the current iteration k or with a delay) since they are simple delays like those used in the 2-state system described in subsection 4.4.1. However, much like the simplified stick-pulling model described in subsection 4.4.2, the coefficients by which the state variable W_s is multiplied are time-varying and are also functions of other state variables or W_s itself (see Γ- and Δ-functions) but not constant like p_w or p_R. This is clearly strongly linked with the fact that the stick-pulling experiment (like the aggregation experiment) is concerned with a distributed manipulation process in which robots modify a shared environment as opposed to, for example, a simple passive sensing process.

4.5.1 Results

We can now compare the results of our microscopic and macroscopic abstractions with those gathered with lower implementation levels, i.e., embodied simulations and real robots. Real robot experiments lasted about 20 minutes (i.e, the autonomy of the on-board batteries) and were repeated three times while those using the embodied simulator lasted 30 minutes (simulated time) and were repeated ten times. As usual, error bars represent standard deviation among different runs with the same setup.

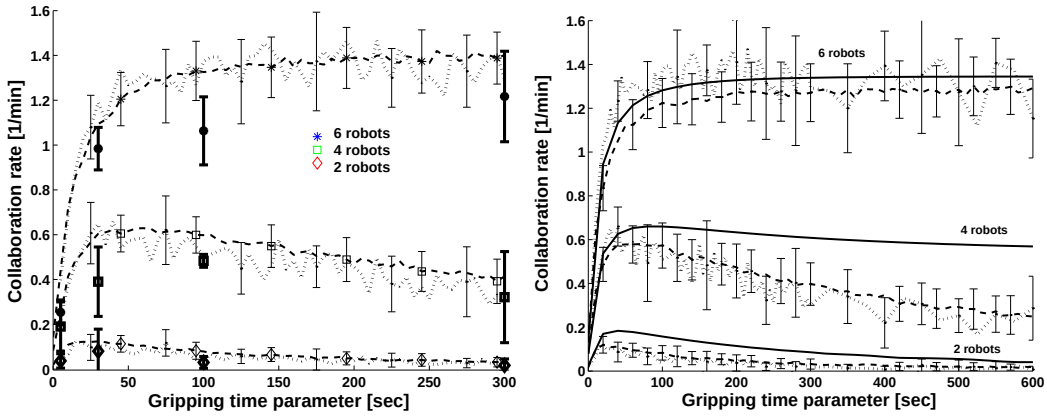

Figure 4.10: Collaboration rate as a function of the gripping time parameter (groups of 2, 4, and 6 robots, arena of radius 40 cm). Left: Real robots ($\tau_g = [5, 30, 100, 300]$ s), embodied simulations ($\tau_g = [0 : 5 : 300]$ s), and microscopic model ($\tau_g = [0 : 5 : 300]$ s). Right: microscopic model, macroscopic models, and embodied simulations ($\tau_g = [0 : 5 : 600]$ s).

Figure 4.10 left compares results obtained with real robots and embodied simulations. Although the real robot experiments have been repeated only three times instead of ten times as carried out with embodied simulations, real robots appear to achieve a slightly lower collaboration rate than their simulated counterparts. Differences between simulated and real robots gripper modules are at the origin of this discrepancy. We discuss this type of problem in detail in subsection 4.6.1.

Furthermore, Figure 4.10 right clearly shows that, in contrast to the macroscopic results, predictions delivered by the microscopic model are in good quantitative agreement with the data collected using the embodied simulator. Much like the results presented in subsection 4.4.2, these problems

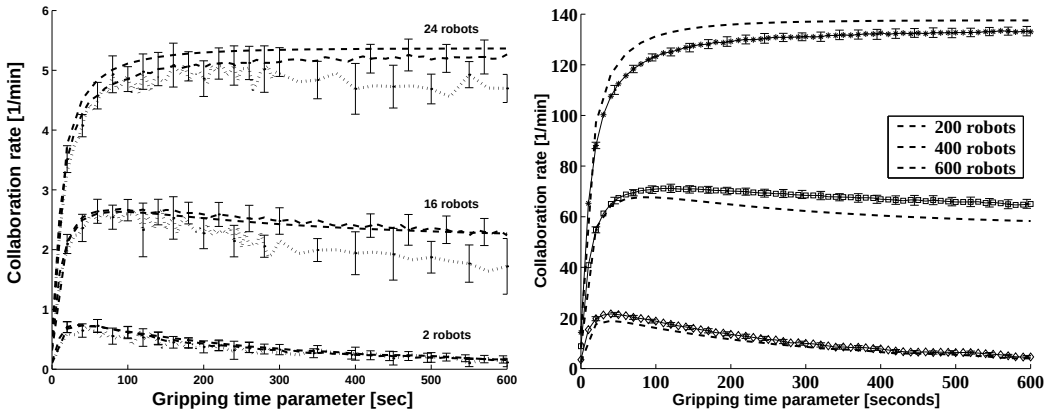

Figure 4.11: Left: Embodied simulations, microscopic model, and macroscopic model results (8, 16, and 24 robots, 16 sticks, and arena of 80 cm in radius). Right: Swarms of 200, 400, and 600 robots in an arena 400 cm in radius.

disappear as soon as we increase the number of robots and sticks by maintaining the same density of both items per unit area. Figure 4.11 left shows a good quantitative agreement among the three simulation levels simply by multiplying the robots and sticks quantities by four and without changing any implementation details. However, as soon as we multiply the quantities by 100 (see Figure 4.11 right), we notice that the problem mentioned in subsection 4.4.2 arises again, but even more accentuated. Subsection 4.6.3 gives a detailed explanation of the discrepancies between microscopic and macroscopic models.

4.5.2 Steady-state analysis

Like the simplified 2-state system described in subsection 4.4.2, the DEs describing the full system are nonlinear and therefore we have to perform the steady-state analysis in the time domain. Using the same method we adopted in subsection 4.4.2 for Equations (4.23) and (4.28) and the approximation in time domain introduced in subsection 4.4.1 for all the delay states of the full system (Equations

(4.24)-(4.27)), we obtain

$$0 = -p_{g_1}\left(M_0 - W_g^* - W_d^*\right) + p_{g_1}\left(M_0 - W_g^* - W_d^*\right)\Gamma^* + p_{g_2}W_g^* \tag{4.31}$$

$$W_a^* = T_a W_s^*[p_{g_1}\left(M_0 - W_g^* - W_d^*\right) + p_{g_2}W_g^* + p_w + p_R] \tag{4.32}$$

$$W_i^* = T_i p_R W_s^* \tag{4.33}$$

$$W_c^* = T_c W_s^*[p_{g_1}\left(M_0 - W_g^* - W_d^*\right) + p_{g_2}W_g^*] \tag{4.34}$$

$$W_d^* = T_d p_{g_2} W_s^* W_g^* \tag{4.35}$$

$$W_g^* = W_0 - W_s^* - W_a^* - W_i^* - W_c^* - W_d^* \tag{4.36}$$

The collaboration rate in steady-state becomes

$$C^* = p_{g_2}W_s^*W_g^* = p_{g_2}W_s^*\left(W_0 - W_s^* - W_a^* - W_i^* - W_c^* - W_d^*\right) \tag{4.37}$$

Simulating the DE system (4.23)-(4.28) long enough until a stationary regime is reached and solving the equation system (4.31)-(4.36) are two alternative options for obtaining the full system's steady-state vector $W^* = [W_s^*, W_a^*, W_i^*, W_c^*, W_d^*, W_g^*]^t$ (t stands for the transposition operator). Both of these operations have to be performed numerically. Indeed, if we try to solve the equation system (4.31)-(4.36) analytically by substitution and introducing all the results in Equation (4.31) we obtain the following transcendental equation.

$$P_0(W_s^*) + P_1(W_s^*)\Gamma^* + P_2(W_s^*)\Gamma^{*2} = 0 \tag{4.38}$$

where P_0, P_1, and P_2 are second order polynomials in W_s^* whose coefficients are functions of all the system parameters, i.e., $p_{ij} = p_{ij}(M_0, W_0, p_{g_1}, p_{g_2}, p_w, p_R, T_a, T_i, T_c, T_d)$, i being the polynomial index ($i = 0, 1, 2$) and j being the term number in a given polynomial ($j = 0, 1, 2$).

