



Université Libre de Bruxelles  
Faculté des Sciences Appliquées  
CODE - Computers and Decision Engineering  
IRIDIA - Institut de Recherches Interdisciplinaires  
et de Développements en Intelligence Artificielle

---

# Towards Cooperation in a Heterogeneous Robot Swarm through Spatially Targeted Communication

---

**Nithin MATHEWS**

Superviseur:

**Prof. Marco DORIGO**

Rapport d'avancement de recherche  
Année Académique 2009/2010





# Abstract

In this work, we study cooperation in a heterogeneous swarm robotic system composed of wheeled and aerial robots. The wheeled robots are able to physically connect to one another autonomously and thus form collective robotic entities. The aerial robots have a privileged overview of the environment since they can fly and attach to metal ceilings. We present a system that enables *spatially targeted communication*. Our system enables aerial robots to establish dedicated communication links with individual wheeled robots or with selected groups of wheeled robots based on their position in the environment. The heterogeneous swarm robotic hardware is currently under development. We therefore demonstrate the proposed approach on an existing multirobot system consisting of only wheeled robots by letting one of the wheeled robots assume the role of an aerial robot. We then go on to show how such a spatially targeted communication link can be used to enable cooperation in the heterogeneous swarm robotic system considered in this work. That is, we show how the aerial robots can communicate with selected groups of wheeled robots and instruct them on how to overcome obstacles in their path by forming morphologies appropriate to the obstacle encountered. We conduct experiments in simulation to quantify separately the benefits of cooperation and of spatially targeted communication. Our approach does not require any form of global information and relies solely on a situated communication modality based on LEDs and camera.



# Acknowledgements

I thank Prof. Marco Dorigo for giving me the chance to work at IRIDIA. I also thank Dr. Anders Lyhne Christensen<sup>1</sup> and Rehan O'Grady for their guidance and support they have given me to develop the ideas and to conduct the research presented in this work. I am also grateful to Eliseo Ferrante for his immense contributions. This work could not have been conducted without the state-of-the-art simulator developed by Carlo Pinciroli. A special thanks to you, Carlo. I would also like to express my gratitude to all my colleagues at IRIDIA.

This work was carried out in the framework of SWARMANOID, a project funded by the Future and Emerging Technologies programme (IST-FET) of the European Commission, under grant IST-022888.

---

<sup>1</sup>I honestly think that it is you, who is the *champ* !



# Statement

This work describes an original research carried out by the author and was only supported by the declared resources or auxiliaries. It has not been submitted either in the same or in different form to this or any other university for the award of any degree. However, the content of this work is partially based on articles that the author, together with other co-authors, is about to make public in the scientific literature. Therefore, this work contains literal or analogous rewriting from the following articles accepted for publication:

N. Mathews, A. L. Christensen, E. Ferrante, R. O'Grady, and M. Dorigo. Establishing Spatially Targeted Communication in a Heterogeneous Robot Swarm. *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2010)*.

N. Mathews, A. L. Christensen, R. O'Grady, and M. Dorigo. Cooperation in a Heterogeneous Robot Swarm through Spatially Targeted Communication. *Proceedings of the 7th International Conference on Ant Colony Optimization and Swarm Intelligence (ANTS 2010)*. Springer Verlag, Berlin, Germany, 2010.



# Contents

<b>Abstract</b>	<b>iii</b>
<b>Acknowledgements</b>	<b>v</b>
<b>Statement</b>	<b>vii</b>
<b>1 Introduction</b>	<b>1</b>
<b>2 Hardware Platforms</b>	<b>5</b>
2.1 Situated Communication Using LEDs and Camera . . . . .	6
<b>3 Related Work</b>	<b>9</b>
3.1 Communication in Multirobot Systems . . . . .	9
3.2 Self-Assembling Robots . . . . .	10
<b>4 Establishing Spatially Targeted Communication</b>	<b>11</b>
4.1 One-To-One Communication . . . . .	11
4.1.1 Probabilistic Model . . . . .	12
4.1.2 One-to-One Controller . . . . .	14
4.1.3 Simulation-based Experiments . . . . .	15
4.1.4 Real Robot Experiments . . . . .	17
4.2 One-To-Many Communication . . . . .	18
4.2.1 One-to-Many Controller . . . . .	20
4.2.2 Simulation-based Experiments . . . . .	21
4.2.3 Real Robot Experiments . . . . .	22
<b>5 Cooperation through Spatially Targeted Communication</b>	<b>25</b>
5.1 Communication Between an Eye-bot and the Foot-bots . . .	25
5.2 Morphology Generation Using Directional Self-Assembly . . .	26
5.3 Example SWARMORPH-Script . . . . .	26
5.4 Simulated Cooperation in the Heterogeneous Swarm . . . . .	28

<b>6 Experiments and Results</b>	<b>31</b>
6.1 Experimental Setup . . . . .	31
6.2 Three Control Strategies . . . . .	32
6.3 Performance Benefits of Cooperation with Spatially Targeted Communication . . . . .	33
6.4 Performance Benefits of Spatially Targeted Communication .	35
<b>7 Conclusions and Future Work</b>	<b>37</b>
<b>Bibliography</b>	<b>39</b>

# Chapter 1

## Introduction

Heterogeneous multirobot systems have been the topic of several studies, see for instance [34, 36]. One of the main motivations for these systems is that the capabilities required for some tasks are difficult to be satisfied by a single type of robot only. Furthermore, while each type of robot can be kept relatively simple, the multirobot system as a whole can have a broad range of capabilities. In this work, we consider a heterogeneous *swarm* system in which aerial robots (hereafter referred to as eye-bots) supervise the activities of ground based wheeled robots (hereafter referred to as foot-bots). The foot-bots are capable of autonomous self-assembly which means that they can make physical connections with one another and form collective robotic entities. By leveraging the collective strength and reach of a self-assembled entity, the foot-bots can solve tasks that individual robots cannot solve alone. In such a system, it can be crucial for the eye-bots to be able to communicate (i.e., send tailored instructions) with particular foot-bots or groups of foot-bots based on their location in the environment. We refer to this type of communication as *spatially targeted communication*. A typical message that an eye-bot might need to send using spatially targeted communication is “connect to the robot on your left”.

Some researchers have used *situated communication modalities* to implement spatially targeted communication. In situated communication modalities, localization information about the sender is implicit in the message delivery mechanism [35]. As an example of situated communication, consider the case in which a robot receives the message “stay away, I am near danger”. This message is only meaningful if the communication is situated, that is, if the receiving robot can estimate the location of the sender. Spatially targeted communication has been achieved using protocols built on top of situated communication modalities. However, existing implementations of spatially targeted communication are either unsuited to larger swarms due to the characteristics of the communication hardware used [31] or rely on some form of global knowledge [34, 36] that is not available on most

swarm robotics systems.

In this work, we first show how eye-bots can establish a spatially targeted communication channel to one or more co-located foot-bots without relying on any form of global knowledge. Instead, our approach relies on a situated communication based on on-board LEDs and camera. We use a binary selection process, whereby the eye-bot initially communicates with all foot-bots within visual range and iteratively eliminates robots, until only the selected robot is left. We provide a probabilistic model to derive an upper bound of the average convergence time. We show how an established one-to-one communication link can be expanded to a one-to-many communication link with a group of co-located wheeled robots. Our approach is applicable to any communication modality that is (i) situated and (ii) supports at least three distinct signals.

Second, we show how an existing spatially targeted communication link can be used to achieve cooperation in the heterogeneous swarm. That is, we use a technique named SWARMORPH-script [6] developed in previous research. This technique allows the eye-bot to send morphology growth instructions to the foot-bots to which a spatially targeted communication has already been established. SWARMORPH-script has been successfully tested on real robotic hardware, does not require any form of global information and relies solely on a situated communication modality based on LEDs and camera.

The task considered in this work is a gap crossing task (see Fig. 1.1) similar to the one considered in [28]. In [28], a team of wheeled robots autonomously self-assembled into different morphologies to overcome different types of obstacles, one of which was a gap obstacle. In that study, however, the response to the presence of each obstacle type was preprogrammed. The wheeled robots did not have the sensory capabilities to estimate the width of a gap. Therefore, on encountering a gap, they would always self-assemble into a four robot line morphology irrespective of the width of the gap. In this work, we propose an approach that enables the foot-bots to adaptively generate the appropriate morphology for the current gap size by cooperating with eye-bots through spatially targeted communication. Depending on the width of the gap, the foot-bots may need to self-assemble into a collective robotic entity to successfully overcome the gap obstacle (if a robotic entity with an inappropriate size or shape attempts to cross an overly wide gap, it will fall in and will not succeed in solving the task). The task requires one or more foot-bots to cross a gap of a priori unknown width. We use four possible gap widths. The foot-bots do not have the sensory apparatus to perceive the width of the gap. When operating alone, that is, without cooperating with the eye-bot, the foot-bots must assume the worst case scenario and self-assemble into a four foot-bot line morphology. When cooperating with an eye-bot attached to the ceiling, however, the eye-bot can estimate the gap width and communicate instructions to the foot-bots using

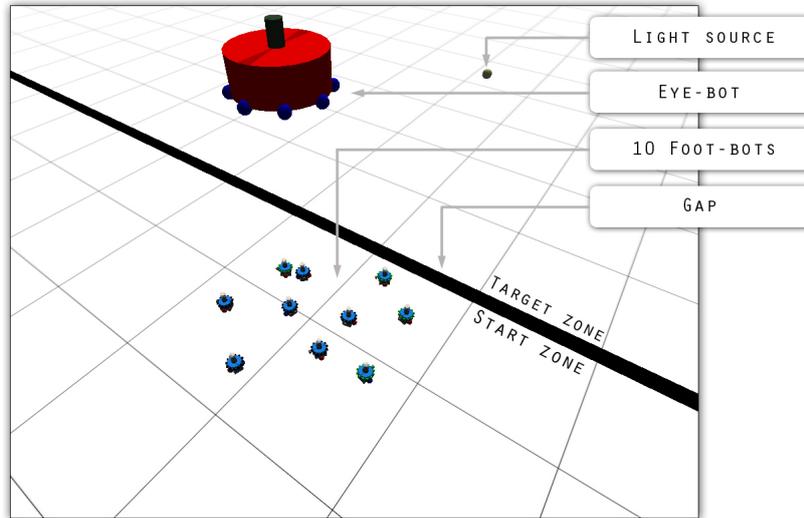


Figure 1.1: A depiction of the arena in which the heterogeneous swarm operates. The dark strip represents the gap which separates the arena into a start zone and a target zone. The circular object shown in the target zone is the light source. An eye-bot and 10 foot-bots are visible in the start zone.

a SWARMORPH-script that will permit the foot-bots to form the simplest possible morphology that will get them over the gap they have encountered.

We present the results of the simulation-based experiments in which we compare the task completion time of the swarm using different control strategies. The simplest strategy is not to cooperate with the eye-bot at all. As in [28], the foot-bots solve the task alone by assuming the worst case scenario and self-assembling into a four foot-bot line morphology as soon as they encounter a gap. We compare this strategy with the cooperative strategy that uses spatially targeted communication and SWARMORPH-script instruction transmission. Finally, we isolate the benefits of spatially targeted communication using a cooperative strategy in which the eye-bots still communicate instructions, but to a randomly selected group of foot-bots rather than to a group of foot-bots selected on the basis of their optimal location.



