

Université Libre de Bruxelles

*Institut de Recherches Interdisciplinaires
et de Développements en Intelligence Artificielle*

Emergent Collective Decisions in a Swarm of Robots

Vito Trianni and Marco Dorigo

Technical Report No.

TR/IRIDIA/2005-21

June 2005

in Proceedings of the 2005 IEEE Swarm Intelligence Symposium (SIS 2005)

Emergent Collective Decisions in a Swarm of Robots

Vito Trianni and Marco Dorigo

IRIDIA - Université Libre de Bruxelles, Bruxelles, Belgium

Abstract

A swarm robotic system is normally characterised by many individuals, each having a partial/limited knowledge about the global pattern of which it constitutes an element. In such a system, decision-making processes may be problematic. However, inspiration can be drawn from insect societies, in which self-organisation plays a crucial role in most of the decisions taken by the colony. In this work, we show how, in a swarm robotic system, a decision can be the result of a collective process: it emerges from the numerous interactions among the individuals and between individuals and environment. We present a task in which a swarm of physically connected, simulated robots has to take a decision whether to pass over a trough or change direction of motion if the gap is too wide to be bridged. We show how such a decision can be collectively taken, based only on a self-organising process.

1 Introduction

Decision-making mechanisms are important features for an intelligent agent, as they make it possible to display different behaviours as a function of the particular environmental situation the agent perceives and in relation to its beliefs and its desires. Individually, a decision is often the result of a process that takes into account information gathered from the environment. For example, animals collect information about the quality of a food source while foraging. Depending on this information, they take the decision to stay in the same area or to search for a more profitable one. Similar behaviours have been studied in robotics, in which a robot takes a decision that depends on its past experience. In these works, decision-making mechanisms are based on short/long term memory or on the integration over time of the agent's perceptions [13, 7, 12].

A more complex case is presented by decisions that have to be taken at a collective level. Societies may entrust their decision-making ability to a few leaders that care about the whole community. This is the case of groups of mammals, characterised by the presence of a few individuals that lead the activities of the others. The situation is different in insect societies, in which decisions are taken collectively. Many examples of collective choice have been studied so far in social insects. These decisions are generally the result of a self-organising process: the decision *emerges* from the numerous interactions among the individuals forming the colony, and from the interactions between individuals and the environment [4]. Therefore, complex decision-making processes can be observed at the collective level, notwithstanding the simple behavioural rules followed by each individual insect. For example, bees are able to collectively select the most profitable foraging site between two different food sources [9], while ants collectively select the shortest path from the nest to a food source, thereby optimising the foraging process [3, 2].

The above examples show how a decision at the collective level can emerge from rather simple behaviours at the individual level. A similar approach characterises

research in *swarm robotics*, which is a novel approach to the design and implementation of robotic systems composed of *swarms* of robots tightly interacting and cooperating to reach their goal. The aim of this paper is to demonstrate how, in a swarm robotic system, complex decision making abilities can emerge from simple individual behaviours.

We study a swarm robotic system composed of a swarm of autonomous mobile robots—called *s-bots*—that have the ability to connect one to the other forming a physical structure—called *swarm-bot*. Therefore, one of the main features of a *swarm-bot* is the ability to solve problems a single individual cannot cope with, due to its limited abilities. In this paper, we study one example of such problems, that is, how to pass over a trough that would block the navigation of a single *s-bot*. In this case, physical connections among *s-bots* serve as support for those that are suspended over the gap, so that the *swarm-bot* as a whole can continue moving.

In a previous work, we studied collective navigation strategies for a *swarm-bot* that had to move in an environment presenting holes and obstacles [10, 11]. The problem consisted in how to coordinate the activity of the *s-bots* for achieving both a coherent navigation of the *swarm-bot* as a whole and an efficient collective avoidance of the hazards, which could be perceived only by a few individuals within the group. In this paper, we show how the same controllers developed for hole avoidance can be used to pass over a trough, whenever it can be bridged by a *swarm-bot*. In particular, we show how a decision is taken collectively by the *swarm-bot*, whether to pass over the gap or change direction of motion and avoid falling. We show how this collective decision emerges purely from the interactions between the *s-bots* and the environment, which enable the *swarm-bot* to roughly estimate the width of the trough to be passed.