Like Equation (4.16), Equation (4.38) (a quadratic equation) can be solved for Γ^* and further for T_g. However, unlike the analysis described in subsection 4.4.2, determining the value of W_s^* that achieves an optimal collaboration rate implies to know what values the delay state variables assume in steady-state, as shown by Equation (4.37). Unfortunately, W_a^*, W_i^*, W_c^*, and W_d^* are in turn nonlinearly coupled with W_s^*, preventing us from finding an optimal value without solving Equation (4.38) for W_s^*. Although we can approximate Γ with a McLaurin series ($\Gamma^* \cong e^{-T_g p_{g_2} W_s^*}$), Equation (4.38) can be solved, to our knowledge, only numerically. This in turn prevents us from formulating analytical expressions for the system (possible) bifurcation points (such as that reported in Equation (4.22)) or superlinear-linear and linear-sublinear regime transitions (such as those reported in [48] for the relative collaboration rate, i.e., the collaboration rate normalized over the total number of robots used in the experiment).

4.6 Discussion

In section 4.5, we have presented results of the stick-pulling experiment obtained at different levels of implementation: real robots, embodied simulations, microscopic and macroscopic models. In this section, we discuss the problems and subtle effects that arise in moving from one (low) level of implementation to a more abstract level. We conclude this section by discussing the usefulness and some common limitations of both the modeling methodologies for robotic manipulation proposed in this chapter and in Chapter 3.

4.6.1 From real robots to embodied simulations

Although in general, as demonstrated in several other tasks [40, 41, 63, 64], the embodied, sensor-based simulator Webots provides very faithful results, simulation is never reality and several effects due to nonlinearities, noise, small heterogeneities among robots and components are simply neglected for reducing the computational cost of simulation. In particular, when a (collaborative) manipulation is performed with miniature robots like the Khepera, these effects can play a major role: grippers are usually endowed with a low number of degrees of freedom and sensors are affected by a high level of noise.

We therefore believe that the discrepancies between results obtained using the embodied simulator and those using real robots shown in Figure 4.10 left have to be attributed to differences in the real gripper module and its corresponding embodied simulation rather than to the different number of (experimental and simulation) runs carried out at the two implementation levels (e.g., three with real robots, ten with Webots). These differences between simulated and real grippers are twofold.

Reliability.— While the simulated gripper never releases the stick unintentionally (in other words, the trigger for releasing is always either a teammate's help or the internal timeout), we have observed that the real gripper sometimes drops the stick early due to the noise in measuring the elevation of the arm. Since the presence of a teammate is assessed through the measurement of the arm elevation sensor[4], if this measurement is noisy, the robot performing grip1 may believe that a teammate is helping and release the stick back into the hole. The error in elevation is correlated with the arm PI controller and is more likely when the arm is in full swing. Although the decision to release the stick is based on redundant sensory measurements, it still occurs from time to time with real robots but never with simulated robots (i.e., in Webots).

Simulated physical embodiment.— In Webots 2, collision detection routines are 2-D and grippers never get entangled (they actually pass through each other) as they may do in the real world. Possible fixes for these discrepancies could be to implement more realistic sensor noise in the gripper in Webots and use a newer version of the Webots simulator (version 3 or higher) which implements

[4] A systematic positive error on the pre-established position is detected by the first robot when the second robot is pulling the stick as well, see also subsection 4.2.3.

more time-consuming 3-D collision detection routines and should therefore eliminate the problem of gripper penetration.

4.6.2 From embodied simulations to microscopic modeling

A further step of abstraction is that which transforms an embodied, sensor-based agent into a much simpler agent. In the general microscopic model presented in this thesis (see section 2.4 in Chapter 2), an agent assumes a random position on the arena at each new iteration instead of being characterized by a trajectory, has a perfectly centered, uniform, and precise range of detection for each object it may encounter in the arena instead of being endowed with individual, heterogeneously distributed, noisy sensors around its body, is characterized by an average speed instead of having a more complex kinematic controller, with, for instance, accelerations tightly coupled with sensory readings, has a behavior-based controller which can be represented by a precise PFSM whose state-to-state transitions follow either precise durations or are triggered by external events but never happen with shorter or longer duration because of control interruptions or because an event has been detected by the sensors earlier or later. In the next three subsections we will discuss the role of some of these important approximations in the accuracy of the predictions delivered by the microscopic model.

A. Object distribution and environmental symmetry

The modeling methodologies for distributed manipulation in robotic systems presented in this thesis assume a uniform distribution of objects on the arena and no relevance of trajectories on the chosen metrics. As long as there are no overlapping detection areas between the objects (e.g., walls, seeds, and robots in the aggregation experiment and wall, sticks, and robots in the stick-pulling experiment), this assumption is correct.

In order to verify this, we considered the stick-pulling experiment and ran several experiments characterized by different stick distributions (see Figure 4.12) using the embodied simulator. In each of these cases, the predictions of both the microscopic and embodied simulator models were as accurate as those shown in Figure 4.10 right, the microscopic model reaching quantitative agreement with the embodied simulations. The current microscopic modeling methodology could be also easily adapted to take into account overlapped detection areas. This should be particularly easy for immobile objects. However, subtle effects due to robot clustering and mutual influence in search and manipulation activities could arise in high density scenarios and, as we will show in subsection 4.6.2, these effects are more difficult to incorporate in the models.

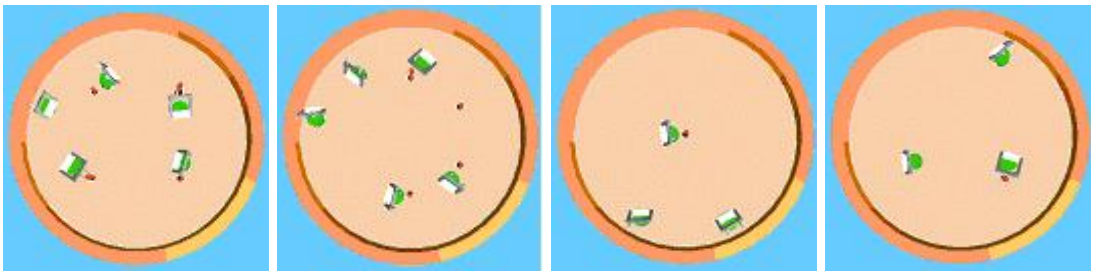

Figure 4.12: Four examples of implemented stick distributions.

B. Model parameters and behavioral granularity

Summarizing a belt of noisy (not necessarily homogeneously distributed sensors) as well as their relative detection and reactive control algorithms with a few simple parameters is not trivial. In most of the models presented in this dissertation, we use the method based on simple geometrical considerations and systematic experiments with one or two real robots. The advantage of the method used to calculate the robots' rates of encounter with objects of different types in the arena (see subsection 2.5.3) is that each type of object is characterized by its own encounter rate, this helps to avoid the use of the smallest object present on the arena to calculate the time granularity of the of system as we did in this chapter. This method, if adapted to the microscopic discrete-time model, may allows us to avoid rounding the durations of the different robot maneuvers (e.g., collision avoidance, seed/stick grasping, success dance) measured in real time units and to choose an arbitrarily small time granularity although more work needs to be done in this direction. For instance, using this approach it is not clear how to calculate the rate at which the robots encounter the wall, in the experiment of [58] the authors using this method ignored the problem by simply neglecting the boundary conditions.