## Chapter 2

# Hardware Platforms

We consider the heterogeneous swarm consisting of the foot-bots and the eye-bots shown in Fig. 2.1a and Fig. 2.1b. At the time of writing, this heterogeneous swarm robotic platform is still under development. Therefore, to evaluate our approach, we use a custom physics based simulator named AR-GoS [30]. We also perform proof-of-concept experiments on a homogeneous robotic platform (see Fig. 2.1c) consisting of wheeled robots only, in which a predesignated wheeled robot assumes the role of the aerial robot. Below, we highlight those features of the heterogeneous and the homogeneous robotic platform, that are essential to this work. Note that these features represent only a subset of the features actually present on the robots. In Sect. 2.1, we also describe how both platforms are related to each other.

The foot-bots (Fig. 2.1a) are capable of moving and connecting to each other using a docking mechanism. They are also equipped with infrared proximity sensors used for obstacle avoidance, colored LEDs to display internal states to the eye-bots and to neighboring foot-bots, an omnidirectional camera to perceive other foot-bots, and an upward pointing camera to perceive eye-bots. We assume that both eye-bots and the foot-bots are able to display at least three distinct colors using their LEDs. Each foot-bot also has a system composed of infrared transceivers and radio [33]. This system provides local communication capabilities between foot-bots.

Fig. 2.1b shows an early prototype of the eye-bot. Once completed development, the eye-bots will be able to fly and will also be equipped with a system of magnets that allows them to attach to a metal ceiling or to metal bars. In this study, the eye-bots are assumed to be attached to the ceiling (thus stationary at an elevated position). They will also be equipped with a pan-and-tilt camera pointing downward. This camera allows them to survey the ground and to detect the foot-bots. Downward facing LEDs are used to communicate internal state information to the foot-bots.

We used a number of autonomous wheeled robots called *s-bots* [24] (see Fig. 2.1c) to conduct our real robot experiments. Each s-bot is equipped

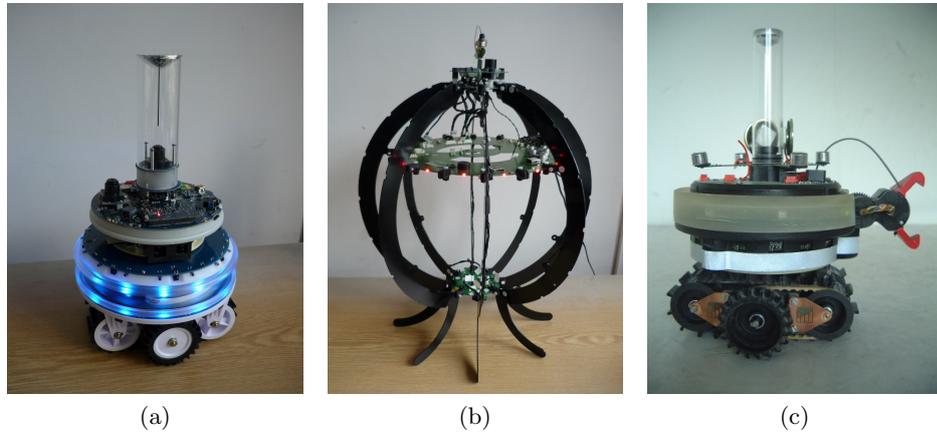


Figure 2.1: The robots considered in this work. (a) An almost final version of the foot-bot. (b) A prototype of the eye-bot. (c) An s-bot. At the time of writing, the robots shown in (a) and (b) are being developed at EPFL within the framework of the Swarmanoid project. More information about the project and the robots are available at <http://www.swarmanoid.org>.

with an XScale CPU running at 400 MHz, a set of actuators including a transparent ring around its chassis containing 8 RGB colored LEDs and a number of sensors including an omnidirectional camera (see Fig. 2.1c). The camera is mounted on the s-bot and points upward at a hemi-spherical mirror mounted at the top end of a transparent tube. The hemi-spherical mirror reflects panoramic images of the s-bot’s vicinity up to a distance of 70 cm, depending on light conditions. An s-bot communicates its internal state to nearby robots using red, green and blue LEDs. The s-bots have been used in several other studies including self-assembly [14], group transport [26], morphology control [7], and negotiation of goal direction [5].

## 2.1 Situated Communication Using LEDs and Camera

According to Kasper Støy [35], communication can be of two different types: situated communication and abstract communication. While in abstract communication (for instance when using wireless Ethernet), the physical signal that transports a message is considered not to have any meaning, whereas situated communication modalities associate a meaning to both the physical properties of the signal that transfers the message and to the content of the message. Consider a human-to-human conversation as an example: a message such as “stay away from me” is instantly meaningful as the location of the speaker together with the content of the message tells the

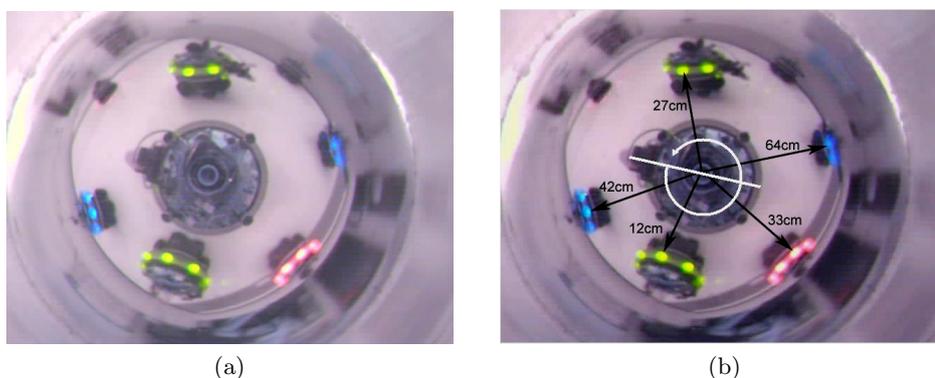


Figure 2.2: Situated communication using LEDs and camera. (a) An image returned by the omnidirectional camera of an s-bot. (b) The post-processed image (i.e., after applying blob detection) showing the estimated relative distance and bearing information of neighboring s-bots.

listener what to do. In this case, neither the location of the speaker nor the content of the message alone would have given any meaning to the listener. Moreover, it is the combination of the both that gives this conversation a meaning.

In this work, we use the onboard LEDs and the camera of the robots in the heterogeneous platform to create a situated communication modality. We use the LED ring of the robots to let them *speak to* each other. In other words, the LED ring is used to send messages to each other. Each color displayed on the LED ring represents a different message. On the other hand, we use the onboard cameras of the robots to detect the colors displayed by neighboring robots. In other words, the robots use the camera to *listen to* each other. The image returned by a camera are post-processed (i.e., by using blob detection or if possible a circle detection techniques to detect the LED rings of neighboring robots) to retrieve the relative localization information (i.e., the distance and the bearing) about all neighboring message sending robots in the vision range. The following instances of communication in the heterogeneous platform are considered in this work:

- An eye-bot perceiving one or multiple foot-bots on the ground using its pan-and-tilt camera
- A foot-bot using its camera pointed to the ceiling to perceive an eye-bot
- A foot-bot using its omnidirectional camera to perceive its neighboring foot-bots

The s-bots platform is very similar to the heterogeneous robotic platform considered in this work as the s-bots possess an identical LEDs and camera-based situated communication modality as the robots in the heterogeneous robotic platform. This communication modality enables the s-bots to retrieve localization information about each message sending neighboring s-bot by using a blob detection technique (see Fig. 2.2). In the LEDs and camera-based situated communication modality, the perception of a message listening robot is reduced to the two-dimensional image returned by the according camera. On such an image, image post-processing techniques can be applied as shown on the s-bots platform to retrieve localization information about message sending robots in the vision range. Given the fact that both the heterogeneous and the s-bots platform possess the identical situated communication modality, the s-bots represent a suitable platform to prove our concept on real robotic hardware. In order to emulate the heterogeneous swarm using the s-bots, however, we let a predesignated s-bot assume the role of the eye-bot.

## Chapter 3

# Related Work

In this chapter, we present a structured review of some existing related works. In Sect. 3.1, we first present the state-of-the-art of communication modalities used in multirobot systems in general by giving preference to i) heterogeneous systems which include aerial robots and wheeled robots and ii) to systems which enable situated communication. In Sect. 3.2, we present existing implementations of systems including modular robots that are able to dynamically self-assemble and therefore reconfigure their morphologies.

### 3.1 Communication in Multirobot Systems

In heterogeneous systems consisting of aerial robots and wheeled robots, wireless Ethernet has previously been used for communication [34, 36]. These studies compensate for the absence of inherent localization information in the wireless Ethernet medium by using global maps in conjunction with additional hardware such as GPS receivers. However, GPS is not available in indoor environments for which our heterogeneous robotic platform is designed.

Pugh and Martinoli [31] were among the first to report on a situated communication modality based on infrared transceivers. In this study, the messages exchanged between robots in the same geometric plane were used by the receiving robot to calculate the relative distance and bearing of the sending robot. However, this technology is not suited to concurrent communication for groups of more than 10 robots.

In [12], an ultrasonic localization system was described in which a team of robots was able to measure the range between each robot pair. However, the approach was subject to severe accuracy problems and did not include any inter-robot communication mechanism. In another study [32], accurate positioning was achieved using time-of-flight evaluation of ultrasonic pulses and a radio frequency communication link. The system was only tested with four robots and it remains unclear how echo effects would affect the

performance if the number of robots in the system is increased.

Some multirobot systems have exploited short-range communication radio technologies [2, 23]. However, these technologies are based on individual robots establishing serial communication links with each other and thus only allow for simultaneous communication between pairs of robots.

### 3.2 Self-Assembling Robots

In the past 50 years, many researchers have designed and studied modular systems whose components—ranging from passive mechanical parts to mobile robots—can reconfigure or self-assemble into physically connected structures [13]. Fukuda *et al.*'s CEBOT system [9, 19] is one of the first implementations of a reconfigurable modular system. The mobile architecture consists of heterogeneous modules with different functions, e.g. to rotate, move, and bend. Various prototypes of the CEBOT system comprising different shapes and connection mechanisms have been studied.

Hirose *et al.* [17] investigated a modular robot and described potential benefits of such systems in the context of autonomous all-terrain locomotion. Yim *et al.* [37] predicted that such systems would be particularly suited to applications in which versatility is critical. “Typically, these are situations in which some information about the environment is not known a priori. Thus, a system cannot be designed specifically for a task, since the task that is needed is not known”.

For an overview of the field of modular self-reconfigurable robotic systems, see [39]. The components of such self-reconfigurable systems can autonomously reorganize into different configurations. Several different hardware architectures (lattice, chain/tree, mobile) and many different implementations and control mechanisms have been proposed [4, 21, 25, 38]. However, in the majority of current implementations, the components are either manually pre-assembled or rely on their environment (be it natural or manmade) to provide the energy required for independent movement. Once assembled, most existing systems are furthermore incapable of autonomously assimilating additional modules.

