Evolution of Direct Communication for a *Swarm-bot* Performing Hole Avoidance

Vito Trianni, Thomas H. Labella, and Marco Dorigo

IRIDIA - Université Libre de Bruxelles - Brussels, Belgium
{vtrianni,hlabella,mdorigo}@ulb.ac.be

Abstract. Communication is often required for coordination of collective behaviours. Social insects like ants, termites or bees make use of different forms of communication, which can be roughly classified in three classes: indirect (*stigmergic*) communication, direct interaction and direct communication. The use of stigmergic communication is predominant in social insects (e.g., the pheromone trails in ants), but also direct interactions (e.g., antennation in ants) and direct communication can be observed (e.g., the waggle dance of honey bee workers). Direct communication may be beneficial when a fast reaction is expected, as for instance, when a danger is detected and countermeasures must be taken. This is the case of hole avoidance, the task studied in this paper: a group of self-assembled robots – called *swarm-bot* – coordinately explores an arena containing holes, avoiding to fall into them. In particular, we study the use of direct communication in order to achieve a reaction to the detection of a hole faster than with the sole use of direct interactions through physical links. We rely on artificial evolution for the synthesis of neural network controllers, showing that evolving behaviours that make use of direct communication is more effective than exploiting direct interactions only.

Keywords: evolutionary robotics, swarm robotics, communication.

1 Introduction

In collective robotics research, the coordination of the activities in a group of robots requires the definition of communication strategies and protocols among the individuals. These strategies and protocols need not, however, be particularly complex. In many cases, simple forms of communication – or no explicit communication at all – are enough to obtain the coordination of the activities of the group [11]. This is the case of *swarm robotics*, that, drawing inspiration from social insects such as ants, termites or bees, focuses on distributed robotic systems characterised by limited communication abilities among robots.

Communication in social insects has been thoroughly studied, identifying different modalities used for the regulation of the colony's activities. The study of the nest building behaviour of termites of the genus *Macrotermes* led Grassé to the introduction of the concept of *stigmergy* [9]. Impressed by the complexity of termites' nests and by their dimension with respect to an individual, Grassé

M. Dorigo et al. (Eds.): ANTS 2004, LNCS 3172, pp. 130–141, 2004.

suggested that the cooperation among termites in their building activities was not the result of either some direct interactions among individuals, nor some other form of complex communication. On the contrary, cooperation could be explained as the result of environmental stimuli provided by the work already done – i.e., the nest itself. Another example of stigmergic communication has been observed in the foraging behaviour of many ant species, which lay a trail of pheromone, thus modifying the environment in a way that can inform other individuals of the colony about the path to follow to reach a profitable foraging area [8]. The concept of stigmergy describes an indirect communication among individuals, which is mediated by the environment [4].

Stigmergy is not the only way of communication that can be observed in social insects. *Direct interactions* – such as antennation, mandibular contact, trophallaxis – account for various social phenomena. For example, in many species of ants such as *Œcophilla longinoda*, recruitment of nest-mates for the exploitation of a food source is performed with a mix of antennation and trophallaxis: when an ant returning from a food source encounters another worker, it stimulates the other ant to follow the laid pheromone trail touching the nest-mate with the antennas and regurgitating a sample of the food source [10].

Some forms of *direct communication* within insect societies have been studied, a well-known example being the waggle dance of honey bees. A bee is able to indicate to the unemployed workers the direction and distance from the hive of a patch of flowers, using a "dance" that gives also information on the quality and the richness of the food source [16]. Direct communication in ants has been reported by Hölldobler and Wilson [10]: ants may use sound signals – called *stridulation* – for recruiting or for help requests. In presence of a big prey, ants of the genus *Aphaenogaster* recruit nest-mates using stridulation. Here, the sound signal does not attract ants, but it serves as a reinforcement of the usual chemical and tactile attractors, resulting in a faster response of the nest-mates.