In the following, we shortly describe the *s-bot* and its simulated model. Then, we sketch the results obtained for the hole avoidance task. Afterwards, we present the experimental setup used to study the ability of passing over a trough. Finally, we report on the obtained results.

2 The *s-bot*

An *s-bot* is a small autonomous mobile robot, shown in the left part of Figure 1. The *s-bot* has a traction system composed of tracks and wheels. Above the traction system, a rotating turret holds many sensory systems and the gripper for making connections with other *s-bots*, as shown in the right part of Figure 1 (for more details, see [6]). In this paper, however, experiments are performed in simulation, using a software based on VortexTM, a 3D rigid body dynamics 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).

The simulated *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, powered by a motorised joint. Connections among *s-bots* can be made using a gripper, which is simulated by dynamically creating a joint between two *s-bots*. The position of the gripper is represented by an arrow painted on the turret.

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 [1] for more details). The




Figure 1: The *s-bot*. Left: one *s-bot*. Right: two *s-bots* connected through their gripper, therefore forming a small *swarm-bot*.

traction sensors are responsible for the detection of the interactions among *s-bots*, and they also mediate the communication about the presence of a hazard. In fact, when holes or obstacles are detected, an *s-bot* can communicate the hazard simply trying to move away from it, therefore generating a traction force that is felt by the other *s-bots*. This force can be exploited for coordinating the activities of the group: it proved to be important in order to evolve collective obstacle and hole avoidance strategies for a *swarm-bot* [10, 11].

The presence of holes/troughs is detected using four *ground sensors*—infrared proximity sensors pointing to the ground—that are integral with the rotating turret. The proximity sensors are evenly distributed around the *s-bot*'s turret, and they are inclined of 30 degrees with respect to the horizontal plane. 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. The value taken by the virtual sensors is computed as the weighted average of the two closest ground sensors (see [10] for more details). Noise is simulated for all sensors, adding a random value uniformly distributed within the 5% of the sensor saturation value.

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,




Figure 2: The simulation model. Left: a simulated *s-bot*, where many details not relevant for our experiments have been omitted. Right: a *swarm-bot* formed by 9 *s-bots* in a square formation.

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 when *s-bots* are connected in a *swarm-bot* formation [1].

3 Hole avoidance

As mentioned above, the work presented in this paper is based on a previous study, where strategies for hole/obstacle avoidance have been evolved for a *swarm-bot* [10]. In this section, we briefly explain the methodology we used and the results obtained, as they are necessary to introduce the results presented in Section 4.

The hole avoidance task is a simple but challenging navigation problem, in which *s-bots* have to explore an arena presenting holes in which they risk to fall (see Figure 3, top). In this situation, a *swarm-bot* is more efficient than a single *s-bot* as it can rely on the cooperation among *s-bots*, and on the physical connections among its components.

3.1 Experimental setup

The *s-bots* are controlled by artificial neural networks, whose parameters are set by an evolutionary algorithm. We make use of a simple evolutionary algorithm that works on a population of 100 binary encoded genotypes. At each generation, the 20 best individuals are selected for reproduction and generate each 5 offspring. Four of them are mutated with a 3% probability of flipping each bit. During the evolution, a genotype is mapped into a neural control structure that is cloned and downloaded to 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. In order to evolve the neural controllers, four *s-bots* are connected in a linear *swarm-bot* formation, and they are placed in the arena shown in Figure 3. The initial orientation of the chassis of each *s-bot* is randomly chosen. The behaviour produced by the evolved controller is evaluated according to a fitness function that takes into account only variables directly accessible to the *s-bots*. The fitness function rewards straight and fast motion of the *s-bots*, and penalises those groups of *s-bots* that do not coordinate their movements or that spend too much time in the vicinity of a hole. This last component is computed simply looking at the activation of the traction and the ground sensors. Additionally, if the behaviour results in a fall of the *swarm-bot* into a hole, the corresponding genotype is penalised (for more details, see [10]).