Self-propelled self-assembling robotic systems, in contrast, are made up of independent autonomous mobile components that are capable of forming physical connections with each other without external direction. Such self-assembling systems are potentially more flexible than pre-connected self-reconfigurable systems. Several architectures have been proposed, which have been implemented with varying degrees of success [3, 8, 10, 14, 17]. However, none of the existing systems display any meaningful control over the morphology of the connected entity formed through the self-assembly process.

## Chapter 4

# Establishing Spatially Targeted Communication

In this chapter, we explain our two-step approach to establish spatially targeted communication. In a first step, we let the eye-bot narrow down the number of potential message receiving foot-bots to a single *seed* foot-bot (i.e., creates a one-to-one communication link). In a second step, we let the eye-bot expand the existing one-to-one communication link with the seed foot-bot to include co-located foot-bots (i.e., to create a one-to-many communication link). In the following sections, we describe both steps of our approach in detail.

### 4.1 One-To-One Communication

We first explain the approach we use to establish a one-to-one communication link between robots. We present a probabilistic model of the approach and use it to derive an upper bound on completion time. We go on to describe the controllers of the simulated and real robots. We then describe the experimental setup used in simulation to study the following: 1) the impact the number of distinctive signals<sup>1</sup> available to the system has on the completion time, and 2) the scalability of our approach. After presenting the results of these simulation-based studies, we prove our concept on the real robots.

Given a set  $C := \{c_1, \dots, c_s\}$  of distinctive signals available to both robot types, where  $s \geq 3$ , a spatially targeted communication can be established by the eye-bot with a particular foot-bot by the means of an iterative selection process. Note that the subset  $C_s := \{c_2, \dots, c_s\}$  of the available distinctive signals is exclusive to the iterative selection process. In what follows, we describe our approach under the assumption that  $C := \{red, blue, green\}$

---

<sup>1</sup>In our system, each different LED color is considered a distinctive signal.

and  $C_s := \{blue, green\}$ .

We assume that the eye-bot has already selected a particular foot-bot with which it wishes to communicate. The eye-bot first attracts the attention of all foot-bots in visual range by signaling  $c_1 = red$ , the SOS signal. All foot-bots able to perceive the SOS signal register to the iterative selection process by replying with  $c_2 = blue$ . The eye-bot responds to this initial registration with a matching handshake using  $c_2 = blue$ . After this handshake, the iterative selection process starts. At each iteration, every foot-bot that is still part of the selection process randomly chooses and illuminates a color from the set  $C_s$ . At each iteration, the eye-bot illuminates its LEDs to match the color chosen by the selected foot-bot with which it wishes to communicate. At the end of every iteration, only those robots whose color match that of the eye-bot remain part of the selection process. The foot-bots which are not part of the selection process do not illuminate any color. This iterative selection process continues until the selected foot-bot is the only illuminated robot. In this case, the eye-bot indicates the termination of the selection process to the foot-bot by repeating  $c_1$  again. The remaining foot-bot acknowledges this by matching the eye-bot's color. The eye-bot and the remaining foot-bot have now established a spatially targeted communication link.

#### 4.1.1 Probabilistic Model

In this section, we introduce a model that formally describes the selection process. Our aim is to provide a model to determine a theoretical upper bound on the average time it takes for the selection process to complete. The model is empirically validated using the data gained from simulation-based experiments in Sect. 4.1.3.

We are interested in a model for the random variable  $T_n$  which is described as the *number of iterations to the end of the selection process* where  $n$  is the number of robots which will be discarded in the selection process. Our second objective is then to find the asymptotic behavior of the expectation  $E[T_n]$  as  $n \rightarrow \infty$  and bounds on its value  $E[T_n] \leq b$ .

Consider the two sets  $R_1 := \{r_1, \dots, r_n\}$  and  $R_2 := \{r_{sel}\}$ : the first set consists of the robots which will be discarded in the selection process, whereas the second set consists of the robot which will eventually be selected. Let  $p_s = \frac{1}{|C_s|}$  be the probability of one robot selecting a particular signal amongst the  $|C_s|$  available signals. If  $n = 1$ , the selection process is reduced to a sequence of Bernoulli trials with parameters  $p_0$  (the probability to leave the process) and  $p_1$  (the probability of staying in the process). The event  $r_1$  emitting a specific signal from  $C_s$  is independent from the event  $r_{sel}$  emitting a specific signal from  $C_s$ . By the product rule, the probability of both  $r_1$  and  $r_{sel}$  selecting a given signal is  $p_s^2$ . Then  $p_1$  is the probability of the two robots selecting an equal color from  $C_s$ :  $p_1 = |C_s| \cdot p_s^2 = |C_s| \cdot \frac{1}{|C_s|^2} = p_s$ .

Obviously,  $p_0 = 1 - p_s$ . In the simple setting  $n = 1$ ,  $T_1$  is a random variable with the geometric probability distribution:

$$P(T_1 = k) = (1 - p_0)^{k-1} \cdot p_0,$$

with mean  $E[T_1] = \frac{1}{p_0}$  and variance  $Var[T_1] = \frac{1-p_0}{p_0^2}$ . However, when  $n > 1$ , the analytical derivation of  $T_n$  and of its moments is a non-trivial task.

To further proceed towards our objectives, we apply the theory of branching processes [20]. A branching process, also called (in its discrete-time version) the Galton-Watson process, is a widely used model to study reproduction and population growth. The process traditionally starts with only one individual (the ancestor) at time or generation  $k = 0$ . At generation 1, the ancestor dies and spawns a number of individuals  $Y$  according to the probability distribution  $P(Y = h) = p_h$ , where  $Y$  takes values in  $0, 1, 2, \dots$  with probability  $p_0, p_1, p_2, \dots$ . The process then goes on: at generation  $k$  there will be  $Z_k$  individuals, which were spawned at generation  $k - 1$  by  $Z_{k-1}$  individuals with the same probability distribution  $p_h$ .

Our selection process can be modeled as a Galton-Watson process that starts with  $n$  individuals instead of 1 and where  $Y$  has probability distribution  $p_h$  defined as:

$$p_h = \begin{cases} 1 - p_s & \text{if } h = 0 \\ p_s & \text{if } h = 1 \\ 0 & \text{if } h > 1, \end{cases}$$

i.e., each individual can only have 1 offspring (itself) with probability  $p_s$  (it matches the color hence it *survives*) or 0 offspring with probability  $1 - p_s$  (it does not match the color hence it *dies*).

In a branching process, the probability of ultimate extinction, i.e.  $P(Z_k = 0)$  for some  $k$  is often considered in studies. If  $P(Z_k = 0) = 1$ , it means that the population will eventually (i.e., for some  $k$ ) become extinct. In our case, it represents the probability of extinction (i.e., termination) of the selection process. Hence, we require this probability to be 1 for the algorithm to be applicable in our case. Fortunately, this is proved to be always true in our case [1]. In [1], a branching process is shown to lead to extinction if  $m = E[Y] < 1$ , where  $m$  is the average number of offspring each individual spawns. In our case,  $m = 0 \cdot p_0 + 1 \cdot p_1 = p_1 = p_s < 1$ . Hence extinction (i.e., eventual termination of the selection process) is guaranteed.

We now return to our original question: given that the process terminates, how long does it take to do so? This question is equivalent to asking what is the probability distribution of the so called *time to extinction*  $T_n$ , i.e.,  $P(T_n = k) = p_k$ . Unfortunately, it is not trivial to derive the closed form of the density function  $p_k$  [16]. However, some general properties can be derived for its mean,  $E[T_n]$ . In particular, in [15, 18, 29] it is shown that,

under some non-strict conditions and for  $n > 3$ , the following two properties hold:

$$E[T_n] \sim \frac{\ln n}{|\ln m|}, n \rightarrow \infty \quad (4.1)$$

$$E[T_n] \leq \frac{\ln n}{|\ln m|} + \frac{2 - m}{1 - m} \quad (4.2)$$

In Sect. 4.1.3, we compare the upper bound on the mean predicted by Eq. 4.2 with the values obtained from simulation-based experiments.

### 4.1.2 One-to-One Controller

We developed two controllers: one for the eye-bots and one for the foot-bots. The controllers are completely distributed and homogeneous, i.e., all foot-bots execute the same controller. Both controllers are behavior-based and are represented as finite state machines (FSMs) in Fig. 4.1. In what follows, we explain both controllers assuming that both aerial and foot-bots use three colors to communicate: red, blue and green.

Fig. 4.1a shows the FSM implemented on the aerial robots consisting of the three states **SOS** (request connection), **SP** (selection process) and **CE** (communication established). The states **SOS** and **CE** are associated with the same predefined color red, whereas the state **SP** is provided with the two remaining colors blue and green for the selection process. Once an eye-bot has determined that it needs to communicate with a particular foot-bot, it enters the **SOS** state. The transition from **SOS** to **SP** is triggered when the selected foot-bot acknowledges the **SOS**. While in state **SP**, the eye-bot keeps matching the color displayed by the foot-bot. If the selected foot-bot is the only robot displaying any color, the eye-bot changes its state to **CE** to confirm the establishment of the communication.

Fig. 4.1b shows the FSM implemented on the foot-bots; it consists of the three states **ACK** (acknowledge), **SP** (selection process), **CE** (communication established) and an end state which causes the foot-bot to terminate the behavior. The state **ACK** is associated with the predefined color blue and the state **CE** is associated with red. The state **SP** is given two colors, namely blue and green. The foot-bot enters the state **ACK** as soon as an **SOS** color is perceived on the aerial robot. In case the **ACK** color is matched by the eye-bot, the transition to state **SP** is triggered. When entering the state **SP**, each foot-bot randomly selects and displays a color from the set of colors provided to the state. At the same time, each foot-bot starts incrementing an internal timer  $t$ . Whenever this timer  $t$  exceeds a fixed threshold  $\tau$ , the foot-bot examines the color displayed on the eye-bot to determine whether to remain in state **SP** or to leave the state and terminate the behavior. The timer mechanism provides the aerial robot sufficient time to perceive,

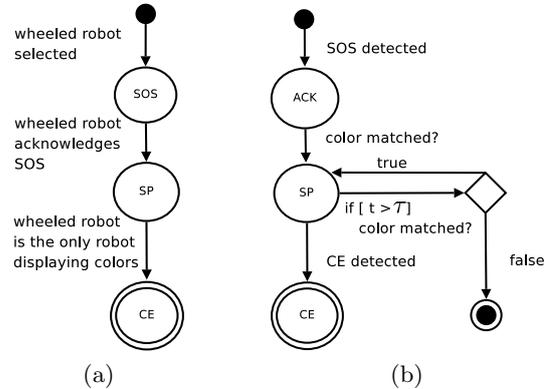


Figure 4.1: The finite state machines running on (a) the eye-bot and (b) the foot-bots.

process and react to the colors displayed by the foot-bots. When a foot-bot is in state **SP** and detects the **CE** color on the eye-bot, the foot-bot can safely assume that it is the robot with which the eye-bot wishes to communicate. In this case, the foot-bot confirms the termination of the selection process by displaying its **CE** color.