The above examples suggest a possible taxonomy of different forms of communications in insect societies that can be borrowed for characterising a collective robotic system. Defining what communication is and classifying its different forms is not trivial, as confirmed by the number of different taxonomies that can be found in the literature [1, 3, 6, 12]. In [12], Matarić distinguishes between indirect or stigmergic, direct and directed communication, on the base of the communication modality (through the environment versus through a "speech act") and of the receiver (unknown versus specified). In [3], Cao et al. introduce three "communication structures" specific for a robotic system: interaction via environment, via sensing and via communication. Defining yet another taxonomy for different communication modalities is out of the scope of this paper. Thus, we borrow the taxonomy introduced in [3], adapting it to the natural examples introduced above. In doing this, we will use the above mentioned terminology, partly borrowed by [12]. Summarising, we will talk of:

Indirect or Stigmergic Communication. A form of communication that takes place through the environment, as a result of the actions performed by some individuals, which indirectly influence someone else's behaviour.

Direct Interaction. A form of communication that implies a non-mediated transmission of information, as a result of the actions performed by some individuals, which directly influence someone else's behaviour.

Direct Communication. A form of communication that implies a non-mediated transmission of information, without the need of any physical interaction.

As described above, all these forms of communication can be observed in biological systems, and in particular in social insects: research in swarm robotics focuses on the application of these simple forms of communication to artificial, autonomous systems. Referring to the above taxonomy, in this paper we will show how direct communication can be beneficial for reinforcing direct interactions. In our work, we study a swarm robotic system composed of a swarm of autonomous mobile robots, called *s-bots*, which have the ability to connect one to the other forming a physical structure – called *swarm-bot* – that can solve problems the single *s-bots* are not able to cope with[1] [5, 13]. The physical connections provide direct interactions among *s-bots* that can be exploited for coordination. Additionally, *s-bots* are provided with a sound signalling system, which can be used for direct communication. In this paper, we show that, using the sound signalling system, *s-bots* can reinforce the information passed through the physical connections, thus achieving a faster reaction.

The rest of this paper is organised as follows. Section 2 describes the problem we are interested in, that is, the hole avoidance task. Section 3 details the experimental setup used to perform the experiments. Finally, Section 4 is dedicated to the obtained results and Section 5 concludes the paper.

2 The Hole Avoidance Task

The hole avoidance task has been defined for studying collective navigation strategies. It can be considered an instance of a broader family of tasks, aimed at the study of all-terrain navigation. This family of tasks includes scenarios in which the robotic system has to face challenges such as avoiding holes or obstacles, passing through narrow passages or over a trough, climbing steep slopes and coping with rough terrain. With respect to these scenarios, the single robot approach may fail due to physical constraints or to limited individual capabilities. Our approach consists in relying on a swarm of robots, that can cooperate to overcome the individual limitations. Here, we address the all-terrain navigation problems making use of self-assembled structures – i.e., the *swarm-bots*.

The hole avoidance task represents a relatively simple problem compared to others in the all-terrain navigation family, but it is still very interesting for the study of collective navigation behaviours for a *swarm-bot*. The *s-bots* are placed in an arena presenting open borders and holes, in which the *swarm-bot* could

[1] This research is carried out within the SWARM-BOTS project, funded by the Future and Emerging Technologies Programme (IST-FET) of the European Community, under grant IST-2000-31010.

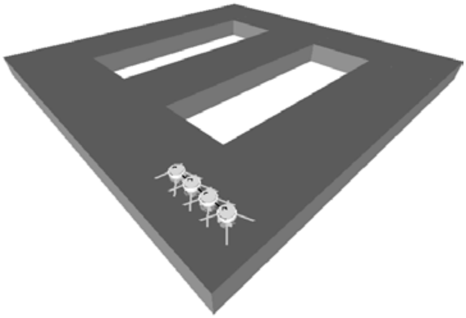

Fig. 1. The hole avoidance task. The picture shows the arena used, which presents open borders and contains two large rectangular holes. A *swarm-bot* formed by four linearly connected *s-bots* is shown.

fall (see Fig. 1). Four *s-bots* are rigidly connected in a linear formation. Their goal is to efficiently explore the arena, avoiding to fall into the holes or out of the borders of the arena.