3.2 Results

Using the above methodology, we performed 10 different evolutionary runs, each starting with a different randomly initialised population. The average fitness values, computed over all the replications, are shown in the bottom part of Figure 3. 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 approaches 0.4, where a value of 1 should be understood as a loose upper-bound to the maximum value the fitness can achieve.¹

¹This maximum value could be achieved only by a *swarm-bot* coordinately moving in a flat environment, without holes. In the arena shown in the top part of Figure 3, the narrow passages result in frequent activations of the ground sensors, and therefore in frequent re-organisations of the *swarm-bot*.

The behaviours produced by the evolved neural network are characterised by an initial coordination phase that leads to a coherent motion of the *swarm-bot*. The *s-bots* start moving in the direction they were initialised, 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 a same direction, the traction forces disappear and the coordinated motion of the *swarm-bot* starts (see also [1]). When an *s-bot* detects an edge, it rotates the chassis and changes the direction of motion in order to avoid falling. This change in direction produces




Figure 3: The hole avoidance task. Top: The picture shows the arena used, which presents open borders and contains two large rectangular holes. A *swarm-bot* composed of four linearly connected *s-bots* is shown, along with their trajectories while avoiding holes. Bottom: The average fitness of the best individual in the 10 replications of the experiments and the average fitness of the population are plotted against the generation number.

a traction force for the other *s-bots*, which triggers a new coordination phase. The *s-bots* eventually choose a new direction of motion that leads the *swarm-bot* away from the edge. In some cases, the reaction of a single *s-bot* may not be sufficient to influence the behaviour of the rest of the group. As a consequence, the *s-bot* may be pushed out of the arena. However, physical connections serve as support for this *s-bot*, while the rest of the group continues to perform hole avoidance and eventually leads the whole *s-bot* to a safer location.

These results are mainly based on the properties of the traction sensor, which proved to be a powerful mechanism for achieving coordination in the *swarm-bot*. In fact, it allows the *swarm-bot* to exploit the complex dynamics arising from interactions among individual *s-bots* and between the *s-bots* and the environment. It also provides robustness and adaptivity features with respect to environmental or structural changes of the *swarm-bot*. Additionally, it provides a way to exploit the direct interactions among *s-bots*—shaped as traction forces—to communicate the presence of a hazard—the hole to be avoided [10]. Finally, traction forces are also at the base of the self-organising process that leads to the collective decision about passing over a trough or avoiding it when it is too wide. In the following section, we will detail this process.

4 Passing over a trough

The controller described above bases its functioning on the perception of holes through the ground sensors, and on the traction forces applied by one *s-bot* to the others. Intuitively, if the perception of holes is masked to the *s-bots*—for example, setting to 0 the activation of the ground sensors—then the *swarm-bot* will sooner or later fall into one of them. However, whenever the hole is small enough to be bridged, one could observe the *swarm-bot* passing on the other side and continuing its exploration of the arena. Therefore, if the *swarm-bot* were able of estimating the size of the hole, it could decide whether to change direction of motion and avoid falling, or to try to pass on the other side of the hole.

In this section, we show how such an estimation of the size of a trough can be collectively performed—and a decision collectively taken—by the *s-bots* forming the *swarm-bot*. We designed a set of experiments in order to test the ability of a *swarm-bot* to bridge a gap of varying size. This test is intended to demonstrate how the simple controllers developed for hole avoidance generalise to a collective decision-making mechanism for discriminating between situations that can be faced by a *swarm-bot* from situations that could be too hazardous even for a large connected structure.