### 4.1.3 Simulation-based Experiments

We carried out experiments with the heterogeneous robotic platform to study the impact that the number of colors available to the selection process has on the number of iterations required for the termination of the selection process (i.e., completion time). We also investigate the scalability of our approach by varying the total number of foot-bots within the visual range of an eye-bot.

#### Experimental Setup

Each simulation run starts with an aerial robot placed in the center of a closed, obstacle-free arena (2 m x 2 m) at a height of 2 m. A number of foot-bots are randomly placed within the visual range of this eye-bot (see Fig. 4.2). The eye-bot is able to perceive all foot-bots within the arena and vice versa. Furthermore, each foot-bot is able to perceive neighboring foot-bots within a radius of 1 m.

Initially, the foot-bots perform a random walk while avoiding other robots and the arena wall. Their green LEDs are on so that they are visible to the eye-bot. The aerial robot then randomly picks a particular foot-bot and starts the one-to-one communication establishment process presented in Sect. 4.1.2. All foot-bots respond to the eye-bot and remain stationary.

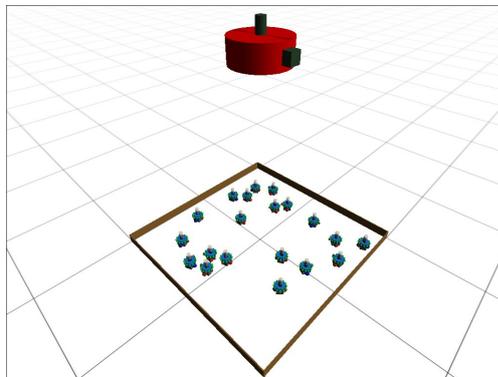


Figure 4.2: A screenshot from simulation including one eye-bot and twenty foot-bots.

## Results

We conducted five series of experiments varying the number of colors used in the selection process from 2 to 6. In each series, we increased the number of foot-bots within the visual range of an eye-bot from 10 to 80 in steps of 5. We ran 1000 replications for each combination of number of foot-bots and colors used. In Fig. 4.3, we plot the mean number of iterations spent on establishing communication between the eye-bot and a particular foot-bot.

The results in Fig. 4.3 show that the number of colors available to the selection process has a significant impact on the number of iterations required to establish a communication link. For instance, in the case of 2 colors and 20 foot-bots, the average number of iterations is 5.5 and the corresponding standard deviation (not shown in Fig. 4.3) is 1.9. On the other hand, in the case of 6 colors and 20 foot-bots, the average is 2.5 iterations and the standard deviation 0.8. The results show that the more colors available, the faster the termination of the selection process.

The results in Fig. 4.3 also show that for all series of experiments, the number of iterations needed for an eye-bot to establish communication with a particular foot-bot scales logarithmically with the number of foot-bots. In Tab. 4.1, we have listed the mean, the standard deviation, the minimum and maximum, and the upper bound as predicted by the model presented in Sect. 4.1.1 for the number of iterations spent by an eye-bot on establishing communication with a particular foot-bot. In all cases, 80 foot-bots were used. For all the different numbers of colors used, the mean number of iterations obtained in our simulation-based experiments are well below the upper bound predicted by the model. Furthermore, an interesting trend is apparent when considering the standard deviations: the more colors used, the lower the standard deviation. Hence, using more colors does not only reduce the number of iterations required, it also makes the number of iter-

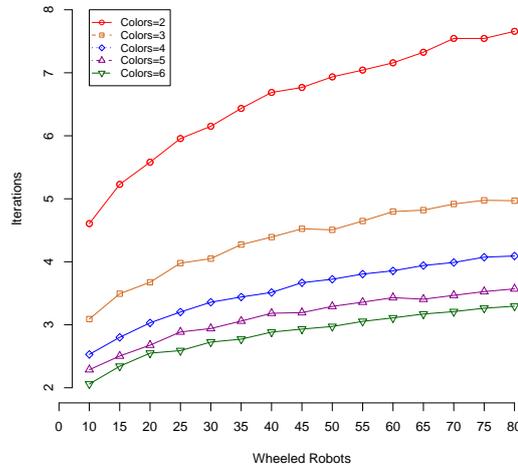


Figure 4.3: Results of the scalability experiments with different number of colors for the selection process. Each data point is the average of 1000 repetitions.

Table 4.1: Scalability results for 80 foot-bots (in number of iterations). 1000 replication were conducted for each experimental setup.

Colors	Mean	st.dev.	Min.	Max.	Upper bound
2	7.557	1.715136	4	17	9.3038
3	4.969	1.220729	3	11	6.4772
4	4.092	0.968749	2	10	5.4852
5	3.570	0.842502	2	8	4.9649
6	3.295	0.740624	2	6	4.6386

ations required more predictable.

Note that we empirically validated only the upper bound predicted by the model (see Eq. 4.2) using the results obtained from simulation. We expect the logarithmic growth of the upper bound to behave similarly for larger groups of robots.

#### 4.1.4 Real Robot Experiments

To confirm the real-world feasibility of our approach, we ran a series of experiments on the s-bot platform. Fig. 4.4 shows snapshots of a sample experiment run using 5 s-bots, in which we let the lone s-bot in the bottom row assume the role of the eye-bot. All robots are stationary and run the control program introduced in Sect. 4.1.2. The timer threshold  $\tau$  is set to 20 control steps (equivalent to 2 seconds). Note that the optimal value of  $\tau$  depends on the underlying hardware and that the value used here has not been fine-tuned. In the example in Fig. 4.4, a total of 3 colors are used by the

controller. The selection process is iterated four times before a one-to-one communication link is successfully established between the s-bot assuming the eye-bot's role and another s-bot. In the sample experiment illustrated in Fig. 4.4, a communication link was established after 9 seconds. We replicated the experiment 10 times using the same setup. On average, 3 iterations were required for the termination of the selection process. The video footage of the experiment shown in Fig. 4.4 and other proof-of-concept experiments can be found at <http://iridia.ulb.ac.be/supp/IridiaSupp2009-006/>.

## 4.2 One-To-Many Communication

In this section, we describe how an already established one-to-one communication link can be expanded to become a one-to-many communication link between an eye-bot and a group co-located foot-bots. We are not interested in which individual foot-bots are selected, but only in how many are selected. We describe how our approach can be used to either grow a group with a "lower-bounded" group size (i.e., the size of a grown group must be equal to or greater than a desired group size) or to grow a group with a size equal to a desired group size. Note that the choice between the two growth types may depend on the application.

Our approach works by iteratively growing a group of foot-bots around a seed robot with which a one-to-one communication link has already been established. In the first iteration, the seed robot is the sole member of the group. In each subsequent iteration, the eye-bot may send a request to increase the size of the group. Only foot-bots that are within visual range of an existing group member process this request. We refer to the robots in this range as *candidate* robots. At this point, robots that are not directly adjacent to the existing group (i.e., they detect other robots between themselves and the group) eliminate themselves as potential candidates. We refer to the remaining candidate robots as the *closest candidate robots*. These closest candidate robots now signal their candidacies to the eye-bot. The eye-bot completes the iteration by granting group membership to some or all the closest candidate robots, depending on the type of growth required. When "lower-bounded" group size is acceptable, the eye-bot can simply grant membership to all of the closest candidate robots. To achieve an exact group size, the eye-bot can request the closest candidate robots to relinquish their candidacies probabilistically. See Fig. 4.6 for an example of the algorithm running on real robots.

Below, we describe the robot controllers under the assumption that the eye-bots and the foot-bots can send and perceive the colors red, green and blue. We also present the results of our simulation-based studies comparing the two growth types. We demonstrate the approach on real robots.

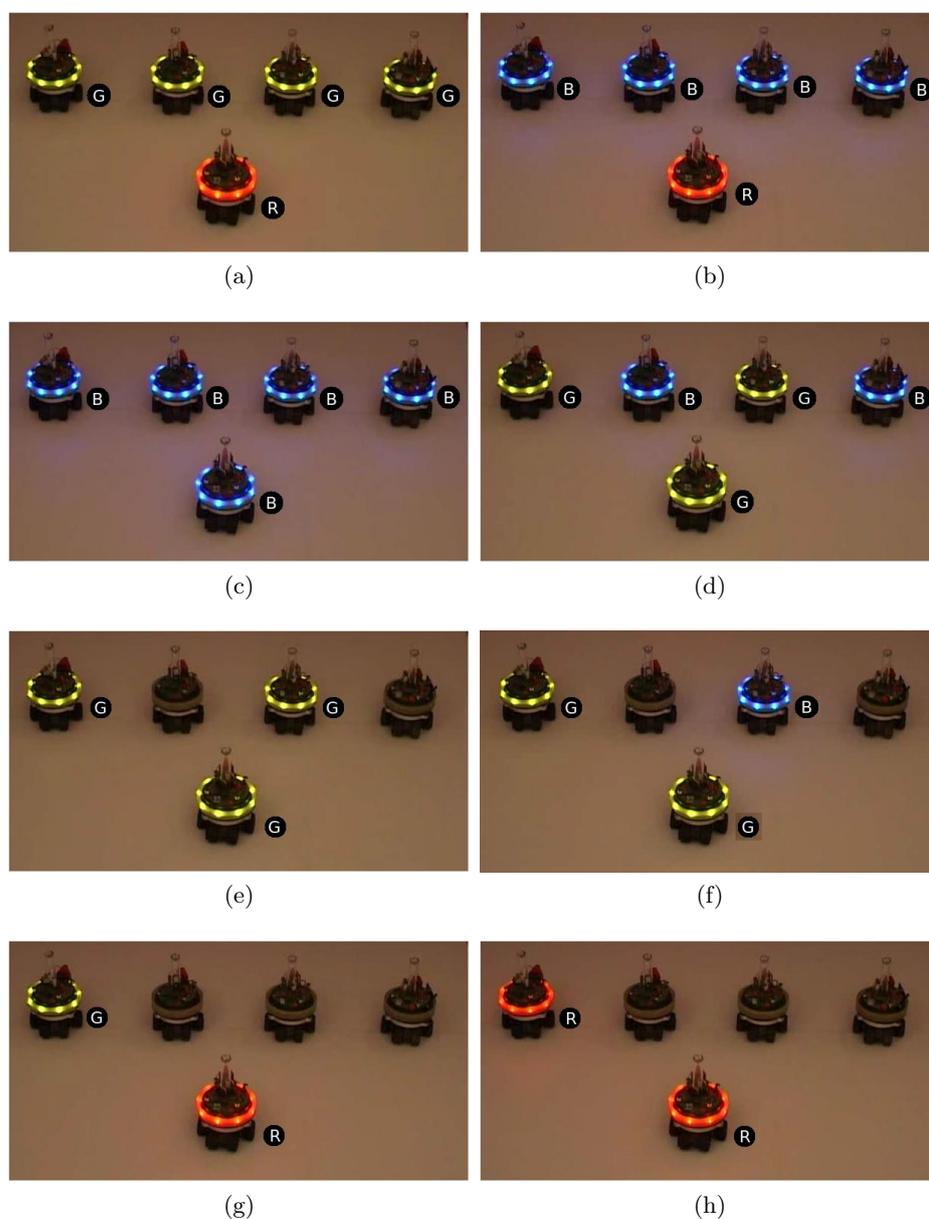


Figure 4.4: Snapshots of an experiment in which we let the lone s-bot in the bottom row assume the role of the eye-bot. This pre-designated s-bot seeks to establish a one-to-one communication link with the s-bot on the top-left. The letters next to the s-bots represent the current color displayed: R=red, G=green and B=blue. (a) An SOS is sent by the pre-designated s-bot. (b) The SOS is acknowledged by the other 4 s-bots. (c) The pre-designated s-bot initiates the selection process SP. (d) SP includes all 4 s-bots, (e) SP includes 2 remaining s-bots, (f) SP includes 2 remaining s-bots, (g) a spatially targeted communication link is established with the selected s-bot and (h) the establishment is confirmed by the selected s-bot.