The control of the *swarm-bot* is completely distributed, and *s-bots* can only rely on local information. The problem consists in how to coordinate the activity of the *s-bots*. In particular, the difficulty of the collective navigation is twofold: (i) coordinated motion must be performed in order to obtain a coherent navigation of the *swarm-bot* as a whole, as a result of the motion of its components; (ii) holes are not perceived by all the *s-bots* at the same time. Thus, the presence of an hazard, once detected, must be communicated to the entire group, in order to trigger a change in the direction of motion.

The complexity of the task justifies the use of evolutionary robotics techniques for the synthesis of the *s-bots*' controller [15]. In a previous work, we studied the hole avoidance problem evolving simple neural controllers that were able to perform coordinated motion and hole avoidance, relying only on direct interactions among *s-bots* [17]. In this paper, we focus on the use of direct communication, in order to reinforce the direct interactions and therefore to obtain more efficient behaviours. In fact, direct communication among *s-bots* speeds up the reaction to the detection of a hole, thus it is beneficial for the efficiency of the navigation.

3 Experimental Setup

As already mentioned, we studied hole avoidance in a previous work [17], obtaining interesting results. In this paper, we aim at improving the obtained results modifying the experimental setup as follows: (i) the simulation model of the *s-bot* is modified, as described in Sec. 3.1; (ii) the controllers include the possibility to actuate the speaker, thus enabling direct communication among *s-bots* (see also Sec. 3.2); (iii) the fitness computation is simplified, taking into account only

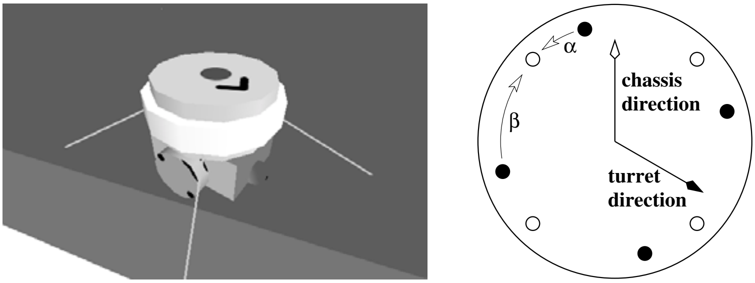

Fig. 2. Left: The simulated model of an *s-bot*. The light rays represent the position of the ground sensors, mounted on the rotating turret. The dark circle painted on the turret indicates that the *s-bot* is emitting a sound signal. Right: Description of the encoding of the sensors integral with the turret to the virtual sensors integral with the chassis. The filled circles indicate the position of the real sensors, while the empty circles refer to the position of the virtual sensor with respect to the direction of the chassis.

variables directly available to each *s-bot*, such as sensor readings or internal state variables (see Sec. 3.3).

In order to test the effectiveness of the use of direct communication among *s-bots*, we performed two sets of experiments: in the first setting only direct interactions were used, while in the second direct communication capabilities were added.

3.1 The Simulation Model

We developed a simulation software based on Vortex$^{\text{TM}}$, a 3D rigid body dynamics simulator that provides primitives for the implementation of detailed and realistic physics-based simulations (see [13] for more details about the simulator). We have defined a simple *s-bot* model that at the same time allows fast simulations and preserves those features of the real *s-bot* that were important for the experiments (see Figure 2 left). This model matches more closely the real *s-bot* than the one used in the previous work [17], both in the geometries and in the sensing abilities.

The *s-bot* has a differential drive motion provided by a traction system composed of four wheels: two lateral, motorised wheels and two spherical, passive wheels placed in the front and in the back, which serve as support. The four wheels are fixed to the chassis, which also holds the cylindrical rotating turret. The turret can rotate around its axis and it holds many sensory systems. Connections among *s-bots* can be made using a virtual gripper, which is modelled by dynamically creating a joint between two *s-bots*. The position of the virtual gripper is represented by an arrow painted on the turret. Finally, the turret also carries a loudspeaker that can be controlled to produce a tone that can be perceived by the other *s-bots*.