4.1 Experimental setup

S-bots are controlled by the same neural network evolved for hole avoidance, described in Section 3. Therefore, the controller takes as input the traction force perceived by the *s-bot* and the readings coming from the four ground sensors. Recall that ground sensors are simple proximity sensors pointing to the ground. These sensors can therefore be used also to estimate the depth of a hole or the width of a nearby trough, as they have an inclination of 30 degrees with respect to a horizontal plane. In fact, if the trough is not too wide, an *s-bot* near the border would perceive the opposite edge, having different perceptions with varying width. However, this applies only for small troughs, having a width of 2-4 cm. In all the other cases, the opposite edge is not perceived and therefore the size of the trough cannot be estimated by a single *s-bot*.

The *swarm-bot* is placed in an arena divided by a trough (see Figure 4). We test *swarm-bots* of different size—4, 9, and 16 *s-bots* connected in a square formation—that have to confront with a trough of width varying from 2 to 30 cm. In each trial, the square structure is rotated choosing every time a new random orientation, indicated by the vector \mathbf{A} and the corresponding angle α in Figure 4. Independently from the direction of the *swarm-bot*'s structure, the *s-bots* are initialised with their chassis aligned in a same random direction, indicated by the vector \mathbf{B} and the corresponding angle β in Figure 4. The angle varies in the range $[-45, 45]$ degrees with respect to the direction perpendicular to the trough. As a consequence of the initial alignment of the chassis, no coordination phase is required at the beginning of the trial, but the *swarm-bot* can directly move in a coherent way toward the trough. These settings let us focus on the ability to pass over the trough rather than on the coordination abilities of the *swarm-bot*.




Figure 4: Experimental setup: a *swarm-bot* composed of 16 *s-bots*, represented as grey circles, has to confront with a trough, represented as a dark rectangle. The initial orientation of the square structure is randomly chosen, and it is indicated by the vector \mathbf{A} and the angle α . The *s-bots* start with the same random orientation of the chassis, indicated within each circle by an arrow parallel to the vector \mathbf{B} and the angle β . The *swarm-bot* is initially positioned at a distance D_e from the first edge of the trough.

We measure the performance of the *swarm-bot* passing over a trough computing the distance covered by the group along the x axis, which is perpendicular to the trough (see Figure 4). In particular, the performance f is given by the maximum distance covered in the direction of the trough during the trial, given by the following equation:

$$f = \frac{\max_{t \in [0, T]} d_x(t)}{D}, \quad d_x(t) = \mathbf{x}(t) - \mathbf{x}(0), \quad (1)$$

where $\mathbf{x}(t)$ is the position of the *swarm-bot* centre of mass on the x -axis at time t , T is the length of the trial and D is the maximum distance the *swarm-bot* can cover in T simulation cycles. If the *swarm-bot* is not able to pass over the trough, the performance f takes values around D_e/D , where D_e is the distance of the first edge of the trough from the *swarm-bot*'s starting position (see Figure 4). In fact, the trough is always reached due to the initialisation of the *swarm-bot*, and

therefore the maximum distance $d_x(t)$ is obtained in the vicinity of the trough. Higher performance values are obtained whenever the *swarm-bot* is able to pass over the trough.

Note that the performance metric f has been explicitly defined to evaluate the behaviour of passing over a trough. Consequently, it assigns a high score to those situations in which the gap is passed, while an avoidance action corresponds to a low value. This low value should not be considered as a failure, but it should be rather used to distinguish in which conditions the *swarm-bot* performs an avoidance or a passing action, as we show in Sections 4.2 and 4.3.

4.2 Results

A qualitative analysis of the behaviour produced by the controllers evolved for hole avoidance when used in an arena presenting small holes reveals that: (i) if the width of the gap is small enough (2-4 cm), an individual *s-bot* does not perceive it as a hazard—the activation of the ground sensors is rather low—and therefore the *swarm-bot* can pass over the trough. Here, physical connections provide the support for the suspended *s-bots*. (ii) If the width of the gap is bigger, the individual *s-bot* perceives the trough via the ground sensors and reacts consequently. However, the *s-bot* may be pushed out of the borders by the actions of the remaining *s-bots* in the formation. In this case, it may reach the opposite side of the trough, bridging the gap and letting also other *s-bots* pass (see Figure 5, left). (iii) If the gap cannot be bridged by the *swarm-bot*, a normal hole avoidance behaviour is performed and the *swarm-bot* will move away from the hole (see Figure 5, right).