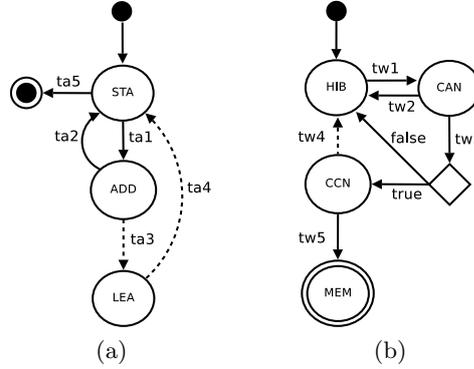


Figure 4.5: The FSMs running (a) on an eye-bot and (b) the foot-bots. Including the transitions in dashed lines results in an exact group size, whereas without them the final group size is lower-bounded to a desired group size.

#### 4.2.1 One-to-Many Controller

We developed one controller for the foot-bots and one controller for the eye-bots. Both controllers are behavior-based and are represented as FSMs in Fig. 4.5.

Fig. 4.5a shows the FSM implemented on the aerial robot. Consider the two states *STA* (stable group size) and *ADD* (add members). The state *STA* displays the color red and the state *ADD* the color green. The transition *ta1* is triggered if the number of foot-bots in red (i.e., in the group) is smaller than the desired group size. The transition *ta2* is triggered when the number of foot-bots displaying red or green (i.e., closest candidate robots) is equal to or greater than the desired group size. When the group size is reached or exceeded, the transition *ta5* terminates the behavior.

Fig. 4.5b shows the FSM implemented on the foot-bots and consists of 4 states *HIB* (hibernate), *CAN* (candidate), *CCN* (closest candidate) and *MEM* (group member). The LEDs are switched off while in state *HIB*, whereas blue is displayed while in state *CAN*, green while in state *CCN*, and red when in state *MEM*. The state transition *tw1* is triggered if the aerial robot illuminates green, at least one foot-bot in the visible range displays red, no other foot-bot displaying blue is perceived closer to the group<sup>2</sup> and no foot-bot in the visible range displays green. The transition *tw2* is triggered if a candidate robot sees another candidate robot closer to the group than itself. While in state *CAN*, a timer  $t$  is incremented. Whenever this timer  $t$  exceeds a

<sup>2</sup>Using its omnidirectional camera, a foot-bot first finds the closest red blob (closest group member). It divides its field of view into 8 equally sized slices and checks for blue blobs (candidate robots) in the slice containing the closest red blob and the two adjacent slices. If all of these blue blobs are further away than the closest red blob, the foot-bot assumes that it is the closest robot to the group.

given threshold  $\tau$ , transition **tw3** is triggered. The triggering of transition **tw3** depends on the outcome of a Bernoulli trial with probability  $p_j = 0.5$ : if successful, the transition is triggered otherwise the timer  $t$  is reset. This timer mechanism provides the candidate robots sufficient time to determine the closest candidate robots. Finally, the transition **tw5** is triggered when the eye-bot grants the membership to the group by displaying red.

The mechanism described so far allows the group size to reach a size larger than required. The state **LEA** and the transitions **ta3** and **ta4** (see Fig. 4.5a) allow the eye-bot to request the closest candidate robots to relinquish their candidacies. The eye-bot displays the color blue while in state **LEA**. If the eye-bot is in state **ADD** and the sum of the group members and the closest candidate robots exceeds the desired group size, the transition **ta3** is triggered. On the other hand, if the eye-bot is in state **LEA** and the sum of the group members and closest candidate robots is below or equal to the desired group size, the transition **ta4** is triggered and the control program returns to state **STA**.

Additionally, we extended the controller of the foot-bots with transition **tw4** as shown in Fig. 4.5b. If a foot-bot is in state **CCN** and the color blue is displayed by the eye-bot, a timer  $t$  is incremented. Whenever this timer  $t$  exceeds a given threshold  $\tau$ , transition **tw4** is triggered. The foot-bot utilizes the outcome of a Bernoulli trial with probability  $p_l = 0.5$  to decide whether to trigger transition **tw4** or to reset the timer  $t$ . These extended control programs allow the eye-bot to grow a group containing an exact number of foot-bots by iterating between the states **STA**, **ADD** and **LEA** until the desired group size is reached.

### 4.2.2 Simulation-based Experiments

We ran two types of experiments using the heterogeneous robotic platform. 1) Experiments to grow lower-bounded groups, and 2) experiments to grow exact group sizes. In order to study the differences between the two cases, we compared the number of iterations required by the eye-bot through the state **STA** before a given group size was reached. For both cases, we varied the required group size between 20, 40, 60 and 80 while keeping the total number of foot-bots at 80. We ran 1000 replications for each varied condition. The experiments were conducted in the experimental setup described in Sect. 4.1.3.

Tab. 4.2 summarizes the results obtained. For both cases studied, the mean number of iterations is shown. For the experiments in which only a lower-bounded group size was required, we also list the mean of the number of excess robots in the final groups. The results clearly show that an exact growth requires up to 5 times more time (for group size 20) than the lower-bounded growth. On the other hand, the lower-bounded growth adds around 11% (for group size 60) to 32% (for group size 20) excess robots to the group.

Table 4.2: Mean time (in number of iterations) to grow lower-bounded and exact group sizes (using 80 foot-bots).

Group size	Lower-bounded		Exact
	Mean	Excess robots	Mean
20	3.668	6.490	18.876
40	5.939	6.820	24.111
60	7.583	6.560	21.152
80	10.016	0	10.053

### 4.2.3 Real Robot Experiments

Fig. 4.6 shows snapshots of a proof-of-concept experiment we ran on the s-bot platform. We placed 4 s-bots in the shape of an arch around a pre-designated s-bot, which we let assume the role of the eye-bot. All the other robots ran the control logic shown in Fig. 4.5b (but without transitions `ta3`, `ta4` and `tw5`). The timer threshold  $\tau$  is set to 20 control steps (equivalent to 2 seconds). Once the one-to-one communication link was established, the expansion of a one-to-many communication link to include a further s-bot took 2 s. The video footage of the experiment shown in Fig. 4.6 and other proof-of-concept experiments can be found at <http://iridia.ulb.ac.be/supp/IridiaSupp2009-006/>.

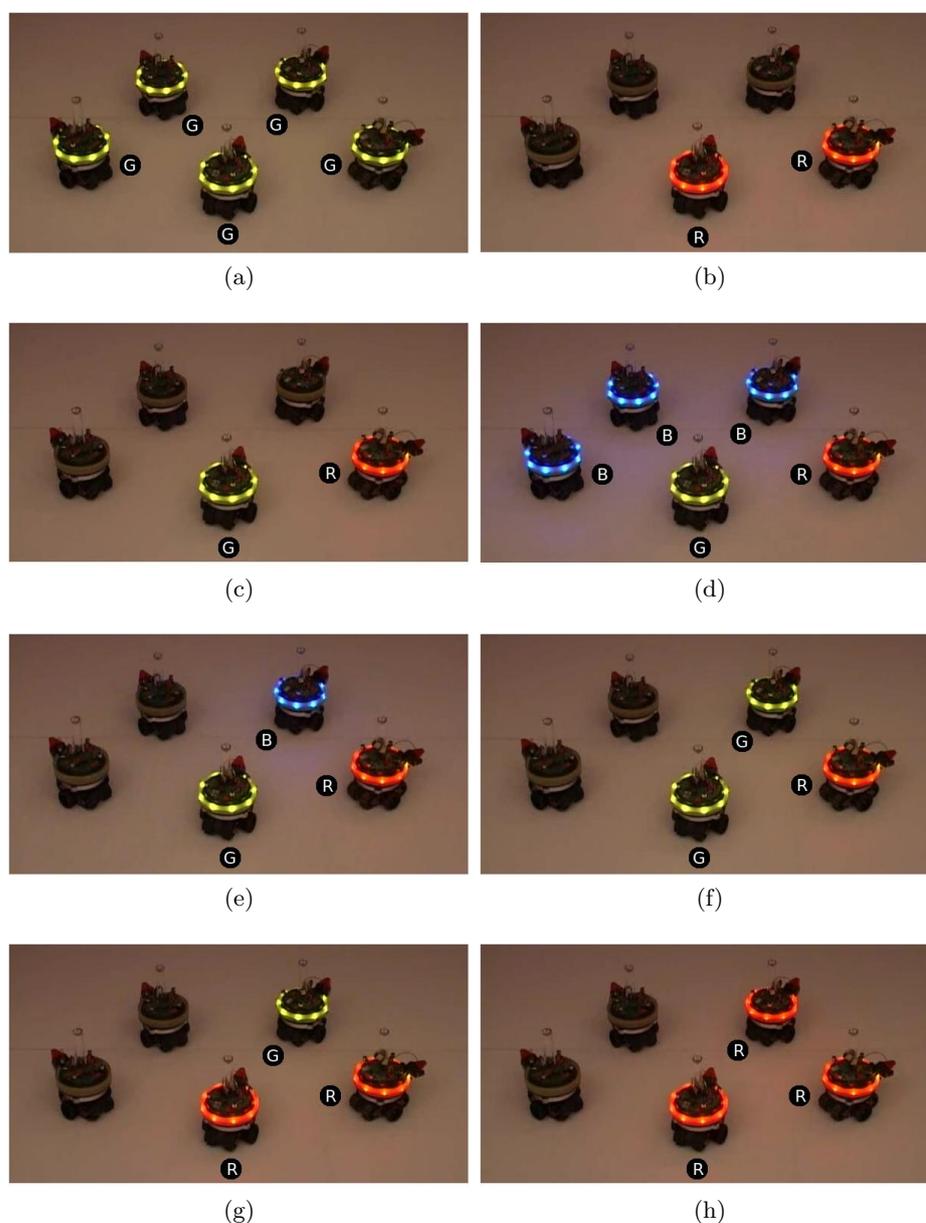


Figure 4.6: Snapshots of an experiment in which we let the s-bot in the center assume the role of the eye-bot. This predesignated s-bot seeks to grow a group of size 2. The letters next to the s-bots represent the current color displayed: R=red, G=green and B=blue. (a) Experiment initialization. (b) A one-to-one communication link is established. (c) The predesignated s-bot requests for more group members. (d) Three s-bots candidate by illuminating blue. (e) The closest candidate robot is determined. (f) The closest candidate robot signals its candidacy by illuminating green. (g) The membership to the group is granted by the predesignated s-bot. (h) The membership is confirmed by the new group member.