Each *s-bot* is provided with a *traction sensor*, which detects the forces that are applied to the junction between the chassis and the rotating turret. Four variables encode the traction force information from four different preferential orientations with respect to the chassis (front, right, back and left, see [2] for more details). Traction sensors are responsible for the detection of the direct interactions among *s-bots*. In fact, an *s-bot* can generate a traction force that is felt by the other *s-bots* connected through their grippers. This force mediates the communication among *s-bots*, and it can be exploited for coordinating the activities of the group: it proved to be important to evolve coordinated motion strategies in a *swarm-bot* and for collective obstacle and hole avoidance [2, 17].

The presence of holes is detected using four *ground sensors* – infrared proximity sensors pointing to the ground – that are integral with the rotating turret. In order to account for the rotation of the turret, we encode the information coming from the ground sensors in four virtual sensors integral with the chassis. As pictured in the right part of Fig. 2, the value taken by the virtual sensors is computed as the weighted average of the two closest ground sensors. In particular, if α and β are the angular differences from the two closest ground sensors, then $\cos^2(\alpha)$ and $cos^2(\beta)$ are the weights for the average. Noise is simulated for all sensors, adding a random value uniformly distributed within the 5% of the sensor saturation value.

Each *s-bot* is also equipped with a loudspeaker and three directional microphones, used to detect the tone emitted by other *s-bots*. Also directional microphones, being integral with the turret, are encoded in three virtual sound sensors integral with the chassis following a procedure similar to the one used for ground sensors. The loudspeaker can be switched on, simulating the emission of a continuous tone, or it can be turned off. Exploiting this equipment, *s-bots* have direct communication capabilities.

S-bots can control the two wheels, independently setting their speed in the range $[-6.5, 6.5]$ rad/s. The virtual gripper is used to connect to another *s-bot*. However, in this work, the *s-bots* stay always assembled in a *swarm-bot* formation, thus connection and disconnection procedures have not been simulated. Finally, the motor controlling the rotation of the turret is actuated setting its desired angular speed proportionally to the difference between the desired angular speed of the left and right wheels. This setting helps the rotation of the chassis with respect to the turret also when one or both wheels of the *s-bot* do not touch the ground [2].

The *swarm-bot* is composed of four *s-bots* rigidly connected to form a chain. It is placed in a square arena of 4 meters side, that presents open borders and two rectangular holes (80×240 cm, see Fig. 1). The dimensions have been chosen to create passages that can be navigated by the *swarm-bot*, no matter its orientation.

3.2 The Controller and the Genetic Algorithm

The *s-bots* are controlled by artificial neural networks, whose parameters are set by an evolutionary algorithm. A single genotype is used to create a group of *s-bots* with an identical control structure – a homogeneous group. Each *s-bot* is controlled by a fully connected, single layer feed-forward neural network – a perceptron. Each input is associated with a single sensor, receiving a real value in the range $[0.0, 1.0]$, which is a simple linear scaling of the reading taken from its associated sensor. Additionally, the network is provided with a bias unit – an input unit whose activation state is always 1.0 – and two output neurons that control the motors of the *s-bot*.

As mentioned above, we performed two sets of experiments, which differ in the form of communication used. Thus, the neural networks controlling the *s-bots* change depending on which sensors and actuators are employed. In all the experiments, traction and ground sensors have been used. Specifically, 4 inputs of the perceptron are dedicated to the traction sensors and 4 other inputs are dedicated to the virtual ground sensors (see Sec. 3.1). If direct communication is used, three more sensors are used, corresponding to the three microphones with which an *s-bot* is endowed. These sensors are connected to three additional neural inputs. Concerning the actuators, the two outputs of the perceptron are used to control the left and the right wheel. Additionally, the same two outputs control the turret-chassis motor, as described in Sec. 3.1. When direct communication is used, the activation of the loudspeaker has been handcrafted, simulating a sort of reflex action: an *s-bot* activates the loudspeaker whenever one of its ground sensors detects the presence of a hole. Thus, the neural network does not control the emission of a sound signal. However, it receives the information coming from the three directional microphones, and evolution is responsible for shaping the correct reaction to the perceived signals.