Using the performance metric described in equation (1), we performed a quantitative analysis to evaluate the ability of a *swarm-bot* in passing over a trough. We performed 100 evaluation trials per experimental setup, systematically varying




Figure 5: Trajectories drawn by a *swarm-bot* composed of 9 *s-bots* in a square formation. Left: the *swarm-bot* is able to pass over a 10 cm wide trough. Right: the *swarm-bot* avoids a 20 cm wide trough, which could be too large to be bridged.




Figure 6: Performance of a *swarm-bot* passing over a trough. Performance is defined according to equation (1). Each box-and-whiskers plot represents 100 evaluation trials. Boxes covers the interquartile range, while whiskers extend to the last data-point within 1.5 times the interquartile range. The small circles are outliers. The dark grey area represents the performance for those distances occupied by the trough.

the *swarm-bot* size and the trough width—i.e., 100 trials for each size/width pair. Each trial lasts $T = 300$ simulation cycles, that correspond to 30 seconds of real time. The results of this analysis are plotted in Figure 6. The plot shows, for each trough width, the performance of the three studied *swarm-bots*. The light grey area that spans over the various trough widths gives an indication of the position of the trough with respect to the performance metric. The bottom edge of the grey area corresponds to the performance of D_e/D achieved when the *swarm-bot* reaches the first edge of the trough. Whenever the gap is bridged and the *swarm-bot* finds itself on the other side of the arena, the performance has higher values than the grey area. If the *swarm-bot* is not able to bridge the gap, than the performance obtained is within the grey area or lower.

From the results shown in Figure 6 it is possible to notice how the performance generally decreases as the width of the gap increases: a good performance can be observed for small gaps, followed by a transition that leads to poor performance for large troughs. Looking at the performance of the 4-individual *swarm-bot*, we notice that for gaps of 2-6 cm the performance is always higher than the grey area, indicating that the *swarm-bot* is systematically passing over the trough. An abrupt change in the performance can be observed for a trough 8-12 cm wide. For these sizes, a transition can be observed, in which the *swarm-bot* stops passing over the trough systematically and sometimes avoids it, depending on its orientation with respect to the trough. For the 12 cm trough the *swarm-bot* is successful only sporadically, while for bigger sizes—14 cm or more—the avoidance behaviour is always performed.

The situation is different for bigger structures. In fact, the bigger the *swarm-bot*, the larger the gap that can be passed. For a 9-individuals *swarm-bot*, the performance drops for gaps 10-18 cm wide. For smaller sizes, the *swarm-bot* is always able to bridge the gap. For bigger sizes, the *swarm-bot* always avoid it.

Concerning the 16 individuals *swarm-bot*, we can notice that the transition starts with a width of 12 cm. However, in this case the performance drop is more graceful, as the structure is large enough to bridge troughs up to 30 cm. In fact, it is possible to notice that there are trials in which the performance is above the grey area for all test conditions.

It should be noticed that in some cases even if the gap is bridged, the *swarm-bot* is not able to efficiently coordinate in order to pass on the other side. In fact, once the gap is encountered and bridged by some of the *s-bots*, a new coordination phase is triggered which generally leads to the choice of a new direction of motion, that may let the *swarm-bot* retrace its steps. Furthermore, the coordination phase over the trough is time-consuming, and the *swarm-bot* may not be able to completely pass over the trough in the limited available time.

