

Establishing Spatially Targeted Communication in a Heterogeneous Robot Swarm

N. Mathews¹, A. L. Christensen², E. Ferrante¹, R. O’Grady¹ and M. Dorigo¹

¹ IRIDIA, Université Libre de Bruxelles, Brussels, Belgium
{nmathews,eferrant,rogrady,mdorigo}@ulb.ac.be

² Instituto de Telecomunicações, Lisbon, Portugal
anders.christensen@iscte.pt

ABSTRACT

We consider a heterogeneous swarm consisting of aerial and wheeled robots. 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 system does not rely on any form of global information. We show how a spatially targeted one-to-one communication link can be established using a simple LED and camera based communication modality. We provide a probabilistic model of our approach to derive an upper bound on the average time required for establishing communication. In simulation, we show that our approach scales well. Furthermore, we show how our approach can be extended to establish a spatially targeted one-to-many communication link between an aerial robot and a specific number of co-located wheeled robots. 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.

Categories and Subject Descriptors

I.2.11 [Distributed Artificial Intelligence]: Multiagent systems

General Terms

Algorithms, Design, Experimentation

Keywords

Swarm robotics, heterogeneous swarms, communication protocol, situated communication

1. INTRODUCTION

In this paper, we consider a heterogeneous swarm system in which aerial robots supervise the activities of ground based wheeled robots. In such a system, it can be crucial for the aerial robots to be able to communicate with particular

Cite as: Establishing Spatially Targeted Communication in a Heterogeneous Robot Swarm, Mathews et al., *Proc. of 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, van der Hoek, Kaminka, Lespérance, Luck and Sen (eds.), May, 10–14, 2010, Toronto, Canada, pp. 939-946

Copyright © 2010, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

wheeled robots or groups of wheeled robots based on their location in the environment. We refer to this type of communication as *spatially targeted communication*. A typical message that an aerial robot might need to send using spatially targeted communication is “be careful, you are about to drive over an edge”.

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 [15]. 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 [12] or rely on some form of global knowledge [14, 16] that is not available on most swarm robotics systems.

In this study, our objective is to give aerial robots the capacity to establish a spatially targeted communication channel to one or more co-located wheeled robots without relying on any form of global knowledge. We demonstrate how an aerial robot can establish a spatially targeted communication link with a particular wheeled robot among a group of wheeled robots using situated communication based on on-board LEDs and camera. We use a binary selection process, whereby the aerial robot initially communicates with all wheeled robots 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.

2. RELATED WORK

In heterogeneous systems consisting of aerial robots and wheeled robots, wireless Ethernet has previously been used for communication [14, 16]. These studies compensate for the absence of inherent localization information in the wireless Ethernet medium by using global maps in conjunction




Figure 1: The two robot types of the heterogeneous swarm robotic platform considered in this study. (a) The prototype of the aerial robot. (b) A CAD model of the wheeled robot. The two robots are being developed at EPFL within the framework of the Swarmanoid project. More information about the project is available at <http://www.swarmanoid.org>.

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 [12] 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 [3], 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 [13], 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, 8]. 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. ROBOTIC PLATFORMS

We consider the heterogeneous swarm consisting of aerial and wheeled robots shown in Fig. 1. 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 ARGoS [11]. We also perform proof-of-concept experiments on a related robotic platform consisting of wheeled robots only, in which a predesignated wheeled robot assumes the role of the aerial robot. Below, we describe the heterogeneous swarm robotic platform currently under development, the related robotic platform, and how both platforms are related to each other.




Figure 2: The s-bot robotic platform. (a) The location of its camera, the hemi-spherical mirror and the LEDs. (b) A sample image returned by the camera showing the computed relative distances and angles (in white) to the neighboring s-bots.

3.1 The Heterogeneous Platform

The aerial robots can fly and are also equipped with a system of magnets that allows them to attach to a metal ceiling or to metal bars. In this study, the aerial robots are assumed to be attached to the ceiling (thus stationary at an elevated position). They are equipped with a pan-and-tilt camera pointing downward. The camera allows them to survey the ground and to detect the wheeled robots. Downward facing LEDs are used to communicate internal state information to the wheeled robots.