The weights of the perceptron's connections are genetically encoded parameters. A simple generational genetic algorithm (GA) is used [7]. Initially, a random population of 100 genotypes is generated. Each genotype is a vector of binary values – 8 bits for each parameters. The genotype is composed of 144 bits in the first setting (using direct interactions only) and 192 bits in the second setting (using direct communication). Subsequent generations are produced by a combination of selection with elitism and mutation. Recombination is not used. At every generation, the best 20 genotypes are selected for reproduction, and each generates 4 offspring. The genotype of the selected parents is copied in the subsequent generation; the genotype of their 4 offspring is mutated with a 3% probability of flipping each bit. One evolutionary run lasts 150 generations.

3.3 The Fitness Computation

During the evolution, a genotype is mapped into a control structure that is cloned and downloaded in all the *s-bots* taking part in the experiment (i.e., we make use of a homogeneous group of *s-bots*). Each genotype is evaluated 5 times – i.e., 5 trials. Each trial differs from the others in the initialisation of the random number generator, which influences both the initial position of the *swarm-bot* within the arena and the initial orientation of each *s-bot*'s chassis. Each trial lasts $T = 200$ simulation cycles, which correspond to 20 seconds of real time.

The behaviour produced by the evolved controller is evaluated according to a fitness function that takes into account only variables accessible to the *s-bots* (see [15], page 73). In each simulation cycle t, for each *s-bot* s belonging to the *swarm-bot* S, the individual fitness $f_s(t)$ is computed as the product of three components:

$$f_s(t) = \omega_s(t) \cdot \Delta\omega_s(t) \cdot \gamma_s(t), \tag{1}$$

where:

- $\omega_s(t)$ accounts for fast motion of an *s-bot*. It is computed as the sum of the absolute values of the angular speed of the right and left wheels, linearly scaled in the interval $[0, 1]$:

$$\omega_s(t) = \frac{|\omega_{s,l}(t)| + |\omega_{s,r}(t)|}{2 \cdot \omega_m}, \tag{2}$$

  where $\omega_{s,l}(t)$ and $\omega_{s,r}(t)$ are respectively the angular speed of the left and right wheel of *s-bot* s at cycle t, and ω_m is the maximum angular speed achievable.
- $\Delta\omega_s(t)$ accounts for the straightness of the motion of the *s-bot*. It is computed as the difference between the angular speed of the wheels, as follows:

$$\Delta\omega_s(t) = \begin{cases} 0 & \text{if } \omega_{s,l}(t) \cdot \omega_{s,r}(t) < 0 \\ 1 - \sqrt{\frac{|\omega_{s,l}(t) - \omega_{s,r}(t)|}{\omega_m}} & \text{otherwise} \end{cases}, \tag{3}$$

  where the difference is computed only if the wheels rotate in the same direction, in order to penalise more any turning-on-the-spot behaviour. The square root is useful to emphasise small speed differences.
- $\gamma_s(t)$ accounts for coordinated motion and hole avoidance. It is computed as follows:

$$\gamma_s(t) = 1 - \max\big(I_s(t), G_s(t)\big), \tag{4}$$

  where $I_s(t)$ is the intensity of the traction force perceived by the *s-bot* s at time t, $G_s(t)$ is the maximum activation among the ground sensors. Both can take values in the interval $[0, 1]$. This component favours coordinated motion as it is maximised when the perceived traction is minimised, which corresponds to a coherent motion of the *swarm-bot*. It also favours hole avoidance because it is maximised if the *s-bots* stay away from the holes.