Further difficulties may arise because of the behavioral granularity captured in the microscopic model. For instance, the controllers of the robots used in the aggregation and stick-pulling case studies can be approximated as a FSM though certain routines (obstacle avoidance and interference) have been implemented with proximal controllers. Proximal controllers (in our case, neural network-based controllers) are controllers that tightly couple actuators with sensors without passing through a distal representation (e.g., case of behavior-based implementations). The parameters used to describe the states corresponding to such routines (in our case, the duration of obstacle avoidance and interference as well as the probability of detecting an obstacle and a teammate) can still be measured in systematic tests with one or two robots, as mentioned above, even if this implies some inaccuracy. To predict the quantitative dynamics of the variables of the systems under investigation in this thesis (e.g., number of active workers for the aggregation task and the collaboration rate for the stick-pulling task) this approximation is quite sufficient. Nonetheless, we expect that a PFSM

will be a poor candidate for describing some robotic systems (e.g., systems characterized by proximal controllers coupled learning mechanisms) as it will only offer an imprecise description of such systems (thus, possibly weaken the current modeling methodologies).

C. Overcrowded arenas

The proposed modeling methodologies achieve quantitatively correct predictions based on the assumption that robots are homogeneously distributed in the arena (equivalently, all locations on the arena are equally likely to be occupied by any robot). As soon as this assumption no longer holds, such as in an overcrowded scenario, the current methodologies reach their limitations and the predictions are no longer quantitatively correct. Another way to explain why in an overcrowded scenario the methodologies fail is that the models as generated by the methodologies (i.e., using solely geometrical considerations) are no longer Markovian in this situation. As the arena becomes crowded with robots (although the robots may still be able to move about and manipulate seeds/sticks) the future transitions between states of the FSMs are contingent on each robot's trajectory and position (and to some extent on the trajectories and the positions of the surrounding robots). For instance, the robots would form clusters (with overlapping detection areas), thus, freeing space around seeds (resp. sticks), which in turn can be exploited by other robots for successful seed manipulation (resp. stick-pulling collaboration). In order to be maintain the Markovian property of the system in this situation, we should consider also states relative to trajectories and positions with an obvious, consequent, net increase of the complexity of the models, if formulated in the same way.

Figure 4.13 left illustrates this effect for the microscopic model applied to the stick-pulling experiment. For each group size we plotted the collaboration rate achieved after optimization of τ_g (using systematic search). The plot shows a clear discrepancy between the predictions of the microscopic model and the embodied simulator of the optimal collaboration rate for group sizes greater than ten robots. While the embodied simulation results reflect this continuing collaboration mitigated by crowding, the microscopic model predicts ever-increasing performance. The continuous increase is due to the fact that the interference area represented by the robot cannot extend over the detection areas of the wall and sticks. An alternative option which allows the robots' detection area to grow until the probability of encountering another robot is one will equally fail to deliver correct predictions since at a certain point the collaboration rate becomes null [48, 66] (collaboration is taken over by intermittent interference maneuvers) whereas in the embodied simulation it continues to be different from zero, as shown in Figure 4.13 left.

Figure 4.13 right shows the problem from another more general perspective. The piecewise linear approximation we do in the model for the probability of encountering another robot per time step (and similarly, the rate at which a given robot encounter another robot) is also about correct for different group sizes until ten robots then, in the saturation phase, the curve diverges from that

measured using embodied simulations.

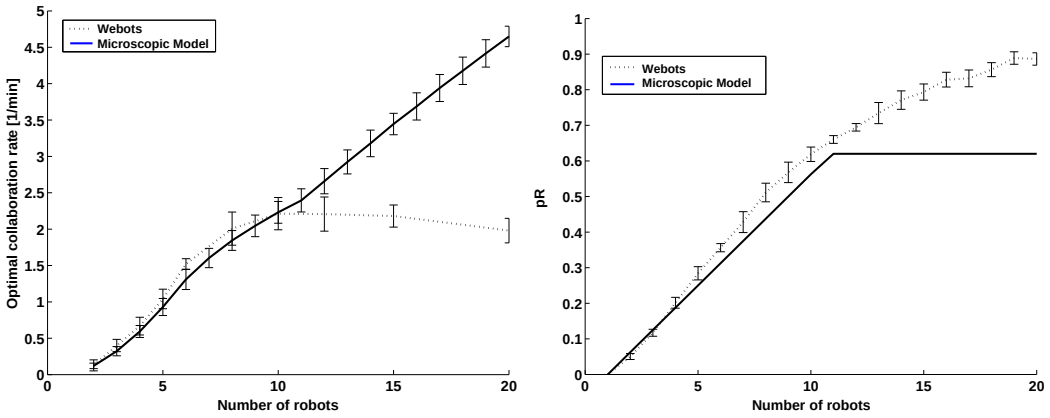

Figure 4.13: Left: Comparing embodied simulations and microscopic model predictions in an overcrowded arena (up to 20 robots in an arena of 40 cm in radius). Right: probability of encountering a robot in the embodied simulator compared with the linear approximation of macroscopic model.

4.6.3 From microscopic to macroscopic models

As seen in subsections 4.4.2 and 4.5.1, the microscopic and macroscopic models, although they rely fundamentally on the same abstraction of the individual agent, may deliver slightly divergent predictions. The discrepancies are due to the manipulation involved in the stick-pulling experiment, to the metric used for evaluating the swarm performance, and more importantly to the way the macroscopic model describes the dynamics of the system. The essence of the problem is combinatorial: the microscopic model correctly calculates the allocation of robots at limited, shared manipulation sites by keeping track of each modification of the environment generated by the actions of the robots. The macroscopic model instead summarizes the states of the robotic swarm and the objects present in the arena (e.g., seeds, sticks) with average quantities. In fact, the proposed macroscopic models of robotic systems in this thesis, describe the dynamics of the averages of the variables of the system as opposed to the dynamics of the variables distinctively. As a consequence, the macroscopic models tend to agree better with the microscopic model and embodied simulations when the system is linear (or at least does not present a high degree of nonlinearity in the robots' interaction with each other and with the environment). For instance, in the simplified models presented in subsections 3.4.1 and 3.4.2 in Chapter 3 and that in subsection 4.4.1 in this chapter, the macroscopic models agree very well with the microscopic model and the embodied simulator because the dynamics of the variables of interest are linear.

In the following paragraphs, we explain the discrepancies between the microscopic and macroscopic models predictions of the outcome of stick-pulling experiment, nevertheless, to some extent, the reasons detailed below are also valid for the macroscopic model presented in Chapter 3.

Let us consider again Figure 4.8, Figure 4.10 right, and Figure 4.11. A first observation tells us that the discrepancies are not due to the additional delay states in the full system since the problem exists in both the simplified and the full system models. As a matter of fact, the main cause for the discrepancies between microscopic and macroscopic models is the nonlinear coupling between equations due to stick manipulation. A second observation tell us that the problem appears to be more important in the full system model than in the simplified one.

Let us consider, for instance, the second term on the right hand sides of Equations (4.11) and (4.23), i.e., the number of sticks available for grip1. In the simplified model, the number of sticks $M(k)$ available for grip1 is $M(k) = M_0 - W_g(k)$ while in the full system model this number is $M(k) = M_0 - W_g(k) - W_d(k)$. The microscopic model considers all the robots belonging to the swarm sequentially at each iteration and each robot's actions on the environment are updated at that iteration. For instance, the first robot in the swarm at iteration k has a probability $p_{g_1} M(k)$ of finding and gripping a stick. When it is the next robot's turn, if the first robot has found and gripped one of the sticks, the following robot's probability of finding a free stick is $p_{g_1}(M(k)-1)$, and so on. More generally, the probability of any robot finding and gripping the jth stick out of $M(k)$ sticks given that $(j-1)$ sticks have already been found by $(j-1)$ different other robots at iteration k is given by $p_{g_1}(M(k)-j+1)$. It is worth noting that the robots that can find a stick for grip1 are only those in the search state, i.e., $W_s(k)$. At the macroscopic level, we try to estimate how many of the $W_s(k)$ robots on average will actually find one of the available $M(k)$ sticks. We can calculate this average number with a binomial distribution X characterized by the parameters $p_{g_1} M(k)$ and $W_s(k)$. Equation (4.39) describes the probability that i sticks are found and gripped at iteration k, i.e., $P(X = i)$. However, the expected value of this binomial distribution is twofold and depends on the ratio between $W_s(k)$ and $M(k)$. Equation (4.40) expresses the expected value of the binomial distribution when the number of searching robots is less than or equal to that of the available sticks, i.e., $W_s(k) \leq M(k)$. Equation (4.41) expresses the expected value of the binomial distribution when the number of searching robots exceeds the number of available sticks, i.e., $W_s(k) > M(k)$.