The wheeled robots are capable of moving and manipulating objects on the ground. They are equipped with infrared proximity sensors used for obstacle avoidance, colored LEDs to display internal states to the aerial robots and to neighboring wheeled robots, an omnidirectional camera to perceive other wheeled robots, and an upward pointing camera to perceive aerial robots. We assume that both aerial robots and the wheeled robots are able to display at least three distinct colors using their LEDs.

3.2 The S-bot Platform

We used a number of autonomous wheeled robots called *s-bots* [9] to conduct our real robot experiments. Each s-bot is equipped 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. 2a). 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. At each control step, the image returned by the omnidirectional camera is processed to detect color blobs. Since the height of the transparent tube and the optical properties of the hemi-spherical mirror are known, situational information such as the relative distance and the angle to each detected color blob (corresponding to a neighboring robot with its LEDs illuminated) can be estimated (see Fig. 2b).

The s-bots represent a suitable platform to prove our con-

cept on real robotic hardware as they possess an identical LEDs and camera-based situated communication modality as the robots in the heterogeneous swarm. This communication modality, when used together with the hemi-spherical mirror, enables both the s-bots and the wheeled robots in the heterogeneous swarm to communicate with neighboring robots. Furthermore, the wheeled robots in the heterogeneous swarm use another camera pointing upward at the ceiling to perceive the LEDs of the aerial robots. The aerial robot, in turn, uses a camera pointing downward to perceive colors displayed by the wheeled robots. In order to emulate the heterogeneous swarm using the s-bots, we let a pre-designated s-bot assume the role of the aerial robot.

4. ONE-TO-ONE COMMUNICATION

In this section, 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 signals¹ 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 aerial robot with a particular wheeled robot 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 $C := \{\text{red}, \text{blue}, \text{green}\}$ and $C_s := \{\text{blue}, \text{green}\}$.

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

¹In our system, each different LED color is considered a distinctive signal.

4.1 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. 4.3.2.

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 [7]. 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

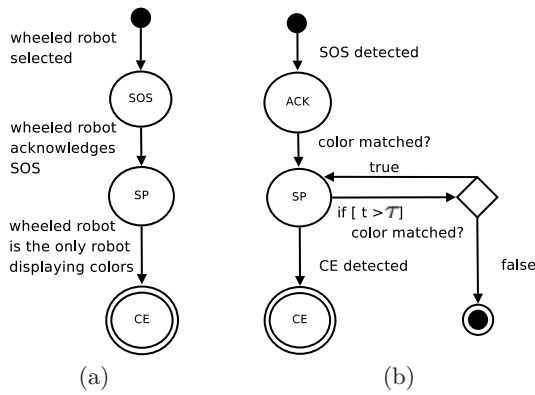


Figure 3: The finite state machines running on (a) the aerial robot and (b) the wheeled robots.

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 [5]. However, some general properties can be derived for its mean, $E[T_n]$. In particular, in [4, 6, 10] it is shown that, under some non-strict conditions and for $n > 3$, the following two properties hold:

$$E[T_n] \sim \frac{\ln n}{|\ln m|}, n \rightarrow \infty \quad (1)$$

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

In Sect. 4.3.2, we compare the upper bound on the mean predicted by Eq. 2 with the values obtained from simulation-based experiments.

4.2 One-to-one Controller

We developed two controllers: one for the aerial robots and one for the wheeled robots. The controllers are completely distributed and homogeneous, i.e., all wheeled robots execute the same controller. Both controllers are behavior-based and are represented as finite state machines (FSMs) in Fig. 3. In what follows, we explain both controllers assuming that both aerial and wheeled robots use three colors to communicate: red, blue and green.