Given the individual fitness $f_s(t)$, the fitness F_θ of a trial θ is computed as follows:

$$F_\theta = \frac{1}{T} \sum_{t=1}^{T_f} \min_{s \in S} f_s(t), \tag{5}$$

where T is the maximum number of cycles and $T_f \leq T$ is the cycle at which the simulation ended, which may be smaller than the maximum allowed if the *swarm-bot* happens to fall into a hole. Averaging the individual components on T rather than on T_f simulation cycles puts an additional selective pressure for the evolution of hole avoidance. Additionally, at each simulation cycle t we select the minimum among the individual fitnesses $f_s(t)$, which refers to the worst-behaving *s-bot*, therefore obtaining a robust overall fitness computation.

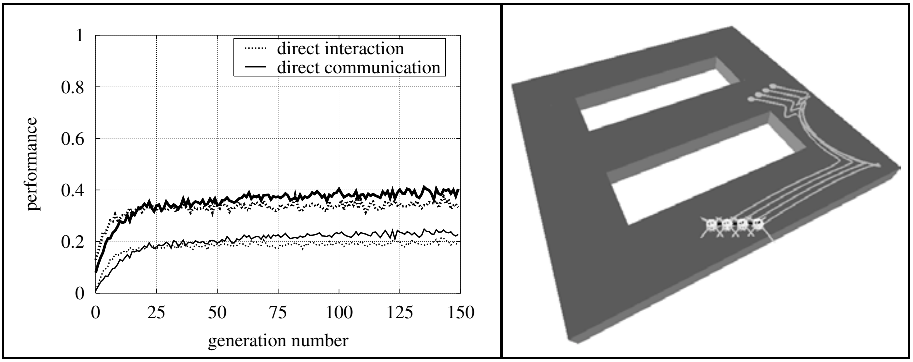

Fig. 3. Left: The average performance of the 10 replications is plotted against the generation number for each experimental setting. Thick lines refer to the best individual of the population, while thin lines refer to the population average. Right: trajectories of the *s-bots* while performing hole avoidance. Movies of this behaviour are available at `http://www.swarm-bots.org/hole-avoidance.html`.

4 Results

For both settings – using only direct interactions (hereafter indicated as DI) and complementing them with direct communication (hereafter indicated as DC) – the evolutionary experiments were replicated 10 times. The average fitness values, computed over all the replications, are shown in Figure 3 left. The average performance of the best individual and of the population are plotted against the generation number. All evolutionary runs were successful. The average fitness value of the best individuals reaches 0.4, where a value of 1 should be understood as a loose upper-bound to the maximum value the fitness can achieve[2]. It is worth noting that the average fitness of DC is slightly higher than in the case of DI. This suggests that the use of direct communication among *s-bots* is beneficial for the hole avoidance task.

A qualitative analysis of the behaviours produced by the two settings reveals no particular differences in the initial coordination phase that leads to a coherent motion of the *swarm-bot* (see Fig. 3 right). In both cases, the *s-bots* start to move in the direction they were positioned, resulting in a rather disordered overall motion. Within a few simulation cycles, the physical connections transform this disordered motion into traction forces, that are exploited to coordinate the group. When an *s-bot* feels a traction force, it rotates its chassis in order to cancel this force. Once the chassis of all the *s-bots* are oriented in the same direction, the traction forces disappear and the coordinated motion of the *swarm-bot* starts (see also [17, 2]).

[2] This maximum value could be achieved only by a *swarm-bot* coordinately moving in a flat environment, without holes. In the arena shown in Fig. 3 right, the narrow passages result in frequent activations of the ground sensors, and therefore in frequent re-organisations of the *swarm-bot*.

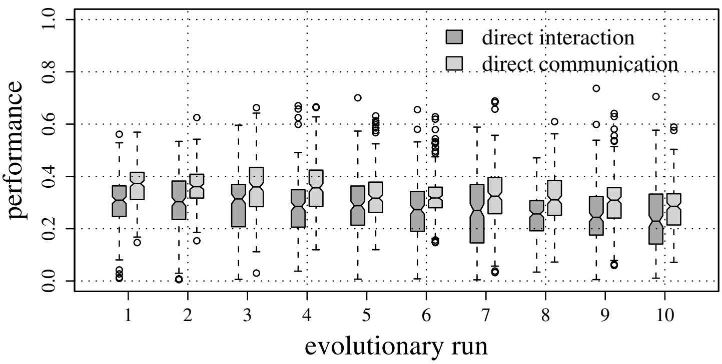

Fig. 4. Post-evaluation analysis performed evaluating 200 times the best individuals obtained from each replication of the experiment. Boxes represent the inter-quartile range of the data, while the horizontal bars inside the boxes mark the median values. The whiskers extend to the most extreme data points within 1.5 of the inter-quartile range from the box. The empty circles mark the outliers.