$$P(X = i) = \binom{W_s(k)}{i} \left(p_{g_1} M(k)\right)^i \left(1 - p_{g_1} M(k)\right)^{W_s(k)-i} \tag{4.39}$$

$$E(X_1) = \sum_{i=0}^{W_s(k)} i \binom{W_s(k)}{i} \left(p_{g_1} M(k)\right)^i \left(1 - p_{g_1} M(k)\right)^{W_s(k)-i} = p_{g_1} M(k) W_s(k) \tag{4.40}$$

$$E(X_2) = \sum_{i=0}^{M(k)} i \binom{W_s(k)}{i} \left(p_{g_1} M(k)\right)^i \left(1 - p_{g_1} M(k)\right)^{W_s(k)-i} \tag{4.41}$$

In order to avoid discontinuities in the system description and to allow closed form expressions ($E(X_2)$ does not have a closed form), in our macroscopic model we use the same expected value in all the cases, i.e., $E(X_1)$, independently of the ratio between $W_s(k)$ and $M(k)$ (see Equations (4.11)

and (4.29)). Thus, depending on the scenario in which the stick-pulling experiment takes place, the approximation of $E(X_2)$ with $E(X_1)$ happens more or less often, generating as a consequence more or less discrepancies between the predictions of microscopic and macroscopic models of the collaboration rate. Moreover, in the simplified model, since $W_s(k) + W_g(k) = W_0$ (robots' conservation law) and $C(k) = p_{g_2} W_s(k) W_g(k)$ (see Equation (4.14)), the error introduced when computing $W_s(k)$ is largely compensated by the error of opposite sign introduced in the value of $W_g(k)$, thus $C(k)$ varies little from its actual value due to a small error. This explains why there is only a minor difference between the microscopic and macroscopic predictions in the simplified model related to the full system model.

To confirm this theoretical explanation, we now provide a simple numerical example. As shown in Equations (4.11) and (4.29), the second term on the right hand sides of the equations is weighted by the probability p_{g_1}. If this term is small, according to the explanation above the error introduced should be attenuated. Table 4.3 illustrates the relative error introduced by approximating $E(X_2)$ with $E(X_1)$ when $W_s(k) > M(k)$ for different values of the probability p_{g_1} and the number of available sticks $M(k)$.

Table 4.3: Influence of the value of p_{g_1} on the approximated value of $E(X)$.

p_{g_1}	Number of available sticks ($M(k)$)	$E(X_1)/E(X_2)$
0.01	1	0.9393
0.01	2	0.9925
0.1	1	0.0805
0.1	9	0.9949

Figure 4.14 further illustrates the explanation above: the disagreement between the microscopic and macroscopic predictions shown in Figure 4.10 right for teams of 2 and 4 robots in an arena of radius 40 cm disappears when the same robot teams are placed in a larger arena of radius 120 cm (thus with a reduced value of p_{g_1}). Similarly, the disagreement between the two models shown in Figure 4.11 right for teams of 200, 400, and 600 robots in an arena of radius 400 cm progressively disappears when the same team is placed in arenas of radii 1200 cm and 5000 cm, respectively.

We did not get similar (large) discrepancies with the macroscopic model of the aggregation experiment because the seeds-to-robots ratio that would trigger such disagreement cannot be violated. This is simply due to the fact that in the aggregation experiment the total number of robots has to be inferior to the number of seeds to prevent the scenario where all the seeds are picked up and cannot be dropped next to any other seeds from happening. In that case, only the nonlinearity of the system could constitute a potential limitation to the proposed models.

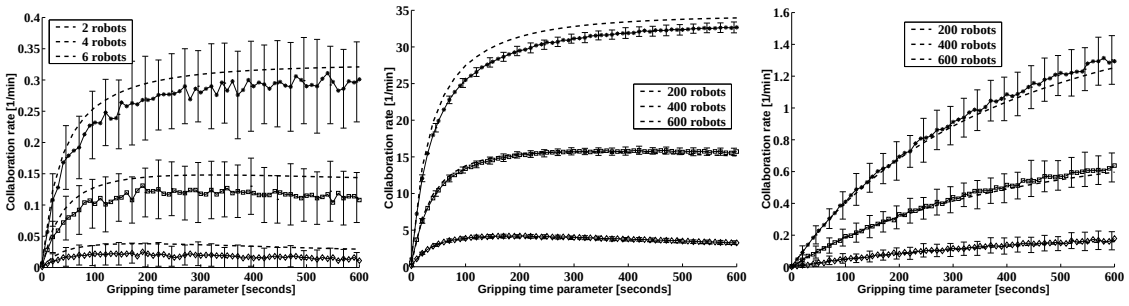

Figure 4.14: Comparing microscopic model and macroscopic model predictions for small p_{g_1} probabilities. Left: 4 sticks, arena of 120 cm in radius. Center: 400 sticks, arena of 1200 cm in radius. Right: 400 sticks, arena of 5000 cm in radius.

4.6.4 Modeling as an optimization tool

The modeling methodology we presented in this chapter can be useful also for optimization purposes. First, like the models presented in the previous chapter, those presented here also deliver results in time lapses that are at least four orders of magnitude shorter than a corresponding embodied simulation. Second, in swarm systems, usually based upon the principles of self-organized collective behavior of multiple interactions among individuals and between individuals and the environment, low level system parameters such as the sensory configuration, body shape, proximal control parameters characterizing reactive behavior, and individuals' heading and positions, may also play an important role, but this is often neglected when we try to optimize the system at a high abstraction level. In other words when searching for optimal solutions, better results may be achieved by searching in a large parameter space (possibly including even the low-level hardware and software parameters) than in a small parameter space. Two concrete examples illustrate this point.

A first example is concerned with collective plume tracing by a group of robots [41]. In the experiment, due to the highly stochastic nature of robot-plume interactions, the unavailability of suitable plume models for the specific task conditions, and the relevance of the robots' heading and positions, the authors were only able to create models of the whole system characterized by several free parameters (see [42] for more details). On the other hand, they were much more successful in system optimization with a *ad hoc* reinforcement learning algorithm (which more easily allows searching in a large parameter space) combined with embodied simulations [43].

A second example comes from additional experiments performed in the framework of the stick-pulling experiment. As Ijspeert *et al.* [48] have shown using systematic search (with Webots), and Li *et al.* [59] using a simple learning algorithm combined with the same microscopic model as that presented here, a homogeneous swarm may not necessarily achieve the best performance. Subswarms of specialists in grip1 and grip2 may, for instance, outperform an optimized homogeneous swarm, depending on the sticks-to-robots ratio. At the macroscopic level, either we know the

numbers of castes in advance and we introduce a new set of DEs (and similarly, DDEs) for each new type of agent involved in the system, or it would be impossible to explore heterogeneous (even less likely, emergent) solutions. Therefore, to explore heterogeneous solutions, in particular those with non pre-established heterogeneity, microscopic models combined with machine-learning algorithms appear to be a more efficient solution.

4.7 Conclusion

In this chapter, we have presented a methodology for generating macroscopic models of a swarm robotic system. The methodology was explicitly designed for distributed manipulation experiments and therefore we have supported the discussion with a case study belonging to this class. The case study is that of stick-pulling consisting of using a group of autonomous robots to pull out sticks standing in holes scattered around an enclosed arena. Each stick (because of it length) requires the collaborative effort of exactly two agents to be removed from the ground.

We have shown that the proposed models can not only deliver quantitatively correct predictions in time periods at least four order of magnitude shorter than other popular simulation tools, such as a sensor-based, embodied simulator, but also allow for better understanding of the system properties, and help to gain counter-intuitive insights.

In a general study that applies to both the modeling methodologies presented in this chapter and that presented in Chapter 3, we have discussed subtle numerical effects and prediction discrepancies, tested the fundamental modeling assumptions, and reviewed the intrinsic limitations of the proposed modeling methodologies.

