

Cooperative hole avoidance in a *swarm-bot*

Vito Trianni^{a,*}, Stefano Nolfi^b, Marco Dorigo^a

^a IRIDIA - Université Libre de Bruxelles, Bruxelles, Belgium

^b Institute of Cognitive Sciences and Technologies - CNR, Roma, Italy

Received 15 October 2004; received in revised form 19 January 2005; accepted 13 September 2005

Available online 28 November 2005

Abstract

In this paper, we study coordinated motion in a swarm robotic system, called a *swarm-bot*. A *swarm-bot* is a self-assembling and self-organising artifact, composed of a swarm of *s-bots*, mobile robots with the ability to connect to and disconnect from each other. The *swarm-bot* concept is particularly suited for tasks that require all-terrain navigation abilities, such as space exploration or rescue in collapsed buildings. As a first step toward the development of more complex control strategies, we investigate the case in which a *swarm-bot* has to explore an arena while avoiding falling into holes. In such a scenario, individual *s-bots* have sensory–motor limitations that prevent them navigating efficiently. These limitations can be overcome if the *s-bots* are made to cooperate. In particular, we exploit the *s-bots*' ability to physically connect to each other. In order to synthesise the *s-bots*' controller, we rely on artificial evolution, which we show to be a powerful tool for the production of simple and effective solutions to the hole avoidance task.

© 2005 Elsevier B.V. All rights reserved.

Keywords: Evolutionary robotics; Swarm intelligence; Swarm robotics; *Swarm-bot*

1. Introduction

The first problem to be considered when trying to control an autonomous robot is to make it move efficiently in a given environment. Depending on the robot, this task can be rather simple (i.e., the motion of a wheeled robot) or particularly complex (i.e., walking for a humanoid robot). Also the environment in which the robot is placed influences the complexity of the problem: a flat terrain is clearly less challenging than a rough terrain with holes and obstacles. An additional source of complexity is found in the *coordinated motion* task, in which the robotic system is composed of a number of independent entities that have to coordinate their actions in order to move coherently.

Coordinated motion is a well studied behaviour in biology, being observed in many different animal species. For example, we can think of flocks of birds flying in a coordinated manner, or of schools of fish swimming in perfect unison. These examples are not only fascinating for the charming patterns

they create, but they also represent interesting instances of self-organised behaviours. Many researchers have provided models for schooling behaviours, and thus replicated real-world behaviours in artificial life simulations (see [2, Chapter 11]). Similarly, groups of artificial fish (*e-boids*) have been evolved to display schooling behaviours, obtaining interesting results [13]. Finally, evolutionary computation has been used to evolve coordinated motion behaviours in small groups of physical robots [10].

Coordinated motion is a problem of fundamental importance within the SWARM-BOTS project,¹ wherein this research is conducted. The project aims at the development of a new robotic system, called a *swarm-bot* [4,8]. A *swarm-bot* is defined as an artifact composed of simpler autonomous robots, called *s-bots*. An *s-bot* has limited acting, sensing and computational capabilities. However, an *s-bot* can create physical connections with other *s-bots*, thereby forming a *swarm-bot* that is able to solve problems the single individual cannot cope with. Coordinated motion is a basic ability that the *swarm-bot* should display: a *swarm-bot* should move

* Corresponding author. Fax: +32 2 650 2715.

E-mail addresses: vtrianni@ulb.ac.be (V. Trianni), s.nolfi@istc.cnr.it (S. Nolfi), mdorigo@ulb.ac.be (M. Dorigo).

¹ A project funded by the Future and Emerging Technologies Programme (IST-FET) of the European Community, under grant IST-2000-31010.

coherently across the environment as a result of the cooperation of the *s-bots* assembled in a single structure [1].

Another basic ability for the *swarm-bot* is coping with rough terrains, holes, gaps or narrow passages. All-terrain navigation is an important feature for an intelligent autonomous system, opening many possible application scenarios, like space exploration or rescue in a collapsed building. Research in this direction has focused mainly on the development of rovers provided with articulated wheels or tracks, such as the *sojourner* [7]. A different approach is presented by reconfigurable robotics, where robots can adopt different shapes in order to cope with varying environmental conditions [3,11,14].

The *swarm-bot* concept puts together the advantages given by autonomous rovers and by self-reconfigurable robots for all-terrain navigation. In fact, similarly to rovers, each *s-bot* is fully autonomous in its control and is capable of moving on moderately rough terrains. On the other hand, whenever individual abilities are not sufficient, the *s-bots* can rely on their ability to display collective behaviours. This allows the *swarm-bot* to exploit features similar to those of reconfigurable robots, such as (i) cooperation among assembled elementary units, (ii) physical support through inter-unit connections and (iii) exploitation of different multi-robot configurations (shape changing). Some pioneer work with comparable features has been done by Hirose et al. [5]. However, to the best of our knowledge, this work remained at the level of a proposal. We significantly improve on this proposal within the SWARM-BOTS project [8].