The differences between the two settings *DC* and *DI* are evident once the hole avoidance behaviour is considered. When an *s-bot* detects an edge, it rotates the chassis and changes the direction of motion in order to avoid falling. When using only direct interactions, this change in direction produces a traction force for the other *s-bots*, which triggers a new coordination phase that ends up in a new direction of motion that leads the *swarm-bot* away from the edge. This simple behaviour exploits the direct interactions among *s-bots* – shaped as traction forces – to communicate the presence of an hazard – the hole to be avoided. However, this strategy may fail as communication via traction is sometimes too weak to be perceived by the whole *swarm-bot*. On the contrary, the evolved controllers that makes use of direct communication react faster to the detection of a hole: the *s-bot* that detects the hole emits a sound signal that is immediately perceived by the rest of the group. Thus, the whole *swarm-bot* starts turning away from the hole, without waiting to perceive a strong traction force. Traction is then exploited again in order to perform coordinated motion.

From the qualitative analysis, the use of direct communication seems to confirm our expectations: direct communication provides a faster reaction to the detection of a hole and therefore a more efficient avoidance behaviour. In order to quantitatively assess the difference in performance between *DC* and *DI*, we performed a post-evaluation analysis and compared the results obtained with the two settings. For each evolutionary run, we selected the 20 best individuals of the final population and we re-evaluated them in 200 trials, each characterised by a different random initialisation. All individuals were tested using the same set of trials. The performance of the re-evaluations was measured using Eq. (5). We selected the individual with best mean performance in the post-evaluations and discarded the other nineteen individuals. A box-plot summarising the performance of these individuals is shown in Fig. 4. It is possible to notice that *DC* generally performs better than *DI*.

On the base of these data, we performed a two-way analysis of variance to test if there is a significant difference in performance between the settings [14]. The analysis considers 2 factors (the settings), 200 blocks (the testing trials) and 10 replications for each combination of factor/block (the evolutionary runs). The applicability of the method was checked looking at the residuals coming from the linear regression modelling of the data: no violation of the hypothesis to use the analysis of variance was found. The result of the analysis allows us to reject the null hypothesis that there is no difference among the two settings (p-value < 0.0001). On the base of the mean performance of the two settings – 0.3316 for DC and 0.2708 for DI – we can conclude that, in the experimental conditions considered, a system that uses direct communication among the *s-bots* performs better than one that exploits only direct interactions.

5 Conclusions

In this paper, we have shown how the use of direct communication in a *swarm-bot* performing hole avoidance can be beneficial for the effectiveness of the group behaviour. Comparing the use of direct communication with the case in which only direct interactions among *s-bots* were possible, we found that the former setting performs statistically better than the latter. It is worth noting that direct communication acts here as a reinforcement of the direct interaction among *s-bots*. In fact, *s-bots* react faster to the detection of the hole when they receive a sound signal, without waiting to perceive a traction strong enough to trigger the hole avoidance behaviour. However, traction is still necessary for avoiding the hole and coordinating the motion of the *swarm-bot* as a whole. Finally, it is important to remark that all controllers were synthesised by artificial evolution, which proved to be an efficient mean for automatically developing behaviours for homogeneous groups of robots.

Acknowledgements

This work was supported by the “SWARM-BOTS” project and by the “ECagents” project, two projects funded by the Future and Emerging Technologies programme (IST-FET) of the European Commission, under grant IST-2000-31010 and 001940 respectively. The information provided is the sole responsibility of the authors and does not reflect the Community’s opinion. The Community is not responsible for any use that might be made of data appearing in this publication. Marco Dorigo acknowledges support from the Belgian FNRS, of which he is a Senior Research Associate, through the grant “Virtual Swarm-bots”, contract no. 9.4515.03, and from the “ANTS” project, an “Action de Recherche Concertée” funded by the Scientific Research Directorate of the French Community of Belgium.