Fig. 3a 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 aerial robot has determined that it needs to com-

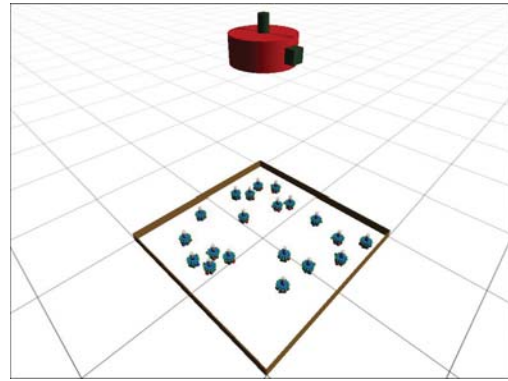


Figure 4: A screenshot from simulation including one aerial robot and twenty wheeled robots.

municate with a particular wheeled robot, it enters the **SOS** state. The transition from **SOS** to **SP** is triggered when the selected wheeled robot acknowledges the **SOS**. While in state **SP**, the aerial robot keeps matching the color displayed by the wheeled robot. If the selected wheeled robot is the only robot displaying any color, the aerial robot changes its state to **CE** to confirm the establishment of the communication.

Fig. 3b shows the FSM implemented on the wheeled robots; it consists of the three states **ACK** (acknowledge), **SP** (selection process), **CE** (communication established) and an end state which causes the wheeled robot 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 wheeled robot 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 aerial robot, the transition to state **SP** is triggered. When entering the state **SP**, each wheeled robot randomly selects and displays a color from the set of colors provided to the state. At the same time, each wheeled robot starts incrementing an internal timer t . Whenever this timer t exceeds a fixed threshold τ , the wheeled robot examines the color displayed on the aerial robot 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, process and react to the colors displayed by the wheeled robots. When a wheeled robot is in state **SP** and detects the **CE** color on the aerial robot, the wheeled robot can safely assume that it is the robot with which the aerial robot wishes to communicate. In this case, the wheeled robot confirms the termination of the selection process by displaying its **CE** color.

4.3 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 wheeled robots within the visual range of an aerial robot.

4.3.1 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)




Figure 5: Results of the scalability experiments with different number of colors for the selection process. Each data point is the average of 1000 repetitions.

at a height of 2 m. A number of wheeled robots are randomly placed within the visual range of this aerial robot (see Fig. 4). The aerial robot is able to perceive all wheeled robots within the arena and vice versa. Furthermore, each wheeled robot is able to perceive neighboring wheeled robots within a radius of 1 m.

Initially, the wheeled robots perform a random walk while avoiding other robots and the arena wall. Their green LEDs are on so that they are visible to the aerial robot. The aerial robot then randomly picks a particular wheeled robot and starts the one-to-one communication establishment process presented in Sect. 4.2. All wheeled robots respond to the aerial robot and remain stationary.

4.3.2 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 wheeled robots within the visual range of an aerial robot from 10 to 80 in steps of 5. We ran 1000 replications for each combination of number of wheeled robots and colors used. In Fig. 5, we plot the mean number of iterations spent on establishing communication between the aerial robot and a particular wheeled robot.

The results in Fig. 5 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 wheeled robots, the average number of iterations is 5.5 and the corresponding standard deviation (not shown in Fig. 5) is 1.9. On the other hand, in the case of 6 colors and 20 wheeled robots, 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. 5 also show that for all series of experiments, the number of iterations needed for an aerial robot to establish communication with a particular wheeled robot scales logarithmically with the number of wheeled robots. In Tab. 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 for the number of iterations spent by an aerial robot on establishing com-

Table 1: Scalability results for 80 wheeled robots (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

munication with a particular wheeled robot. In all cases, 80 wheeled robots 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 iterations required more predictable.

Note that we empirically validated only the upper bound predicted by the model (see Eq. 2) using the results obtained from simulation. We expect the logarithmic growth of the upper bound to behave similarly for larger groups of robots.

4.4 Real Robot Experiments

To confirm the real-world feasibility of our approach, we ran a series of experiments on the s-bot platform. Fig. 6 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 aerial robot. All robots are stationary and run the control program introduced in Sect. 4.2. The timer threshold τ is set to 20 control steps (equivalent to 2 seconds). Note that the optimal value of τ depends on the underlying hardware and that the value used here has not been fine-tuned. In the example in Fig. 6, 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 aerial robot’s role and another s-bot. In the sample experiment illustrated in Fig. 6, 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. 6 and other proof-of-concept experiments can be found at <http://iridia.ulb.ac.be/supp/IridiaSupp2009-006/>.