In summary, we aim at studying all-terrain navigation as the result of the cooperation among *s-bots*, which can self-assemble and build structures that can cope with hazardous situations like avoiding a hole or passing over a trough. In such cases, rigid connections serve as support for those *s-bots* that are suspended over the gap. This approach to all-terrain navigation also has a natural counterpart in ants of the species *Aecophylla longinoda* [6], which are able to build chains connecting one to the other, creating bridges that facilitate the passage of other ants.

In this paper, we study an instance of the family of all-terrain navigation tasks, that is, *hole avoidance*. A *swarm-bot* has to perform coordinated motion in an environment that presents holes too large to be traversed. Thus, holes must be recognised and avoided, so that the *swarm-bot* does not fall into them. The challenges issued by this task are described in Section 2, along with the experimental set-up used for our experiments. Sections 3 and 4 are dedicated to the description of the obtained results. Finally, Section 5 concludes the paper.

2. Evolution of hole avoidance behaviours

The hole avoidance task has similar aspects to a common obstacle avoidance scenario, in which there are zones that should not be traversed. However, this task presents challenges that are not found in obstacle avoidance. Above all, two important differences should be highlighted: (i) the failure in avoiding a hole leads to a situation from which the




Fig. 1. (a) Two *s-bots* physically connected. This illustrates how two robots can cooperate to traverse a gap neither could navigate alone. The *s-bot* is provided with a traction system comprising wheels and tracks, a rotating turret holding the rigid and the flexible grippers, and many sensors. (b) The simulated *s-bot* model. The turret is transparent to show the chassis (centre sphere), the motorised wheels (lighter spherical wheels) and the passive wheels (darker spherical wheels). The position of the gripper is shown with an arrow painted on the *s-bot*'s body. Ground sensors are displayed as lines exiting from the *s-bot*.

robot cannot recover, while a collision with an obstacle is not harmful as long as the robot is not damaged; (ii) the sensory configuration of the *s-bot* and its dynamics make it difficult to perceive and avoid holes with particular shape and dimension, despite the fact that obstacles with identical two dimensional characteristics of shape and size could be easily detected.² These difficulties led us to the choice of exploiting the cooperation among the *s-bots* assembled in a *swarm-bot* configuration. However, this choice introduces new challenges. First, *s-bots* should coordinate their overall motion. Second, *s-bots* have to recognise the presence of a hole, communicate it to the whole group and re-organise to choose a safer direction of motion.

The controllers for the *s-bots* are obtained using *artificial evolution*. There are multiple motivations behind this choice [9]. In particular, in a distributed multi-robot context such as the one considered within the SWARM-BOTS project, hand-crafting the controllers may be too complex. Here, artificial evolution can bypass this difficulty, as it directly tests the behaviour displayed by the robots embedded in their environment. Furthermore, artificial evolution can exploit the richness of solutions offered by the complex dynamics resulting from robot–robot and robot–environment interactions [4].

Fig. 1(a) shows two connected *s-bots*.³ In this paper, however, experiments are performed in simulation, using a software based on VortexTM, a 3D rigid body dynamics engine. In the following sections, we describe the experimental set-up: we give details about the simulation model in Section 2.1, about the evolutionary algorithm in Section 2.2 and about the fitness function used for the evolution of hole avoidance behaviours in Section 2.3.

2.1. The simulation model

We have defined a simple *s-bot* model that both allows fast simulations and preserves those features of the real *s-bot* that

²Experiments performed using a single *s-bot* revealed that the narrow corners of a hole are hard to perceive, and therefore to avoid (data not shown).

³Details regarding the hardware and simulation of the *swarm-bot* can also be found in the project web-site (www.swarm-bots.org).

we considered most important (see Fig. 1b). The simulated *s-bot* is composed of a cylindrical turret (radius: 6 cm, height: 6 cm), connected to a chassis by a motorised hinge joint that allows relative rotation between the two bodies. The chassis is modelled as a sphere (radius: 1.4 cm) to which four spherical wheels are connected (radius: 1.5 cm). The lateral wheels are connected to the chassis by a motorised joint and a suspension system and are responsible for the motion of the *s-bot*. The front and back wheels are passive and they can rotate in any direction. The gripper that allows connections between *s-bots* is simulated by creating a link between the two turrets. The gripping position is indicated by an arrow painted over the turret. In this work, connections are established at the beginning of the simulation and are never released.