References

1. T. Balch and R. C. Arkin. Communication in reactive multiagent robotic systems. *Autonomous Robots*, 1(1):27–52, 1995.
2. G. Baldassarre, S. Nolfi, and D. Parisi. Evolution of collective behavior in a team of physically linked robots. In R. Gunther, A. Guillot, and J.-A. Meyer, editors, *Applications of Evolutionary Computing - Proceedings of the Second European Workshop on Evolutionary Robotics (EvoWorkshops2003: EvoROB)*, pages 581–592. Springer-Verlag, Berlin, Germany, 2003.
3. Y. U. Cao, A. S. Fukunaga, and A. B. Kahng. Cooperative mobile robotics: Antecedents and directions. *Autonomous Robots*, 4:1–23, 1997.
4. M. Dorigo, E. Bonabeau, and G. Theraulaz. Ant algorithms and stigmergy. *Future Generation Computer Systems*, 16(8):851–871, 2000.
5. M. Dorigo, V. Trianni, E. Şahin, R. Groß, T. H. Labella, G. Baldassarre, S. Nolfi, J.-L. Deneubourg, F. Mondada, D. Floreano, and L. M. Gambardella. Evolving self-organizing behaviors for a swarm-bot. *Autonomous Robots*, 17(2–3):223–245, 2004.
6. G. Dudek, M. Jenkin, and E. Milios. A taxonomy of multirobot systems. In T. Balch and L. E. Parker, editors, *Robot Teams: From Diversity to Polymorphism*. A K Peters Ltd., Wellesley, MA, 2002.
7. D. E. Goldberg. *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley, Reading, MA, 1989.
8. S. Goss, S. Aron, J.-L. Deneubourg, and J. M. Pasteels. Self-organized shortcuts in the argentine ant. *Naturwissenchaften*, 76:579–581, 1989.
9. P. P. Grassé. La reconstruction du nid et les coordinations interindividuelles chez *Bellicositermes natalensis et Cubitermes sp.* La théorie de la stigmergie: Essai d'interprétation du comportement des termites constructeurs. *Insectes Sociaux*, 6:41–81, 1959.
10. B. Hölldobler and E. O. Wilson. *Journey to the Ants: A Story of Scientific Exploration*. Bellknap Press/Harvard University Press, Cambridge, MA, 1994.
11. C. R. Kube and H. Zhang. Task modelling in collective robotics. *Autonomous Robots*, 4:53–72, 1997.
12. M. J. Matarić. Using communication to reduce locality in distributed multiagent learning. *Journal of Experimental and Theoretical Artificial Intelligence, Special Issue on Learning in DAI Systems*, 10(3):357–369, 1998.
13. F. Mondada, G. C. Pettinaro, A. Guignard, I. V. Kwee, D. Floreano, J.-L. Deneubourg, S. Nolfi, L. M. Gambardella, and M. Dorigo. SWARM-BOT: A new distributed robotic concept. *Autonomous Robots*, 17(2–3):193–221, 2004.
14. D. C. Montgomery. *Design and Analysis of Experiments*. John Wiley & Sons, Inc., New York, NY, 5th edition, 1997.
15. S. Nolfi and D. Floreano. *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*. MIT Press/Bradford Books, Cambridge, MA, 2000.
16. T. Seeley. *The Wisdom of the Hive*. Harvard University Press, Cambridge, MA, 1995.
17. V. Trianni, S. Nolfi, and M. Dorigo. Hole avoidance: Experiments in coordinated motion on rough terrain. In F. Groen, N. Amato, A. Bonarini, E. Yoshida, and B. Kröse, editors, *Intelligent Autonomous Systems 8*, pages 29–36. IOS Press, Amsterdam, The Netherlands, 2004.