5. 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 aerial robot and a group co-located wheeled robots. We are not interested in which individual wheeled robots 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.

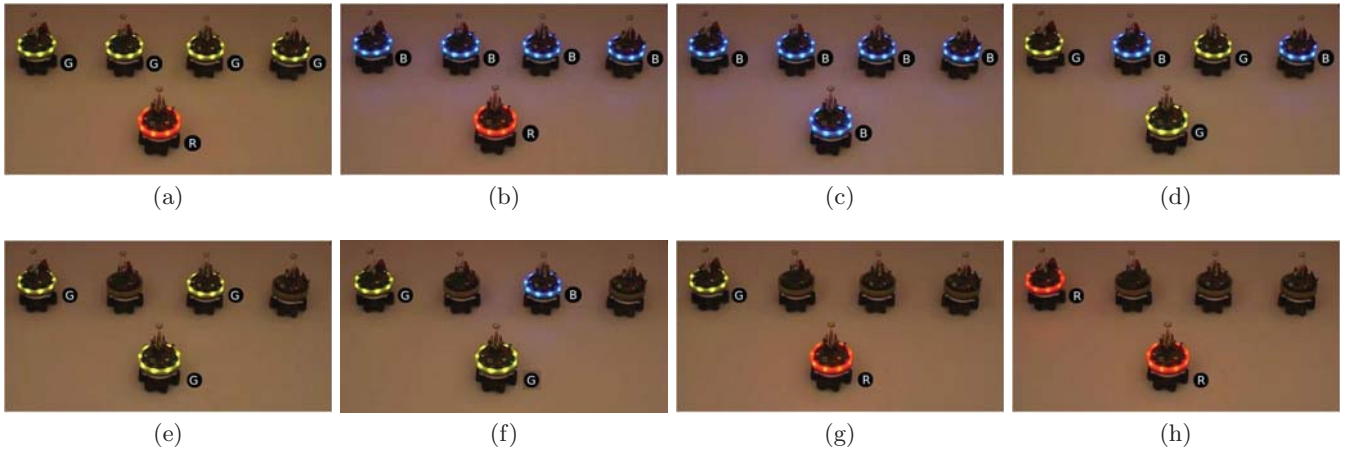


Figure 6: Snapshots of an experiment in which we let the lone s-bot in the bottom row assume the role of the aerial robot. This predesignated s-bot seeks to establish a one-to-one communication link with the s-bot on the top-left. The letters next to the s-bots represent the current color displayed: R=red, G=green and B=blue. (a) An SOS is sent by the predesignated s-bot. (b) The SOS is acknowledged by the other 4 s-bots. (c) The predesignated 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.

Our approach works by iteratively growing a group of wheeled robots 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 aerial robot may send a request to increase the size of the group. Only wheeled robots 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 aerial robot. The aerial robot 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 aerial robot can simply grant membership to all of the closest candidate robots. To achieve an exact group size, the aerial robot can request the closest candidate robots to relinquish their candidacies probabilistically. See Fig. 8 for an example of the algorithm running on real robots.

Below, we describe the robot controllers under the assumption that the aerial robots and the wheeled robots 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.

5.1 One-to-many Controller

We developed one controller for the wheeled robots and one controller for the aerial robots. Both controllers are behavior-based and are represented as FSMs in Fig. 7.

Fig. 7a 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 wheeled robots in red (i.e., in the group) is smaller than the desired group size. The transition ta2 is triggered when the number of wheeled robots 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. 7b shows the FSM implemented on the wheeled robots 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 wheeled robot in the visible range displays red, no other wheeled robot displaying blue is perceived closer to the group² and no wheeled robot in the visible range displays green. The transition tw2 is triggered if a candidate robots 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 given threshold τ , 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 aerial robot 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. 7a) allow the aerial robot to request the closest candidate robots to relinquish their

²Using its omnidirectional camera, a wheeled robot 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 wheeled robot assumes that it is the closest robot to the group.