Abstraction in general allows researchers to understand the role of key system parameters, generalize, and analyze underlying principles and sometimes even enable mathematical tractability and re-situate intuitive considerations, as we have shown in subsections 4.4.1 and 4.4.2. Optimization processes conducted in large parameter spaces are often more fruitful and accurate than those conducted in small parameter spaces. This further explains why high level abstract models (like those presented in this chapter and the previous one) that allow search of optimal solutions for systems consisting of a large number of parameters are relevant.

Finally, it appears that even if we are able to produce quantitatively correct analytical models, nonlinearities and complexity of a real system often prevent us from going further in the analysis (and indirectly in the optimization process) with the mathematical tools currently available.

Chapter 5

Conclusion

This thesis contributes to research in swarm intelligence, a new computational and behavioral metaphor inspired by nature. The swarm intelligence approach implies distributed control strategies within large groups of simple agents (e.g., minimalist robots and sensors) and relies solely on agent-agent and agent-environment interactions to provide robust, flexible, goal-oriented, and sometimes emergent collective behaviors.

5.1 The Contributions

The main objective of this work is to strengthen the bridge between experimental work in biology (and modeling of natural systems) and the artificial systems inspired by the natural systems, by offering a set of mathematical tools for studying the artificial systems. The specific contributions of this thesis are a series of modeling methodologies for investigating two classes of experiments within the swarm intelligence framework: distributed sensing (in swarm-based sensor networks) and distributed manipulation (in swarm-based robotics).

In the first class of experiments, we have shown that artificial swarm sensor systems, in general, and event discovery and information spread in those systems, in particular, present similar characteristics as those of some natural systems and processes such as the spread of genes or infectious diseases in animal populations and scouting and foraging in social insect colonies. Thus, as a novelty, we have cast distributed sensing (in artificial systems) in the well-studied framework of human populations and social insects dynamics and we have shown that the artificial systems can be described mathematically in similar manners as the natural systems. We have considered two categories of models: embodied and non-embodied models respectively. While the embodied models take into account the physical characteristics of the agents, the non-embodied models consider the agents as mobile bodiless points. Moreover, within each category of models we have distinguished two subcategories: spatial and nonspatial models. The spatial models make use either of the trajectories of the agents, the correlation between the positions occupied over consecutive time steps, or the spatial

distribution of the agents resulting from their movement pattern (and the boundary conditions of the arena). A main contribution of this section of the thesis resides in the fact that, as opposed to models of natural systems where the main goal is to obtain models that agree with observation or experiments, our models were designed such that they can be used not only for investigating the influence of the different parameters of the system on its behavior, but also for optimization purposes. For instance, given some particular characteristics of a mobile swarm sensor network (e.g., an enclosed arena, a network size, spatial distribution of the sensors) what is the (optimal) control parameter (e.g., communication range) that maximizes information dissemination in the network? Finally, we have validated most of the proposed models with simulation results.

The second class of experiments is concerned with distributed manipulation in swarm robotics. In this latter class, we have investigated two case studies of non-collaborative and collaborative distributed manipulation respectively and have proposed two incremental methodologies for modeling the systems. We have formalized the following general approach: first the swarm-based (distributed manipulation) system is represented by a finite state machine (FSM) then, when the state transitions of the FSM occur probabilistically, the dynamics of the variables of the systems are described mathematically using discrete-time models, when the transitions occur at given rates, continuous-time models are used. We have completed our work in this last research axis by analyzing the weaknesses and limitations of the proposed methods and the assumption upon which both have been constructed. Here, using sensor-based simulation results or real robots experiment results, we have illustrated the accuracy and the computational efficiency (i.e., much higher speed ratio) of the proposed models in predicting the overall behavior of the systems.

Finally, another originality of this thesis is that throughout this work, we have defined a clear way of calculating the different parameters of each proposed model using solely the characteristics of the system: e.g., radius of the arena, speed, sensor range, and communication range of the agents.

5.2 Future Issues in Modeling Artificial Swarm Systems

This work is concerned with facilitating the transition between the study of social insects and the application of the study to design artificial systems such as swarm sensor networks and swarm-based robotics. Nonetheless, as far as modeling artificial swarm systems is concerned, some questions remains unanswered and constitute areas for future investigations.

5.2.1 Time discretization and behavioral granularity

One such problems is concerned with the calculation of the granularity of the discrete-time models. In fact, in the proposed discrete-time models the transition from one state to another happens with a probability (as opposed to a fixed rate in continuous-time models). Hence, the time step of the

discrete-time model is in number of iterations as opposed to time units in the continuous-time model. Unfortunately, as detailed in section 4.6.2, more investigation is required to find the most efficiency way of calculating the time step of the discrete-time models such that an appropriate level of detail (about the system) is always taken into account without deteriorating the computational efficiency of the models.

5.2.2 Moving from one level of implementation to another

In section 4.6 we have discussed several difficulties that may arise when moving from a (low) level of implementation (e.g., real robots and embodied simulations) to a more abstract (high) level (e.g., microscopic and macroscopic models) and vice versa. It is worth noting that all the abstraction steps were hand-coded, exploiting a combination of engineering, heuristic, and systematic tests. Although we have not compared the (microscopic and macroscopic) modeling predictions with those obtained at lower levels of implementation using statistical tests (e.g., as done in [63, 65] using a Mann-Whitney test), we do not believe that we will gain any further insight into how to resolve the difficulties in moving from one level of implementation to another in this way. However, statistical tests combined with algorithms that can systematically explore possible (abstraction) options and, at the same time, verify the (statistical) influence of each of them on a chosen metric may allow to achieve fundamental breakthroughs in the modeling methodology. First attempts in this direction, for different purposes, have been proposed by Goldberg and Matarić [33] (automatic generation of behavioral FSM for robotic coordination) and Zhang *et al.* [106] (automatic generation of Probability Density Functions for automatic design of sensory configurations).

5.2.3 Models that take into account more complex features

Other future challenges in modeling artificial swarm systems are concerned with overcoming of current limitations of and extending the current modeling methods to other classes of experiments and systems in order to create a general methodology applicable to all swarm-based systems. Boundary conditions (e.g., overcrowded scenarios), non-behavior-based controllers (e.g., proximal, neural controllers), more sophisticated individual capabilities (e.g., learning, peer-to-peer wireless communication, and navigation capabilities), complex distributed sensing tasks (e.g., mapping, searching), and more complex systems (e.g., mixed human-robot societies) are examples of problems that still need to be addressed in modeling artificial swarm systems. As an example of a simple modeling approach that takes into account more complex features of a system, we think that the correlated random walk model can be used to fully describe the movement pattern of reactive embodied agents, for instance, by approximating the step length distribution function with measurement of the distance traveled by the agents between consecutive obstacle encounters. This will provide another faster tool than the sensor-based simulator for the study of complex swarm-based systems of mobile

embodied agents.

5.2.4 Trend in computing: toward massively distributed systems

Finally, in a more general note, it appears that the forty-year trend in the computing industry is away from centralized, high unit cost, low unit volume products toward distributed, low unit cost, high unit volume products. We expect the next step in this process to be the emergence of massively distributed or swarm-based systems as the technology becomes increasingly available. The size of these networks and the diversity of the applications to handle will necessitate the use of statistically correct rather than deterministic algorithms for resource accounting, fault detection and correction, and system management in general. As a consequence, new complex issues will arise such as system heterogeneity, diversity in object representation and resource management, synchronization, etc., that the modeling approaches will have to take into account, although each represents a potential challenge to modeling.

Bibliography

[1] F. R. Adler and D. M. Gordon. Information Collection and Spread by Networks of Patrolling Ants. *The American Naturalist* 140, No. 3, 1992, pp. 373-400.

[2] W. Agassounon, A. Martinoli, and R. M. Goodman. A Scalable, Distributed Algorithm for Allocating Workers in Embedded Systems. *Proc. of the IEEE Conf. on System, Man and Cybernetics*, Tucson, AZ, Oct. 2001, pp. 3367-3373.