4.3 Discussions

The behaviour presented above can be considered conservative, as the avoidance is in general preferred to the passing over the trough. This is not surprising because the behaviour was evolved explicitly for the hole avoidance task. Therefore, a trough can be estimated too large to be bridged even when the *swarm-bot* is big enough to pass over it. However, looking at the performance shown in Figure 6, we can notice that the *swarm-bots* perform reasonably well with respect to their physical constraints. In fact, given the size of a 4-individual *swarm-bot*, the maximum width of a trough that can be bridged is about 14 cm. Our results show that from this width on, the *swarm-bot* always performs an avoidance action, while the *swarm-bot* is able to pass over narrower troughs, even if not systematically. A similar situation can be observed for the case of 9 and 16 *s-bots*, which are respectively characterised by the maximum width of 22 and 30 cm.

Whether a trough is avoided or bridged depends on multiple factors, among which the orientation of the *swarm-bot* and its direction of motion when it first approach the trough. In fact, the collective behaviour of passing over a trough relies on a delicate balance between the forces exerted by the *s-bots* that touch the ground and the missing influence of those *s-bots* that are suspended over the gap. The suspended *s-bots* cannot influence the behaviour of the group and the dynamics of the *swarm-bot* are governed by fewer *s-bots*. Every *s-bot* that perceives a hole will react trying to change its direction of motion and trying to influence the behaviour of the whole group by exerting a traction force. However, the bigger the size of the *swarm-bot*, the bigger the inertia of the physical structure. Once the *swarm-bot* reaches an edge, its inertia will cause some *s-bots* to be pushed out, over the gap. In fact, few *s-bots* have a small effect on the overall behaviour of the group. When a sufficient number of *s-bots* is suspended out of the arena, the forces exerted by those *s-bots* that reach the edge can be perceived by the whole group, and they will trigger a change in the direction of motion of the *swarm-bot* in order to avoid falling. If some of the suspended *s-bots* reach the other side of the trough, they start again to have an influence on the rest of the group. First, they align with the current direction of motion, and afterwards they contribute to the gap passing behaviour pulling the whole structure on the other side of the gap. This emergent behaviour can be considered self-organised, as it depends on the interactions among individuals and on clear feedback loops: the conformist tendency of the *s-bots* in following the average direction of the group constitutes a positive feedback, while the tendency to avoid a hole of the individual *s-bots* and the missing influence of the suspended *s-bots* constitute the negative feedback.

In conclusion, the collective behaviour of passing over a trough can be considered an emergent decision-making mechanism that allows a *swarm-bot* to discriminate between those troughs that are small enough to be safely bridged and those that

are not. We observed that the width of the troughs that can be traversed varies, depending on the size of the *swarm-bot*: the bigger the size, the wider the trough. Therefore, it is possible to conclude that through a self-organising process, the *swarm-bot* is able to collectively estimate the width of the trough, and consequently it is able to take the correct decision about the way to move.

5 Conclusions

Collective decisions are an important issue whenever a swarm robotic system is taken under consideration. Collective decisions allow to keep a low complexity of the individual behaviours, while obtaining more complex behaviours at the group level. The work presented in this paper shows one particular case in which a complex decision—such as passing over a trough or avoiding to confront with it—can be collectively taken relying only on simple behavioural rules. These rules followed by each *s-bot* do not contain any reference to the behaviour of passing over a trough. However, they result in a self-organising process that allow an estimation of the size of the trough and therefore an emergent decision-making process.

We claim that similar self-organised behaviours could be exploited for other problems requiring a collective decision-making process. However, designing a self-organising control system for a swarm robotic system is not a trivial task. From an engineering perspective, the design problem is generally decomposed into two different phases: (i) the behaviour of the system should be described as the result of interactions among individual behaviours, and (ii) the individual behaviours must be encoded into controllers. Both phases are complex because they attempt to decompose a process (the global behaviour or the individual one) that is a result of dynamical interactions among its sub-components (interactions among individuals or between individual actions and the environment) [5]. These dynamical aspects are in general difficult to be predicted by the observer: referring to the case discussed in this paper, predicting the collective dynamics of passing over a trough is difficult, and even understanding them through the observation of the evolved behaviour requires a considerable effort.