Figure 7: The FSMs running (a) on an aerial robot and (b) the wheeled robots. 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.

candidacies. The aerial robot displays the color blue while in state **LEA**. If the aerial robot 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 aerial robot 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 wheeled robots with transition **tw4** as shown in Fig. 7b. If a wheeled robot is in state **CCN** and the color blue is displayed by the aerial robot, a timer t is incremented. Whenever this timer t exceeds a given threshold τ , transition **tw4** is triggered. The wheeled robot utilizes the outcome of a Bernoulli trial with probability $p_i = 0.5$ to decide whether to trigger transition **tw4** or to reset the timer t . These extended control programs allow the aerial robot to grow a group containing an exact number of wheeled robots by iterating between the states **STA**, **ADD** and **LEA** until the desired group size is reached.

5.2 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 aerial robot 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 wheeled robots at 80. We ran 1000 replications for each varied condition. The experiments were conducted in the experimental setup described in Sect. 4.3.1.

Tab. 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 2: Mean time (in number of iterations) to grow lower-bounded and exact group sizes (using 80 wheeled robots).

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

5.3 Real Robot Experiments

Fig. 8 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 predesignated s-bot, which we let assume the role of the aerial robot. All the other robots ran the control logic shown in Fig. 7b (but without transitions **ta3**, **ta4** and **tw5**). The timer threshold τ 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. 8 and other proof-of-concept experiments can be found at <http://iridia.ulb.ac.be/supp/IridiaSupp2009-006/>.

6. CONCLUSIONS

In this paper, we have demonstrated a novel approach to establish spatially targeted communication between aerial robots and one or more co-located wheeled robots. We showed that LEDs and cameras can be used as a situated communication modality to establish spatially targeted communication. We presented a probabilistic model that gives an upper bound on the average time required to establish a spatially targeted communication link between an aerial robot 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 aerial robot 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.

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 use the approach presented in this work in real-world scenarios in which the aerial robots and the wheeled robots collaborate in order to perform one or more tasks.

7. ACKNOWLEDGEMENTS

This research was carried out in the framework of Swarmanoid, a project funded by the Future and Emerging Technologies programme (IST-FET) of the European Commis-




Figure 8: Snapshots of an experiment in which we let the s-bot in the center assume the role of the aerial robot. 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.

sion under grant IST-022888. Marco Dorigo acknowledges support from the Belgian F.R.S.-FNRS, of which he is Research Director.

8. REFERENCES

- [1] K. B. Athreya and P. E. Ney. *Branching Processes*. Springer-Verlag, Berlin, Germany, 1972.
- [2] D. H. Barnhard, J. T. McClain, B. J. Wimpey, and W. 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] R. Grabowski and P. 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.
- [4] P. Haccou, P. Jagers, and V. Vatutin. *Branching Processes: Variation, Growth, and Extinction of Populations*. Cambridge University Press, Cambridge, UK, 2005.
- [5] T. E. Harris. *The Theory of Branching Processes*. Springer-Verlag, Berlin, Germany, 1963.
- [6] P. Jagers, F. 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.
- [7] D. Kendall. Branching processes since 1873. *Journal of London Mathematics Society*, 41:385–406, 1966.
- [8] 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.
- [9] 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.
- [10] A. 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.
- [11] C. Pinciroli. Object retrieval by a swarm of ground based robots driven by aerial robots. Mémoire de DEA, Université Libre de Bruxelles, Brussels, Belgium, 2007.
- [12] J. Pugh and A. 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.
- [13] F. Rivard, J. Bisson, F. Michaud, and D. 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.
- [14] A. T. Stentz, A. Kelly, H. Herman, P. Rander, O. Amidi, and R. Mandelbaum. Integrated air/ground vehicle system for semi-autonomous off-road navigation. In *Proceedings of AUVSI Unmanned Systems Symposium*, 2002.
- [15] K. 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.
- [16] R. T. Vaughan, G. S. Sukhatme, F. J. Mesa-Martinez, and J. 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.