## Chapter 5

# Cooperation through Spatially Targeted Communication

We achieve cooperation in the heterogeneous swarm by letting the eye-bot establish a communication link (as explained in Chapter 4) with a group of foot-bots with an appropriate size (minimum number of foot-bots required to cross the gap) and appropriately located (near the gap). We then let the eye-bot use the spatially targeted communication link to transmit instructions to these selected foot-bots to form the shortest line morphology (i.e., target morphology) that will allow the foot-bots to cross the gap. In what follows, we describe SWARMORPH-script [6], which is the enabling mechanism for this transmission of instructions and how the SWARMORPH-script can be used by the selected foot-bots to generate of the target morphology.

### 5.1 Communication Between an Eye-bot and the Foot-bots

After the eye-bot has selected a suitable number of foot-bots, it sends instructions to the foot-bots on how to self-assemble the target morphology. These instructions are sent using SWARMORPH-script [6]. SWARMORPH-script is a language for distributed self-assembly and morphology control for autonomous self-assembling robots. Each instruction in SWARMORPH-script corresponds to a basic robot behavior.

Before transmitting the SWARMORPH-script for the target morphology, the eye-bot first translates the script into a binary string. The string is then sent to the foot-bots using a protocol developed based on LEDs and camera-based communication modality [6]. A foot-bot that receives such a binary string can, in turn, translate it back to a SWARMORPH-script and

execute this received control logic. In this manner, the foot-bots do not need to have any a priori knowledge about possible morphologies required or even possible tasks.

The transmission of a SWARMORPH-script requires the three colors red, green and blue. Each time a bit is transmitted, the eye-bot changes the illumination of its LEDs. The color green represents a ‘0’ bit, blue represents a ‘1’ and red represents a repeat bit. We rely on acknowledgment to distinguish adjacent bits. A receiving foot-bot acknowledges the receipt of each bit by lighting up its LEDs to match the color of the sender’s LEDs. Once the receipt of a bit has been acknowledged, the sending robot transmits the next bit. This acknowledgement mechanism necessitates our use of the dedicated color for a repeat bit. When transmitting a substring of two or more bits of the same value, every other bit will be represented by the color red (starting from the second bit).

## 5.2 Morphology Generation Using Directional Self-Assembly

The foot-bots which received the SWARMORPH-script make use of a directional self-assembly mechanism similar to the one proposed in [27] to generate the target morphology described by the script. Using directional self-assembly, a foot-bot can invite other foot-bots to connect to it by illuminating its LEDs in a particular pattern. Foot-bots connected to each other communicate using the range and bearing communication system to control the morphology generation process. Typically, morphologies are initiated by a single robot. In this study, the *seed* foot-bot selected by the eye-bot (see Chapter 4) is always the morphology initiating robot.

We refer to these patterns as *connection slots* and to a foot-bot displaying a connection slot as an *extending foot-bot*. Whenever a free (unconnected) foot-bot sees a connection slot, it attempts to connect to the connection slot at extending foot-bot. By changing the LEDs it is illuminating, an extending foot-bot can change the location on its LED ring to which a connecting foot-bot will attach. Once a foot-bot has connected to the extending foot-bot, it uses the range and bearing communication system to receive instructions from the extending foot-bot on how to either extend the morphology or terminate morphology growth [6]. This communication modality is also used by neighboring foot-bots to notify each other about their internal states while they are connected.

## 5.3 Example SWARMORPH-Script

Script 1 is an example of a complete SWARMORPH-script for crossing a gap in a 3 foot-bot line. The seed foot-bot first rotates in place to face the light

source. Then it opens a new connection slot to its rear in order to invite one of the other recruited foot-bots to physically connect. When a foot-bot has connected, the seed has to wait for a notification from the newly connected foot-bot that indicates that the morphology is complete before starting to cross the gap. When a foot-bot connects to the seed, it receives rule ID 0. This tells the newly connect foot-bot that it is the middle foot-bot in the line morphology and that it has to receive a connection from the last foot-bot in order to complete the morphology. The middle foot-bot therefore opens a connection slot to its rear. When the last foot-bot connects to the middle foot-bot, the middle foot-bot sends it ID 1 to tell the newly connected foot-bot that it is the last foot-bot in the morphology. The middle foot-bot then notifies the seed that the morphology is complete and starts to move across the gap. The foot-bot that connected to the morphology last starts to move across the gap immediately after the middle foot-bot has sent it rule ID 1.

---

**Script 1:** A script sent by the eye-bot to form a line of foot-bots of length three.

---

```

if seed then
  // Instructions for seed foot-bot:
  RotateToFaceLightSource();
  OpenConnSlot(rear);
  SendRuleID(0);
  WaitNotify(rear);
  while no new notification from rear do
    Phototaxis();
  end
end
if not seed then
  SearchForConnSlot();
  ReceiveRuleID();
  if receivedruleid = 0 then
    // Instructions for middle foot-bot:
    OpenConnSlot(rear);
    SendRuleID(1);
    Notify(front);
    while no new notification from rear do
      Phototaxis();
    end
    Notify(front);
    Disconnect();
  end
  if receivedruleid = 1 then
    // Instructions for last foot-bot:
    while gap not crossed do
      Phototaxis();
    end
    Notify(front);
    Disconnect();
  end
end

```

---

The foot-bot that connected to the morphology last is also the last foot-bot to cross the gap. When it has crossed the gap (when its infrared ground sensors register solid group again), it sends a notification using the range

and bearing communication system to the middle foot-bot and disconnects from the morphology. When the middle foot-bot receives the notification, it forwards the notification to the seed and disconnects. When the seed receives the notification, it knows that the other foot-bots have disconnected.

## 5.4 Simulated Cooperation in the Heterogeneous Swarm

The heterogeneous swarm of robots presented in this work is still under development. Therefore, we have implemented the proposed approach in simulation. However, note that the two core parts of the approach proposed in this work, namely establishing spatially targeted communication [22] and the SWARMORPH-script [6], have been successfully tested on real robotic hardware.

In Fig. 5.1, we provide a sequence of screenshots from the simulation, in which the heterogeneous swarm successfully manages to solve the gap crossing task. We assume that the eye-bot is capable of detecting the gap and estimating the width of the gap using the onboard image processing software. Using this knowledge, the eye-bot generates a SWARMORPH-script (see Sect. 5.3) which can be sent to the foot-bots to generate the shortest line morphology required to cross the gap. Once the eye-bot detects a foot-bot which is close enough to the gap, but yet far enough to be circled by other foot-bots during the self-assembly process, the eye-bot sends an SOS signal to establish a spatially targeted communication link with this and other co-located foot-bots. The communication link is then used by the eye-bot to transmit the SWARMORPH-script to this selected group of foot-bots.

In this experiment, 10 foot-bots are randomly placed in the start zone and perform a phototaxis behavior and do not have any a priori knowledge about the environment or the gap. Each foot-bot performs the phototaxis until an eye-bot is detected, in which case they stop moving. Subsequently, each foot-bot participates in the process of establishing spatially targeted communication initiated by the eye-bot. If a communication link has been successfully created with the eye-bot, a SWARMORPH-script is received through the link. In a subsequent step, the received script is executed. On the other hand, if a foot-bot is not selected, it performs an anti-phototaxis behavior.

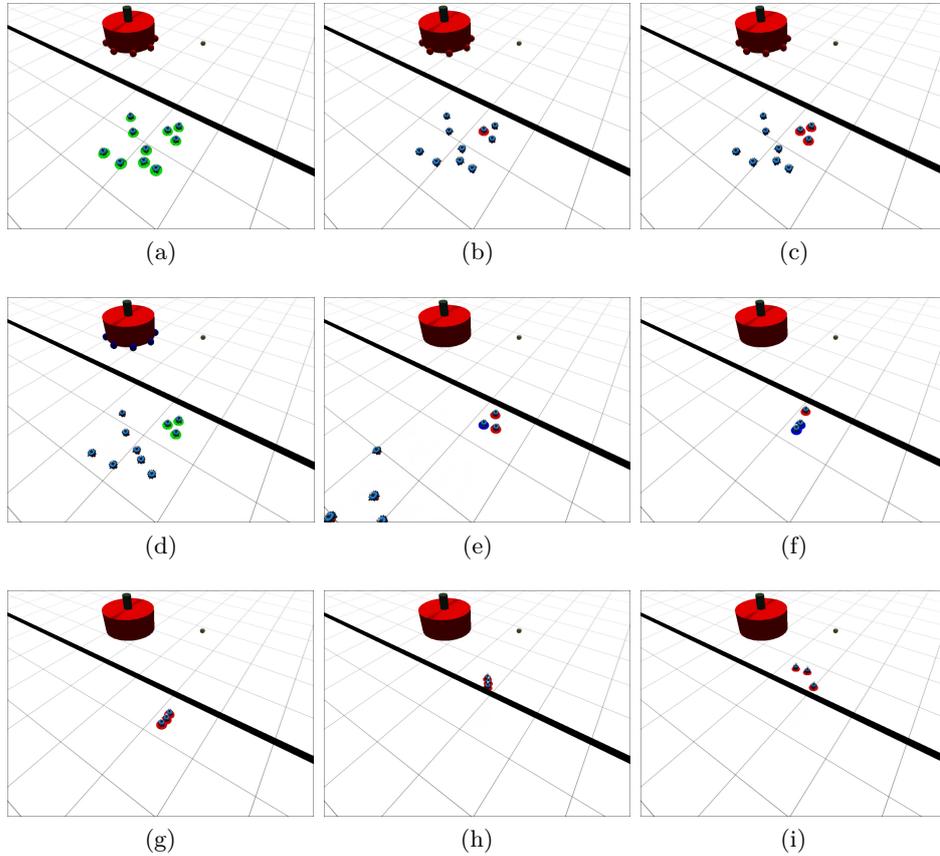


Figure 5.1: Screenshots from the simulation showing cooperation in the heterogeneous swarm. (a) A group of 10 foot-bots move towards the light source and do not have any a priori knowledge about the gap they are reaching. A stationary eye-bot, i.e., it is attached to the ceiling, is assumed to have detected the gap and estimated the width of the gap. (b) The eye-bot establishes spatially targeted communication link to a well located seed foot-bot. (c) This communication link is expanded to include two further co-located foot-bots. (d) The SWARMORPH-script required to form the smallest possible line morphology is transferred through the spatially targeted communication link. (e) The seed foot-bot has aligned to the light source, and has opened a connection slot on its rear. (f) A first foot-bot has attached to the seed foot-bot. (g) The line morphology has been generated in the received script is completed and the foot-bots move towards the light source. (h) The foot-bots have successfully managed to cross the gap. (i) The foot-bots have disconnected from each other and are now located in the target area.



## Chapter 6

# Experiments and Results

We run simulation-based experiments using three different control strategies of the heterogeneous swarm. For each combination of gap size and control strategy, we ran 100 experimental trials. By comparing the task completion times of the three strategies, we first analyse the benefits of cooperation through spatially targeted communication, and then isolate the benefits of spatially targeted communication. Videos of experiments conducted using all three control strategies are available online at: <http://iridia.ulb.ac.be/supp/IridiaSupp2010-007/>.