Each *s-bot* is provided with a *traction sensor* placed at the turret–chassis junction. It detects the direction and the intensity of the traction force that the turret exerts on the chassis. The traction sensor, integrating all the pulling/pushing forces created by the movement of the connected *s-bots*, provides an indication of the average direction toward which the *swarm-bot* is trying to move as a whole.⁴ Besides the traction sensor, we also make use of four *ground sensors*, which are infrared proximity sensors evenly distributed around the chassis of the *s-bot* and pointed toward the ground. Ground sensors are used to perceive the presence of a hole in the vicinity of the *s-bot*.

Each *s-bot* can control its wheels independently. The maximum angular speed has been set to 10 rad/s, which corresponds to a maximum speed of the *s-bot* of 0.15 m/s. In addition, the motor controlling the rotation of the turret is used. Its angular speed is set to half of the difference between the angular speed of the left and right wheels. This motor assists the rotation of the chassis with respect to the turret even when one or both wheels of the *s-bot* are not in contact with the ground [1].

We designed a square arena (side 3 m) that contains four square holes (side 60 cm, see Fig. 2(b)) evenly distributed. The *swarm-bot* consists of a linear structure made of four *s-bots*, which are rigidly connected by means of their grippers. Each *s-bot* is controlled by a simple perceptron, a neural network connecting its sensory inputs to the motor outputs. The network has eight sensory inputs: four are dedicated to the readings coming from the ground sensors, and the other four encode the intensity and direction of traction (for more details, see [1]). The neural network is provided with one bias unit and two outputs that control the two wheels and the turret/chassis motor. This perceptron has a total of 18 connections for which weights are evolved.

2.2. The evolutionary algorithm

A simple generational evolutionary algorithm is used for the synthesis of neural controllers. The initial population is composed of $\mu = 100$ randomly generated genotypes. Each genotype is binary encoded, and is mapped into a neural

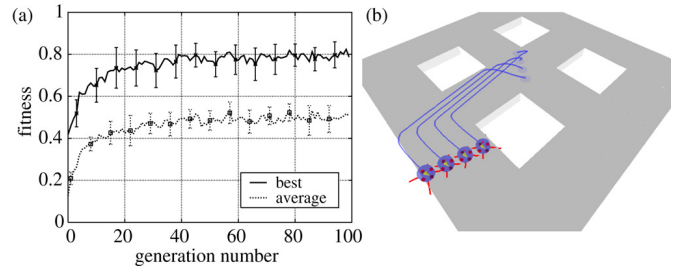


Fig. 2. Hole avoidance results: (a) Average fitness over 10 replications of the experiment. (b) Trajectories displayed by a *swarm-bot* performing hole avoidance.

network controller for a single *s-bot*. Each weight of the neural network ranges in the interval $[-10, 10]$ and is represented in the genotype by eight bits. Therefore, the genotype length corresponds to $L = 18 \times 8 = 144$ bits. In every generation, all genotypes of the population are evaluated using the fitness function defined in the following section. The best $\lambda = 20$ genotypes of each generation are allowed to reproduce, each generating $\mu/\lambda = 5$ offspring, which are exact copies of the parent. Afterwards, each offspring is mutated — i.e., each bit has a probability $2/L$ of being flipped. Parents are not copied to the offspring population (no elitism). No recombination operator is applied. The evolutionary experiment lasts 100 generations. This algorithm is very simple and straightforward, and we found that it is sufficient to evolve simple and efficient controllers for groups of robots [4].

2.3. The fitness evaluation

The neural network controller obtained from a genotype is cloned and downloaded in each of the $n = 4$ *s-bots* involved in the experiment, so that all *s-bots* are *homogeneous* in their control. The fitness of the genotype is computed by measuring the performance of the corresponding group of four robots. The fitness \mathcal{F} of a genotype is a random variable, because of the random initialisation of the positions and orientations of the *s-bots*. Its expected value F can be estimated by evaluating the behaviour of the *swarm-bot* for a number M of trials and then averaging the obtained values. Therefore, in each trial e we compute a sample F_e of the random variable \mathcal{F} . In these experiments, we use the sampling size $M = 5$.