We believe that evolutionary robotics techniques are the tools to be exploited to obtain self-organising behaviours in a group of robots. Evolution bypasses the problem of decomposition at both the level of finding the mechanisms that lead to the global behaviour and at the level of implementing those mechanisms in a controller for the *s-bots*. In fact, evolution relies on the evaluation of the system as a whole, that is, on obtaining the desired global behaviour starting from the definition of the individual ones. Moreover, evolution can exploit the richness of possible solutions offered by the dynamic agent-environment interactions, that could not be apparent *a priori* to the experimenter [8].

In future works, we will continue the research on the evolution of self-organising behaviours related to collective decision-making processes. We will study how to obtain collective decision mechanisms in a swarm robotic system that is able to take into account not only environmental cues, but also temporal ones. This will allow the synthesis of behaviours that change in relation to the persistence of a perceptual cue for a certain amount of time [12]. A similar situation can be studied in the context of the behaviour of passing over a trough: in this case, *s-bots* should first search for a passage that could lead them to the opposite side of the arena. In case such a passage does not exist, the *s-bots* should recognise that they have to self-assemble in a *swarm-bot* in order to cope with the gap. The absence of a passage could be discovered only by means of temporal cues, such as the persistence of the perception of the trough.

References

- [1] G. Baldassarre, D. Parisi, and S. Nolfi. Coordination and behaviour integration in cooperating simulated robots. In S. Schaal, A. Ijspeert, A. Billard, S. Vijayakumar, J. Hallam, and J.-A. Meyer, editors, *From Animals to Animats VIII. Proceedings of the 8th International Conference on Simulation of Adaptive Behavior*, pages 385–394. MIT Press, Cambridge, MA, 2004.
- [2] R. Beckers, J.-L. Deneubourg, and S. Goss. Modulation of trail laying in the ant *Lasius niger* (Hymenoptera: Formicidae) and its role in the collective selection of a food source. *Journal of Insect Behavior*, 6:751–759, 1993.
- [3] R. Beckers, J.-L. Deneubourg, S. Goss, and J. M. Pasteels. Collective decision making through food recruitment. *Insectes Sociaux*, 37:258–267, 1990.
- [4] S. Camazine, J.-L. Deneubourg, N. Franks, J. Sneyd, G. Theraulaz, and E. Bonabeau. *Self-Organization in Biological Systems*. Princeton University Press, Princeton, NJ, 2001.
- [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. IJ.34.
- [6] 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.
- [7] S. Nolfi. Evolving robots able to self-localize in the environment: The importance of viewing cognition as the result of processes occurring at different time scales. *Connection Science*, 14(2):231–244, 2002.
- [8] S. Nolfi and D. Floreano. *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*. MIT Press/Bradford Books, Cambridge, MA, 2000.
- [9] T. D. Seeley. *The Wisdom of the Hive*. Harvard University Press, Cambridge, MA, 1995.
- [10] V. Trianni, T. H. Labella, and M. Dorigo. Evolution of direct communication for a swarm-bot performing hole avoidance. In M. Dorigo, M. Birattari, C. Blum, L. M. Gambardella, F. Mondada, and T. Stützle, editors, *Ant Colony Optimization and Swarm Intelligence – Proceedings of ANTS 2004 – Fourth International Workshop*, volume 3172 of *Lecture Notes in Computer Science*, pages 131–142. Springer Verlag, Berlin, Germany, 2004.
- [11] V. Trianni, S. Nolfi, and M. Dorigo. Cooperative hole avoidance in a *swarm-bot*. *Robotics and Autonomous Systems*, 2005. To appear.
- [12] E. Tuci, V. Trianni, and M. Dorigo. Feeling the flow of time— through sensory-motor coordination. *Connection Science*, 16(4):301–324, 2004.
- [13] T. Ziemke and M. Thieme. Neuromodulation of reactive sensorimotor mappings as a short-term memory mechanism in delayed response tasks. *Adaptive Behavior*, 10(3-4):185–199, 2002.