[3] W. Agassounon and A. Martinoli. Efficiency and Robustness of Threshold-Based Distributed Allocation Algorithms in Multi-Agent Systems. *Proc. of the First Int. Joint Conf. on Autonomous Agents and Multi-Agents Systems*. Bologna, Italy, July 2002. pp. 1090-1097.

[4] W. Agassounon and A. Martinoli A. A Macroscopic Model of an Aggregation Experiment using Embodied Agents in Groups of Time-Varying Sizes. *Proc. of the IEEE Conf. on System, Man and Cybernetics*, Hammamet, Tunisia, Oct. 2002. IEEE Press.

[5] W. Agassounon. Distributed Information Retrieval and Dissemination in Swarm-Based Networks of Mobile, Autonomous Agents. *IEEE Swarm Intelligence Symposium*. Purdue University, Purdue, IN, 2003. To appear.

[6] W. Agassounon, A. Martinoli, and K. Easton. Threshold-Based, Scalable, Distributed Allocation in Swarm Robotics: A Case Study in Distributed Manipulation. *Autonomous Robots Journal*, 2003. Submitted.

[7] A. Asama, M. K. Habib, I. Endo, K. Ozaki, A. Matsumoto, and Y. Ishida. Functional Distribution among Multiple Mobile Robots in an Autonomous and Decentralized Robot System. *Proc. of IEEE Int. Conf. on Robotics and Automation*, 1991, pp. 1921-1926.

[8] N. T. J. Bailey. The Mathematical Theory of Infectious Diseases and its Applications. Griffin, London, 1975.

[9] T. Balch and R. C Arkin, Communication in Reactive Multi-Agent Robotic Systems. *Autonomous Robots*, **1**, 1994, pp. 1-25.

[10] R. Beckers, J.-L. Deneubourg, S. Goss, and J. M. Pasteels. Collective Decision Making through Food Recruitment. *Insectes Sociaux*, **37**, 1990, pp. 258-267

[11] R. Beckers R., Holland O. E., and Deneubourg J.-L. From Local Actions to Global Tasks: Stigmergy and Collective Robotics. In Brooks R. and Maes P., editors, *Proc. of the Fourth Workshop on Artificial Life*, Boston, MA, July, 1994, The MIT Press, pp. 181-189.

[12] G. Beni and J. Wang. Swarm Intelligence. *Proc. of the Seventh Annual Meeting of the Robotics Society of Japan*, Tokyo, Japan, 1989, pp. 425-428.

[13] A. Billard, A. J. Ijspeert, and A. Martinoli. Adaptive Exploration of a Frequently Changing Environment by a Group of Communicating Robots. *Proc. of the Fifth European Conf. on Artificial Life*, Lausanne, Switzerland, 1999, pp. 596-605.

[14] 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. Special Issue on Adaptive Robots, Torras C., editor, *Connection Science*, Vol. 11, No. 3/4, 1999, pp. 359-379.

[15] E. Bonabeau, G. Theraulaz, and J.-L.Deneubourg. Fixed Response Thresholds and Regulation of Division of Labour in Insect Societies. *Bulletin of Mathematical Biology*, 1998, Vol. 60, 753-807.

[16] E. Bonabeau, M. Dorigo, and G. Theraulaz. Swarm Intelligence: from Natural to Artificial Systems. SFI Studies in the Science of Complexity, Oxford University Press, New York, NY, 1999.

[17] P. Bovet and S. Benhamou. Spatial Analysis of Animals' Movements using a Correlated Random Walk Model. *Journal of Theoretical Biology*, **131**, 1988, pp. 419-433.

[18] A. Cai, T. Fukuda, and F. Arai. Information Sharing among Multiple Mobile Robots for Co-operation in Cellular Robotic System. *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robotic Systems*, Grenoble, France, 1997, pp. 1768-1773.

[19] R. Chauvin and P. Janin. Facteurs de Direction et d'Excitation au cours de l'Accomplissement d'une Tâche chez *Formica polyctena*. *Insectes Sociaux*, Vol. 22, 1975, pp. 199-206.

[20] L. Chrétien. Organisation Spatiale du Matériel provenant de l'Excavation du Nid chez Messor Marbarus et des Cadavres chez Lasius Niger (*Hymenopterea : Formicidea*). PhD Dissertation, Université Libre de Bruxelles, 1996.

[21] M. L. Cody. Finch flocks in the Mohave Desert. *Theoretical Population Biology*, **2**, 1971, pp. 142-158.

[22] D. W. Crumpacker and J. S. Williams. Density, Dispersion and Population Structure in *Drosophila Peudoobscura*. *Ecol Mon.*, **43**, 1973, pp. 498-538.

[23] A. K. Das, R. Fierro, V. Kumar, J. P. Ostrowoski, J. Spletzer, and C. J. Taylor. A Vision-Based Formation Control Framework. Special Issue on Advances in Multi-Robot Systems, T. Arai, E. Pagello, and L. E. Parker, editors, *IEEE Trans. on Robotics and Automation*, Vol. 18, No. 5, pp. 813-825.

[24] A. Drogoul and J. Feber. From Tom Thumb to the Dockers: Some Experiments with Foraging Robots. *Proc. of the* 2^{nd} *Int. Conf. on Simulation of Adaptive Behavior*, From Animals to Animats 2. 1993, pp. 451-459.

[25] G. Dudek, P. Freedman, and S. Hadjres. Mapping Unknown Graph-Like Worlds. *Journal of Robotics Systems*, **13**, 1996, pp. 539-559.

[26] D. B. Dusenberry. Ranging Strategies. *Journal of Theoretical Biology*, **139**, 1989, pp. 309-316.

[27] K. Easton and A. Martinoli. Efficiency and Optimization of Explicit and Implicit Communication Schemes in Collaborative Robotics Experiments. *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, Lausanne, Switzerland, Sept.-Oct. 2000, pp. 2795-2800.

[28] R. A. Fisher. The Wave of Advance of Advantageous Genes. *Ann. Eugen.*, **7**, 1937, pp. 353-369.

[29] J. Fredslund and M. J. Mataric. General Algorithm for Robot Formations Using Local Sensing and Minimal Communication. Special Issue on Advances in Multi-Robot Systems, T. Arai, E. Pagello, and L. E. Parker, editors, *IEEE Trans. on Robotics and Automation*, Vol. 18, No. 5, pp. 837-846.

[30] B. P. Gerkey and M. J. Mataric. Sold! Auction Methods for Multi-Robot Coordination. Special Issue on Advances in Multi-Robot Systems, T. Arai, E. Pagello, and L. E. Parker, editors, *IEEE Trans. on Robotics and Automation*, Vol. 18, No. 5, pp. 758-768.

[31] B. P. Gerkey, R. T. Vaughan, K. Sty, A. Howard, G. S. Sukhatme, and M. J. Matarić. Most Valuable Player: A Robot Device Server for Distributed Control. *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, Wailea, Hawaii, IROS-2001. Available at http://playerstage.sourceforge.net.

[32] B. P. Gerkey and M. J. Mataric. MURDOCH: Publish/Subscribe Task Allocation for Heterogeneous Agents. *Proc. of Autonomous Agents*, Barcelona, Spain, June 2000, pp. 203-204.

[33] D. Goldberg and M. J. Mataric. Coordinating Mobile Robot Group Behavior Using a Model of Interaction Dynamics. *Proc. of the third Int. Conf. on Autonomous Agents*, Seattle, Washington, pp. 100-107.

[34] D. M. Gordon. Ants at Work: How an Insect Society is Organized. The Free Press, 1999.

[35] S. Goss, J.-L. Deneubourg, J. M. Pasteels, and G. Josens. A Model of Non-Cooperative Foraging in Social Insects. *American Naturalist*, **134**, 1989, pp. 273-287.

[36] M. Grossglauser and D. Tse. Mobility Increases the Capacity of Ad-hoc Wireless Networks. *Proceedings INFOCOM 2001*, pp. 1360-1369.