### 6.1 Experimental Setup

The heterogeneous swarm is located in the obstacle-free environment shown in Fig. 1.1. The environment consists of a start zone, a target zone, and a gap that separates the two zones. A light source perceivable by the foot-bots is located in the target zone. At the start of each experiment, 10 foot-bots are placed at random positions with random orientations within a square area of 2 m x 2 m in the start zone. The foot-bots use their light sensors to detect the light source in the target zone. They use their proximity sensors to avoid collisions with other robots. They use the ground sensors placed on their chassis to detect if they are close to the gap. An eye-bot is placed on the ceiling in the start zone and is able to perceive all the foot-bots in the start zone. The eye-bot can estimate the width of the gap by using its pan-and-tilt camera and the on-board image processing software.

To reach the target zone, the foot-bots may need to connect to each other to form a collective morphology, such as a line morphology [6]. Note that the required minimal length of such a line morphology (i.e., the number of foot-bots in the line) which guarantees a safe crossing of the gap depends on the width of the gap. In this study, we vary the width of the gap between 5 cm, 10 cm, 15 cm and 25 cm. These different gap widths require the foot-bots to form a line morphology of 1, 2, 3 and 4 foot-bots respectively.

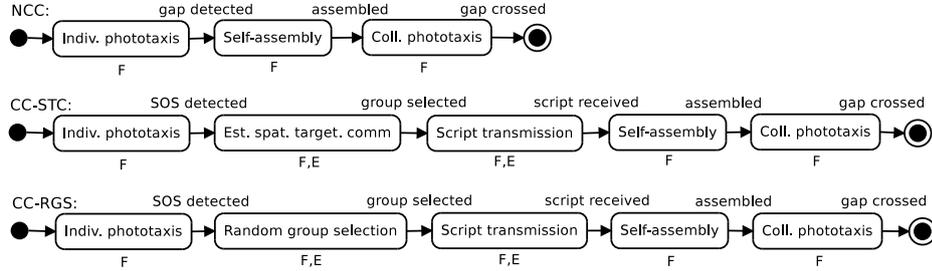


Figure 6.1: Decomposition of control strategies into phases. Phases only involving foot-bot are marked ‘F’, phases involving foot-bot eye-bot cooperation are marked ‘F.E’. i) NCC: non-cooperative control, ii) CC-STC: cooperative control with spatially targeted communication and iii) CC-RGS: cooperative control with random group selection. NB ‘Indiv. phototaxis’ = ‘Individual phototaxis’, ‘Coll. phototaxis’ = ‘Collective phototaxis’, ‘Est. spat. target. comm.’ = ‘Establishing Spatially Targeted Communication’.

The task is considered to be completed when the final foot-bot of the line morphology has crossed the gap and reached the target zone.

## 6.2 Three Control Strategies

**Non-cooperative control (NCC):** The foot-bots are provided with a SWARMORPH-script for the formation a 4 foot-bot line when a gap is encountered (regardless of the width of the gap). The foot-bots do not communicate nor do they attempt to cooperate with an eye-bot. They initially move towards the light source until one of the foot-bots detects the gap. The foot-bot retreats approximately 40 cm from the gap and becomes the seed by inviting other foot-bots to connect to its rear. Other foot-bots that are able to see the invitation stop moving towards the light source and attempt to connect to the morphology. Once the line of 4 foot-bots is formed, the foot-bots move towards the light source to cross the gap.

**Cooperative control with spatially targeted communication (CC-STC):** The foot-bots do not have a priori knowledge about the task or the target morphology. The foot-bots initially move towards the light while avoiding the gap until the eye-bot illuminates its red LEDs. Spatially targeted communication link is then established with optimally situated foot-bots as described in Section 4. The eye-bot selects a foot-bot that is approximately 40 cm away from the gap to be the seed leaving enough room for the free foot-bots to maneuver

around the morphology during the self-assembly process. Then, a certain number of immediate neighbors of the seed robot are selected to form the target morphology. The selected foot-bots receive the SWARMORPH-script from the eye-bot and follow the instructions in the script to self-assemble into the target morphology. Once the morphology is formed, the foot-bots move towards the light to cross the gap. The non-selected foot-bots move away from the light.

**Cooperative control with random group selection (CC-RGS):** The foot-bots do not have a priori knowledge about the task or the target morphology. The foot-bots initially move towards the light until the eye-bot illuminates its red LEDs. The eye-bot then randomly selects a seed foot-bot (i.e., without considering its location in the environment with respect to the gap) and further foot-bots (i.e., without considering their relative distance to the seed foot-bot) required to form the target morphology. The selected foot-bots receive the SWARMORPH-script from the eye-bot and follow the instructions in the script to self-assemble into the target morphology. Once the morphology is formed, the foot-bots move towards the light to cross the gap. The non-selected foot-bots move away from the light.

### 6.3 Performance Benefits of Cooperation with Spatially Targeted Communication

To analyse the performance benefits of cooperation with spatially targeted communication, we compare the task execution times of strategies NCC and CC-STC. The results are shown in Fig 6.2. In the case of NCC, we have only plotted the results of the narrowest gap of 5 cm, as the task completion times between the various gap widths did not prove to be significantly different for the NCC strategy.

In all the experiments, the foot-bots solved the task. According to the results in Fig 6.2, the median task completion times of CC-STC are 507, 2590 and 4032 simulation steps for gap width 5 cm, 10 cm and 15 cm, respectively. Compared to the median task completion time of NCC (4340 simulation steps), the mean completion times for CC-STC were respectively 88%, 40% and 7% lower in environments with gaps that can be crossed by 1, 2 or 3 physically connected foot-bots. This is due to the fact that in CC-STC, the length of the line is optimal with respect to the width of the gap. The cooperation with the eye-bot avoids the inclusion of excess robots in the morphology. For the widest gap, namely the gap of 25 cm that can only be crossed by four or more physically connected foot-bots, NCC is, in general, faster than CC-STC. Intuitively, this could have been expected given that both control strategies (i.e., CC-STC and NCC) form a line of

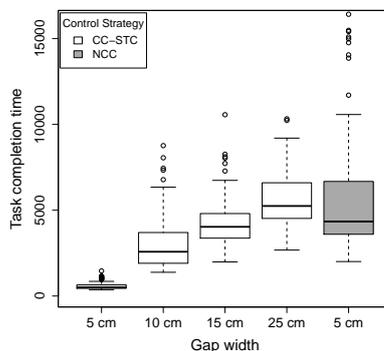


Figure 6.2: Box-and-whisker plot showing task completion times (in simulation steps) of two different control strategies.

four robots close to the gap, but in the case of CC-STC, instructions have to be first received from the eye-bot before the self-assembly process can start and therefore requires more time. However, the NCC strategy has several outlier trials that take very long to complete. This is because in the NCC strategy all foot-bots try to participate in the construction of the morphology. Thus there are still foot-bots trying to join the morphology even after the morphology is completed, which can interfere (sometimes severely) with the collective phototaxis of the complete morphology.

In Fig. 6.3, we have plotted a breakdown of how much time is spent in the different phases of each control strategy: (i) establishing spatially targeted communication (CC-STC), (ii) transmitting the SWARMORPH-script (CC-STC) (iii) self-assembly (CC-STC), (iv) self-assembly (NCC). The results show that with the increasing size of the morphology, and therefore with the increasing length of the SWARMORPH-script that has to be transmitted, the transmission time increases. However, this communication overhead of CC-STC would become negligible if a communication modality with higher bandwidth (such as WiFi) could be utilized. The results also show that when a line of equal length is formed in both control strategies, as in the case of 4 foot-bots, the self-assembly process of CC-STC requires on average 39% more time than that of NCC. This can be explained by the fact that in NCC there are more robots attempting to connect to a connection-inviting foot-bot which in turn leads to faster morphology formation. On the other hand, CC-STC deals with the resources optimally by only allocating precisely the required number of robots needed for the self-assembly process. The decision involving this trade-off between faster morphology formation times and optimal resource allocation may depend on the task and/or the priorities of the system.

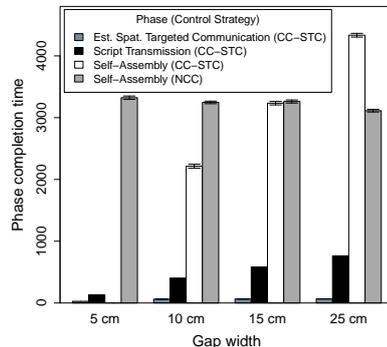


Figure 6.3: Bar-plot showing a breakdown of the time spent in different phases of two different control strategies (bars appended with the standard deviation).

## 6.4 Performance Benefits of Spatially Targeted Communication

To isolate the performance benefits of spatially targeted communication, we compare the task completion timings of strategies CC-STC and CC-RGS. Note that both control strategies use the same iterative selection process to select the seed foot-bot. However, the selection of further foot-bots required in the target morphology is different. Therefore, in order to maintain objectivity in the comparison, in this set of experiments the time spent to select the non-seed foot-bots was omitted for both control strategies. The results are plotted in Fig. 6.4.

As the results show, CC-STC was on average faster than CC-RGS in all cases studied independent of the width of the gap. The explanation for these results is that a morphology formed next to the gap in most cases require less time to reach and cross the gap than a morphology formed at a random place in the environment. We expect that this difference in terms of task completion time would be even greater for larger start zones.

Additionally, we also studied the difference in completion times between CC-STC and CC-RGS by conducting additional experiments (not described above) in a walled environment: the foot-bots were placed in the start zone within an area of 2 m x 2 m surrounded by walls on three sides to adjoin the gap on the fourth side. We found that the presence of the walls had no significant impact on the completion time of CC-STC in which the eye-bot selects the seed and the group in favorable locations (i.e., always close to the gap and away from the walls). For the CC-RGS control strategy, on the other hand, the presence of walls had a significant negative impact on performance. When the randomly selected seed happened to be located close to one of the walls, it could be difficult or even impossible for foot-bots

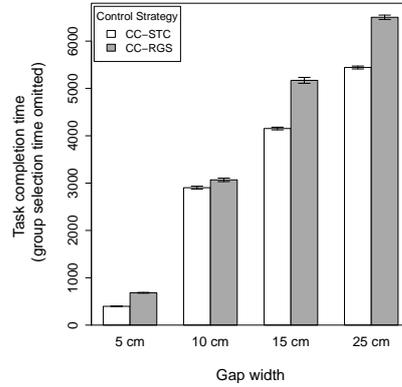


Figure 6.4: Completion times (in simulation steps) of strategies CC-STC and CC-RGS minus the time taken to form the group (bars appended with the standard deviation).

to physically connect to an inviting foot-bot. As a result, the task was not solved in our experiments with the CC-RGS control strategy in 13%, 29% and 34% of the experiments for the line morphology composed of 2 foot-bots, 3 foot-bots and 4 foot-bots, respectively.

It should be noted that we have only considered tasks that can be solved by all strategies. We have not considered environments with gaps too wide to be crossed, nor have we considered the presence of obstacles or other features that would require a cooperative strategy to solve the task.