[37] P. Gupta and P. R. Kumar. The Capacity of Wireless Networks. *IEEE Transactions on Information Theory*, 46(2), 2000, pp. 388-404.

[38] K. P. Hadeler and F. Rothe. Travelling Wave-Fronts in Nonlinear Diffusion Equations. *Journal of Math. Biol.*, **2**, 1975, pp. 251-263.

[39] M. Hashimoto, F. Oba, and T. Egushi. Dynamic Control Approach for Motion Coordination of Multiple Wheeled Mobile Robots Transporting a Single Object. *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robotic Systems*, 1993, pp. 1944-1951.

[40] 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*, October, 2000, Knoxville, TN, pp. 261-270.

[41] A. T. Hayes, A. Martinoli, and R. M. Goodman. Distributed Odor Source Localization. Special Issue on Artificial Olfaction, H. T. Nagle, Gardner, and K. Persaud, editors, *IEEE Sensors*, Vol. 2, No. 3, 2002, pp. 260-271.

[42] A. T. Hayes. Self-Organized Robotic System Design and Autonomous Odor Localization. Unpublished doctoral manuscript, Caltech, CA, June, 2002. See also http://www.coro.caltech.edu/People/athayes/athayes.html.

[43] A. T. Hayes, A. Martinoli, and R. M. Goodman. Swarm Robotic Odor Localization: Off-Line Optimization and Validation with Real Robots. Special issue on Bio-Inspired Robotics, McFarland D., editor, *Robotica*, 2003. To appear.

[44] O. E. Holland and C. Melhuish. Stigmergy, Self-Organization, and Sorting in Collective Robotics. *Artificial Life*, Vol. 5, 1999, pp. 173-202.

[45] B. Hölldobler and E. O Wilson. Journey of the Ants: A Story of Scientific Exploration. The Harvard University Press, Cambridge, Ma, 1994.

[46] S. Ichikawa and F. Hara. An Experimental Realization of Cooperative Behavior of Multi-Robot System. *Distributed Autonomous Robotic Systems*, 1994, pp. 224-234.

[47] S. Ichikawa and F. Hara. Experimental Characteristics of Multiple-Robots Behaviors in Communication Network Expansion and Object-Fetching. *Distributed Autonomous Systems*, **2**, eds. Springer-Verlag, Berlin, 1996, pp. 224-234.

[48] 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*, 2001, Vol. 11, No. 2, pp. 149-171.

[49] P. M. Kareiva and N. Shigesada. Analyzing Insect Movement as a Correlated Random Walk. *Oecologia*, (Berlin) 56, 1983, pp. 234-238.

[50] S. Kazadi, A. Abdul-Khaliq, and R. M. Goodman. On the Convergence of Puck Clustering Systems. *Robotics and Autonomous Systems*, Vol. 38, No. 2, pp. 93-117, 2002.

[51] I. D. Kelly, D. A. Keating. Flocking by the Fusion of Sonar and Active Infrared Sensors on Physical Autonomous Mobile Robots. *Proc. of the Third Conf. on Mechatronics and Machine Vision in Practice*, Gaimardes, Portugal, 1996, Vol. 1, pp. 1-4.

[52] D. G. Kendall. Mathematical Models of the Spread of Infection. *Mathematics and Computer Science in Biology and Medicine*, Medical Research Council, London, 1965, pp. 213-225.

[53] A. N. Kolmogoroff, I. G. Petrovsky, and N. S. Piscounoff. Etude de l'Equation de la Diffusion avec croissance de la Quantité de Matière et son Application à un Problème Biologique. *Bulletin de l'Université d'Etat à Moscou*, (série interne) A, **1**(6), 1937, pp.1-25.

[54] 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*, Vol. 30, No. 1-2, 2000, pp. 65-84.

[55] C. R. Kube and E. Bonabeau. Cooperative Transport by Ants and Robots. *Robotics and Autonomous Systems*, Vol. 30, No. 1-2, 2000, pp. 85-101.

[56] D. Kurabayashi, J. Ota, T. Arai, and E. Yoshida. Cooperative Sweeping by Multiple Mobile Robots. *Proceedings of IEEE International Conference on Robotics and Automation*, 1996, pp. 1744-1749.

[57] K. Lerman, A. Galstyan, A. Martinoli, and A. J. Ijspeert. A Macroscopic Analytical Model of Collaboration in Distributed Robotic Systems. *Artificial Life*, Vol. 7, No. 4, pp. 375-393, 2001.

[58] K. Lerman and A. Galstyan. Mathematical Model of Foraging in a Group of Robots: Effect of Interference.*Autonomous Robots*, Vol. 13, 2002, pp. 127-141.

[59] L. Li, A. Martinoli, and Y. Abu-Mostafa. Emergent Specialization in Swarm Systems. In H. Yin, N. Allinson, R. Freeman, J. Keane, and S. Hubbard, editors, *Proc. of the Third Int. Conf. on Intelligent Data Engineering and Automated Learning*, August 2002, Manchester, UK. Lecture Notes in Computer Science (2002), Vol. 2412, pp. 261-266.

[60] R.-L. Luther. Räumliche Fortpflanzung Chemischer Reaktionen. Z. für Elecktrochemie and angew. Physikalische Chemie 12(32), 1906, pp. 506-600. English translation: R. Arnold, K. Showalter, and J.J. Tyson: Propagation of Chemical Reactions in Space. *Journal of Chem. Educ.* 1988.

[61] M. Maris and R. de Boekhorst. Exploiting Physical Constraints: Heap Formation through Behavioral Error in a Group of Robots. *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, Osaka, Japan, vol. 3, 1996, pp. 1655-1660.

[62] A. Martinoli and F. Mondada. Collective and Cooperative Group Behaviours: Biologically Inspired Experiments in Robotics. In Khatib O. and Salisbury J. K., editors, *Proc. of the Fourth Int. Symp. on Experimental Robotics*, Stanford, June, 1995, Lecture Notes in Control and Information Sciences, pp. 3-10.

[63] A. Martinoli, A. J. Ijspeert, and F. Mondada. Understanding Collective Aggregation Mechanisms: From Probabilistic Modelling to Experiments with Real Robots. *Robotics and Autonomous Systems*, Vol. 29, pp. 51-63, 1999.

[64] A. Martinoli, A. J. Ijspeert, and L. M. Gambardella. A Probabilistic Model for Understanding and Comparing Collective Aggregation Mechanisms. In D. Floreano, F. Mondada, and J.-D Nicoud, editors, *Proc. of the Fifth European Conf. on Artificial Life*, Lausanne, Switzerland, Sept. 1999, Lectures Notes in Computer Science, pp. 575-584.

[65] A. Martinoli. Swarm Intelligence in Autonomous Collective Robotics: From Tools to the Analysis and Synthesis of Distributed Control Strategies. Unpublished doctoral manuscript, EPFL Ph.D. Thesis No. 2069, Lausanne, Switzerland, October 1999. Downloadable at: http://www.coro/caltech/edu/people/alcherio/am_pub.

[66] A. Martinoli and K. Easton. Modeling Swarm Robotic Systems. In Siciliano B. and Dario P., editors, *Proc. of Eighth Int. Symp. on Experimental Robotics*, Sant'Angelo d'Ischia, Italy, July 2002. Springer Tracts in Advanced Robotics (2003), pp. 285-294.

[67] A. Martinoli and K. Easton. Optimization of Swarm Robotic Systems via Macroscopic Models. In A. Schultz, L. Parker, and F. Schneider, editors, *Proc. of the Int. Workshop on Multi-Robot Systems*, Naval Research Laboratory, Washington, DC, March 2003, pp. 181-192.

[68] A. Martinoli, K. Easton, and W. Agassounon. Modeling Swarm Robotic Systems: A Case Study in Collaborative Distributed Manipulation. International Journal of Robotic Research, 2003. To appear.

[69] M. J. Matarić. Minimizing Complexity in Controlling a Mobile Robot Population. *Proc. of the IEEE Int. Conf. on Robotics and Automation*, 1992, pp. 830-835.

[70] M. J. Matarić. Interaction and Intelligent Behavior. Unpublished doctoral manuscript, Dept. of Electrical Engineering and Computer Science, MIT, Cambridge, MA, 1994.

[71] M. J. Matarić. Using Communication to Reduce Locality in Distributed Multi-Agent Learning. *Journal of Experimental and Theoretical Artificial Intelligence*, **10**, No. 3, 1998, pp. 357-369.

[72] C. Melhuish, O. Holland, and S. Hoddell. Collective Sorting and Segregation in Robots with Minimal Sensing, *Proc. of the Fifth Conf. on Simulation of Adaptive Behavior*, Zurich, Germany, 1998.

[73] R. M. Metcalfe and D. R. Boggs. Ethernet: Distributed Packet Switching for Local Computer Networks. *Communication of the ACM*, **19**, No. 5, July 1976, pp 395-404.

[74] O. Michel. Webots: Symbiosis Between Virtual and Real Mobile Robots. In J.-C Heuding, editor, *Proc. of the First Int. Conf. on Virtual Worlds*, Paris, France, July 1998, Springer Verlag, pp. 254-263. See also http://www.cyberbotics.com/webots/.

[75] N. Miyata, J. Ota, T. Arai, E Yoshida, D. Kurabayashi, J. Sasaki, and Y. Aiyama. Cooperative Transport with Regrasping of Torque-Limited Mobile Robots. *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robotic Systems*, 1996, pp. 304-309.

[76] D. Mollison. Possible Velocities for a Simple Epidemic. *Adv. Appl. Prob.*, **4**, 1972, pp.233-258.

[77] D. Mollison, Spatial Contact Models for Ecological and Epidemic Spread. *J. R. Statist. Soc.*, B 39, No. 3, 1977, pp. 283-326.

[78] F. Mondada, E. Franzi, and P. Ienne. Mobile Robot Miniaturization: A Tool for Investigation in Control Algorithms. In T. Yoshikawa and F. Miyazaki, editors, *Proc. of the Third Int. Symp. on Experimental Robotics*, Kyoto, Japan, October 1993, Lecture Notes in Control and Information Sciences, Springer Verlag, pp. 501-513.

[79] J. D. Murray. Mathematical Biology. Springer-Verlag, Heidelberg, 1989.

[80] F. Noreils, An Architecture for Cooperative and Autonomous Mobile Robots, *Proc. of the IEEE Int. Conf. on Robotics and Automation*, 1992, pp. 2703-2709.

[81] G. F. Oster and E. O. Wilson. Caste and Ecology in the Social Insects. *Monographs in Population Biology*, No. 12, Princeton University Press, Princeton, NJ.

[82] S. W. Pacala, D. M. Gordon, and H. C. J. Godfray. Effects of Social Group Size on Information Transfer and Task Allocation, *Evolutionary Ecology*, vol. 10, 1996, pp. 127-165.

[83] L. E. Parker. ALLIANCE: An Architecture for Fault Tolerant Multi-Robot Cooperation, *IEEE Transaction on Robotics and Automation*, 1998, vol. 14, No. 2., pp. 220-240.

[84] L. E. Parker. Lifelong Adaptation in Heterogeneous Multi-Robot Teams: Response to Continual Variation in Individual Robot Performance. Autonomous Robots, Vol. 8, pp. 239-267, 2000.

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

[86] J. M. Pasteels, J.-L. Deneubourg, and S. Goss. Self-Organizing Mechanisms in Ant Societies. I. Trail Recruitment to Newly Discovered Food Sources. In J. M. Pasteels and J.-L. Deneubourg, eds. From individual to collective behavior in social insects. *Experientia Supplementum.* **54**. Birkhauser, Basel, 1987, pp. 155-175.

[87] G. H. Pyke. Are Animals Efficient Harvesters? *Animal Behaviour*, **26**, 1978, pp. 247-250.

[88] D. J. Stilwell and J. S. Bay. Optimal Control for Cooperative Mobile Robots Bearing a Common Load. *Proc. of the IEEE Int. Conf. on Robotics and Automation*, 1994, pp. 58-63.

[89] K. Singh and K. Fujimura. Map Making by Cooperative Mobile Robots. *Proc. of the IEEE Int. Conf. on Robotics and Automation*, 1993, pp. 254-259.

[90] J. G. Skellam. Random Dispersion in Theoretical Populations. *Biometrica*, **38**, 1951, pp. 196-218.

[91] L. Steels. Cooperation between Distributed Agents through Self-Organization. *Decentralized Artificial Intelligence*, North-Holland, Amsterdam, 1990, pp. 175-196.

[92] D. J. Stilwell and J. S. Bay. Optimal Control for Cooperative Mobile Robots Bearing a Common Load. *Proc. of the IEEE Int. Conf. on Robotics and Automation*, 1994, pp. 58-63.

[93] K. Sugawara and M. Sano. Cooperative Acceleration of Task Performance: Foraging Behavior of Interacting Multi-Robot Systems. *Physica* D, Vol. 100, 1997, pp. 343-354.

[94] K. Sugawara, M. Sano, I. Yoshihara, and K. Abe. Cooperative Behavior of Interacting Robots. *Artificial Life and Robotics*, No. 2, 1998, pp. 62-67.

[95] H. Takagi and L. Kleinrock. Optimal Transmission Ranges for Randomly Distributed Packet Radio Terminals. *IEEE Transaction on Communication*, Vol. COM-32, No. 3, 1984, pp. 246-257.

[96] G. Theraulaz, E. Bonabeau, and J.-L Deneubourg. Response Threshold Reinforcement and Division of Labour in Insect Societies. *Proc. of Royal Society of London*, Series B, Vol. 265, 1998, pp. 327-332.

[97] G. Theraulaz, E. Bonabeau, and J.-L Deneubourg. Les Insectes Architectes ont-ils leur Nid dans la Tête?. *La Recherche*, No.313, 1998, pp. 84-90.

[98] J. Wang. On Sign-Board based Inter-Robot Communication in Distributed Robotic Systems. *Proc. of IEEE Int. Conf. on Robotics and Automation*, 1994, pp. 1045-1050.

[99] J. Wang and S. Premvuti. Resource Sharing in Distributed Robotic Systems based on a Wireless Medium Access Protocol (CSMA/CD-W). *Robotics and Autonomous Systems*, vol. 19, No. 1, 1999, pp. 33-56.

[100] T. Weigel, J.-S Gutmann, M. Dietl, A. Kleiner, and B. Nebel. CS Freiburg: Coordinating Robots for Successful Soccer Playing. Special Issue on Advances in Multi-Robot Systems, T. Arai, E. Pagello, and L. E. Parker, editors, *IEEE Trans. on Robotics and Automation*, Vol. 18, No. 5, pp. 685-699.

[101] E. O. Wilson and B. Hölldobler. Dense Heterarchies and Mass Communication as the Basis of Organization in ant colonies. *Trends in Ecology and Evolution*, **3**, 1978, pp. 65-68.

[102] E. O. Wilson. The Relation between Caste Ration and Division of Labour in Ant Genus Pheidole (*Hymnoptera Formidacae*). *Behav. Ecol. Sociobiol.* **16**, 1984, pp.89-98.

[103] B. Yamauchi. Decentralized Coordination for Multi-robot Exploration. Robotics and Autonomous Systems, Vol. 29, No. 1, 1999, pp. 111-118.

[104] E. Yoshida and T Arai. Performance Analysis of Local Communication by Cooperating Mobile Robots. *IEICE Transaction on Communication*, Vol. E83-B, No. 5, 2000, pp. 1048-1059.

[105] S. Yuta and J. Iijima. State Information Panel for Inter-Processor Communication in an Autonomous Mobile Robot-Controller. *Proc. of the IEEE Workshop on Intelligent Robot Systems*, 1990.

[106] Y. Zhang, A. Martinoli, and E. K. Antonsson. Evolutionary Design of a Collective Sensory System, *Proc. of the AAAI Spring Symp. Series on Computational Synthesis*, Stanford, CA, March 2003. pp. 283-290.