## Chapter 7

# Conclusions and Future Work

In this work, we have demonstrated a novel approach to establish spatially targeted communication between eye-bots and one or more co-located foot-bots. We showed that LEDs and cameras can be used as a situated communication modality to establish spatially targeted communication. We presented a probabilistic model which gives an upper bound on the average time required to establish a spatially targeted communication link between an eye-bot and a specific wheeled robot. We also showed how such a one-to-one communication link can be expanded to a one-to-many communication link between an eye-bot and a group of co-located robots.

In simulation, we demonstrated that the approach scales well and that it remains within the bounds predicted by the model. On real robotic hardware, we demonstrated the approach through a series of proof-of-concept experiments. Although experiments were performed using on-board LEDs and on-board cameras, any scalable, situated communication modality that allows robots to communicate their internal state to nearby robots could be used.

We have also demonstrated how eye-bots and foot-bots can cooperate. We have shown how, by cooperating, the swarm was able to carry out different instances of a gap crossing task without requiring any global information and without a priori knowledge about the task. Compared to a non-cooperative strategy, the cooperative strategy was shown to be more efficient in terms of resource allocation as the eye-bot recruited only the robots necessary based on the width of a gap. Furthermore, the cooperative strategy led to faster task completion times in the environment in which fewer than four connected robots could cross the gap.

We demonstrated the benefits of cooperation through spatially target communication. When the eye-bot recruited foot-bots based on the location in the environment and based on their mutual proximity, the time required

to self-assemble and to cross the gap was lower than when robots were randomly selected.

In our ongoing work, we are experimenting with prototypes of the heterogeneous swarm, and we expect to have results in the near future. Given the limited throughput of the LEDs and camera based communication modality, we are considering to follow up an established spatially targeted communicated link with a high bandwidth communication modality, such as a standard WiFi link, to actually let the robots communicate to each other.

Our long-term goal is to repeat the experiments in this work on real robotic hardware. In ongoing research, we are investigating other cooperation mechanisms between eye-bots and foot-bots, in particular where the cooperation is more bidirectional. In this study, the foot-bots passively received instructions from the eye-bots. In the future, foot-bots on the ground could ask an eye-bot to find additional robots for a given task, and multiple eye-bots could allocate and share groups of foot-bots dynamically.

# Bibliography

- [1] Krishna B. Athreya and Peter E. Ney. *Branching Processes*. Springer-Verlag, Berlin, Germany, 1972.
- [2] D. H. Barnhard, J. T. McClain, Bradley J. Wimpey, and Walter D. Potter. Odin and Hodur: Using bluetooth communication for coordinated robotic search. In *IC-AI*, pages 365–371. CSREA Press, Athens, GA, USA, 2004.
- [3] H. B. Brown, J. M. V. Weghe, C. A. Bererton, and P. K. Khasla. Millibot trains for enhanced mobility. *IEEE/ASME Transactions on Mechatronics*, 7(4):452–461, 2002.
- [4] Z. Butler, K. Kotay, D. Rus, and K. Tomita. Generic decentralized control for lattice-based self-reconfigurable robots. *International Journal of Robotics Research*, 23(9):919–937, 2004.
- [5] Alexandre Campo, Shervin Nouyan, Mauro Birattari, Roderich Groß, and Marco Dorigo. Negotiation of goal direction for cooperative transport. In *Ant Colony Optimization and Swarm Intelligence: 5th International Workshop, ANTS 2006*, volume 4150 of *Lecture Notes in Computer Science*, pages 191–202. Springer-Verlag, Berlin, Germany, 2006.
- [6] A. L. Christensen, R. O’Grady, and M. Dorigo. SWARMORPH-script: A language for arbitrary morphology generation in self-assembling robots. *Swarm Intelligence*, 2(2-4):143–165, 2008.
- [7] A. L. Christensen, Rehan O’Grady, and Marco Dorigo. Morphology control in a self-assembling multi-robot system. *IEEE Robotics & Automation Magazine*, 14(4):18–25, 2007.
- [8] R. Damoto, A. Kawakami, and S. Hirose. Study of super-mechano colony: concept and basic experimental set-up. *Advanced Robotics*, 15(4):391–408, 2001.
- [9] T. Fukuda, M. Buss, H. Hosokai, and Y. Kawauchi. Cell structured robotic system CEBOT: control, planning and communication methods. *Robotics and Autonomous Systems*, 7(2-3):239–248, 1991.

- [10] T. Fukuda, M. Buss, H. Hosokai, and Y. Kawauchi. Cell structured robotic system CEBOT: Control, planning and communication methods. *Robotics and Autonomous Systems*, 7(2-3):239–248, 1991.
- [11] T. Fukuda and S. Nakagawa. Approach to the dynamically reconfigurable robotic system. 1(1):55–72, 1988.
- [12] Robert Grabowski and Pradeep Khosla. Localization techniques for a team of small robots. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, volume 2, pages 1067–1072. IEEE Press, Piscataway, NJ, USA, 2001.
- [13] R. Groß and M. Dorigo. Self-assembly at the macroscopic scale. *Proceedings of the IEEE*, 96(9):1490–1508, 2008.
- [14] Roderich Groß, Michael Bonani, Francesco Mondada, and Marco Dorigo. Autonomous self-assembly in swarm-bots. *IEEE Transactions on Robotics*, 22(6):1115–1130, 2006.
- [15] Patsy Haccou, Peter Jagers, and Vladimir Vatutin. *Branching Processes: Variation, Growth, and Extinction of Populations*. Cambridge University Press, Cambridge, UK, 2005.
- [16] Theodore E. Harris. *The Theory of Branching Processes*. Springer-Verlag, Berlin, Germany, 1963.
- [17] S. Hirose, T. Shirasu, and E. F. Fukushima. Proposal for cooperative robot “Gunryu” composed of autonomous segments. *Robots and Autonomous Systems*, 17:107–118, 1996.
- [18] P. Jagers, F.C. Klebaner, and S. Sagitov. On the path to extinction. *Proceedings of the National Academy of Science of the United States of America*, 104:6107–6111, 2007.
- [19] Y. Kawauchi, M. Inaba, and Fukuda T. A principle of distributed decision making of cellular robotic system (CEBOT). pages 833–838. IEEE Press, Piscataway, NJ, 1993.
- [20] D.G. Kendall. Branching processes since 1873. *Journal of London Mathematics Society*, 41:385–406, 1966.
- [21] E. Klavins, R. Ghrist, and D. Lipsky. A grammatical approach to self-organizing robotic systems. *IEEE Transactions on Automatic Control*, 51(6):949–962, 2006.
- [22] Nithin Mathews, Anders Lyhne Christensen, Eliseo Ferrante, Rehan O’Grady, and Marco Dorigo. Establishing spatially targeted communication in a heterogeneous robot swarm. In *Proceedings of 9th Inter-*

- national Conference on Autonomous Agents and Multiagent Systems (AAMAS 2010)*. In press.
- [23] J. T. McClain, B. J. Wimpey, D. H. Barnhard, and W. D. Potter. Distributed robotic target acquisition using bluetooth communication. In *Proceedings of the 42nd annual Southeast regional conference*, pages 291–296. ACM, New York, NY, USA, 2004.
- [24] F. Mondada, L. M. Gambardella, D. Floreano, S. Nolfi, J.-L. Deneubourg, and M. Dorigo. The cooperation of swarm-bots: Physical interactions in collective robotics. *IEEE Robotics & Automation Magazine*, 12(2):21–28, 2005.
- [25] S. Murata, E. Yoshida, A. Kamimura, H. Kurokawa, K. Tomita, and S. Kokaji. M-tran: Self-reconfigurable modular robotic system. *IEEE-ASME Transactions on Mechatronics*, 7(4):431–441, 2002.
- [26] Shervin Nouyan, Roderich Groß, Marco Dorigo, Michael Bonani, and Francesco Mondada. Group transport along a robot chain in a self-organised robot colony. In T. Arai, R. Pfeifer, T. Balch, and H. Yokoi, editors, *Intelligent Autonomous Systems 9 – IAS 9*, pages 433–442. IOS Press, Amsterdam, The Netherlands, 2006.
- [27] R. O’Grady, A. L. Christensen, , and M. Dorigo. SWARMORPH: multi-robot morphogenesis using directional self-assembly. *IEEE Transactions on Robotics*, 25(3):738–743, 2008.
- [28] R. O’Grady, A. L. Christensen, C. Pinciroli, and M. Dorigo. Robots autonomously self-assemble into dedicated morphologies to solve different tasks (extended abstract). In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2010)*. In press.
- [29] Anthony G. Pakes. Asymptotic results for the extinction time of markov branching processes allowing emigration, I. Random Walk Decrements. *Advances in Applied Probability*, 21(2):243–269, 1989.
- [30] Carlo Pinciroli. Object retrieval by a swarm of ground based robots driven by aerial robots. Mémoire de DEA, Université Libre de Bruxelles, Brussels, Belgium, 2007.
- [31] Jim Pugh and Alcherio Martinoli. Relative localization and communication module for small-scale multi-robot systems. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 188–193. IEEE Press, Piscataway, NJ, USA, 2006.

- [32] Frédéric Rivard, Jonathan Bisson, François Michaud, and Dominic Létourneau. Ultrasonic relative positioning for multi-robot systems. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 323–328. IEEE Press, Piscataway, NJ, USA, 2008.
- [33] James F. Roberts, Timothy S. Stirling, Jean-Christophe Zufferey, and Dario Floreano. 2.5d Infrared Range and Bearing System for Collective Robotics. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'2009)*. IEEE Press, Piscataway, NJ, USA, 2009.
- [34] Anthony (Tony) Stentz, Alonzo Kelly, Herman Herman, Peter Rander, Omead Amidi, and Robert Mandelbaum. Integrated air/ground vehicle system for semi-autonomous off-road navigation. In *Proceedings of AUVSI Unmanned Systems Symposium*, 2002.
- [35] Kasper Støy. Using situated communication in distributed autonomous mobile robotics. In *SCAI '01: Proceedings of the Seventh Scandinavian Conference on Artificial Intelligence*, pages 44–52. IOS Press, Amsterdam, The Netherlands, 2001.
- [36] Richard T. Vaughan, Gaurav S. Sukhatme, Francisco J. Mesa-Martinez, and James F. Montgomery. Fly spy: Lightweight localization and target tracking for cooperating air and ground robots. In *Proceedings of the International Symposium on Distributed Autonomous Robot Systems*, pages 315–324. Springer-Verlag, Berlin, Germany, 2000.
- [37] M. Yim, D. G. Duff, and K. D. Roufas. Polybot: a modular reconfigurable robot. volume 1, pages 514–520. IEEE Press, Piscataway, NJ, 2000.
- [38] M. Yim, K. Roufas, D. Duff, Y. Zhang, C. Eldershaw, and S. B. Homans. Modular reconfigurable robots in space applications. *Autonomous Robots*, 14(2-3):225–237, 2003.
- [39] M. Yim, W. M. Shen, B. Salemi, D. Rus, M. Moll, H. Lipson, E. Klavins, and G. S. Chirikjian. Modular self-reconfigurable robot systems. *IEEE Robotics & Automation Magazine*, 14(1):43–52, 2007.