The fitness function is designed to favour coordinated motion, exploration of the arena and a fast reaction to the detection of a hole. The fitness estimation F_e in each trial is given by the average of two components, F_{e_1} and F_{e_2} (see below). In order to compute the fitness components, we divide each trial e into two sub-trials, e_1 and e_2 . In the former, we test the controller for its ability to perform coordinated motion in a flat environment. Here the *s-bots* start connected in a linear formation, with the orientation of their chassis randomly initialised. They are rewarded for the ability to move as far as possible from their initial position. Note that this indirectly implies an ability to display coordinated movements. Therefore, the fitness estimation F_{e_1} is computed as the distance covered

⁴ This particular kind of sensor proved to be of fundamental importance for the evolution of coordinated motion in a *swarm-bot* [1].

Table 1
Performance of the best individuals for each replication of the experiment, averaged over 100 trials

| Replication | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 |
|--------------|------|------|------|------|------|------|------|------|------|------|
| Fitness avg. | 0.66 | 0.65 | 0.65 | 0.61 | 0.58 | 0.64 | 0.69 | 0.64 | 0.66 | 0.65 |
| Fitness std. | 0.14 | 0.12 | 0.14 | 0.19 | 0.10 | 0.16 | 0.12 | 0.15 | 0.13 | 0.16 |

The mean value and the standard deviation are reported.

by the group:

$$F_{e_1} = \frac{\|\mathbf{X}(0) - \mathbf{X}(T_{e_1})\|}{D}, \quad (1)$$

where $\mathbf{X}(t)$ is the coordinate vector of the centre of mass of the group, and D is the maximum distance achievable. The sub-trial e_1 lasts $T_{e_1} = 150$ simulation cycles, each cycle corresponding to 100 ms of real time.

In sub-trial e_2 , the *s-bots* are positioned in the centre of the arena with holes, and start connected in a linear formation. Their chassis are all initialised with the same random orientation. In this way, there is no need for a coordination phase at the beginning of the sub-trial, the focus being on hole avoidance. Also the chain is randomly oriented at the beginning of each sub-trial. The sub-trial lasts $T_{e_2} = 200$ cycles.⁵ The fitness estimation F_{e_2} is given by the product of two sub-components: the *survival* sub-component F_s and the *exploration* sub-component F_x . The former rewards only those behaviours that reach the end of the trial without resulting in a fall. It is computed as follows:

$$F_s = \begin{cases} 1, & \text{if } T_s = T_{e_2}; \\ 0, & \text{otherwise;} \end{cases} \quad (2)$$

where T_s is the number of cycles the *swarm-bot* “survived” without falling into a hole. This sub-component introduces a strong selective pressure towards safe behaviours, as it penalises every fall, even if it happens at the end of the sub-trial.⁶ The second sub-component is designed to favour those genotypes that are able to better explore the arena. In this case, the arena is virtually divided in 25 square zones (side: 60 cm). The genotype is rewarded for the number of visited zones during the sub-trial, formalised as follows:

$$F_x = \frac{z(T_s)}{Z(T_{e_2})}, \quad (3)$$

where $z(T_s)$ is the number of visited zones at cycle T_s and $Z(T_{e_2}) = 5$ corresponds to the maximum number of zones that can be visited in T_{e_2} cycles. A zone is considered visited if the *swarm-bot*'s centre of mass lies within the corresponding square area. Passing multiple times over a visited zone does not correspond to any additional reward.

⁵ A longer time is needed in this sub-trial in order to let the *swarm-bot* interact with holes and edges of the arena as much as possible.

⁶ However, we cannot ensure that a *swarm-bot* that does not fall within 200 simulation cycles will never fall in the future. A trade-off value must be chosen which maximises the selective pressure while minimising the duration of the evolutionary experiment, and therefore the required computation time.

3. Obtained results

We performed 10 replications of the evolutionary experiment, every time starting with a different population of randomly generated genotypes. The average fitness values of the best individuals and of the whole populations, computed over all the replications, are plotted against the generation number in Fig. 2. The plot indicates that the evolutionary experiments were successful: the average fitness value of the best individuals reaches 80% of the maximum achievable value, which cannot normally be reached due to the particular experimental set-up.⁷

We tested the performance of the controllers evolved in the 10 different replications of the evolutionary experiment. We evaluated the 10 best individuals of the last generation obtained in the different replications, averaging their fitness F_e over 100 different trials. The corresponding results are shown in Table 1. All individuals perform reasonably well. It can be noted that their performance is lower than 0.8, which is the average performance achieved at the end of the 10 replications of the evolutionary experiment, as shown in Fig. 2a. This is due to the small sampling size M used for the estimation of the fitness during the evolution ($M = 5$ samples). In fact, a small sampling size usually leads to an over-estimation of the fitness of the best individual during the evolution. Thus, the post-evaluation analysis with a larger sampling size ($M = 100$ trials, in this case) gives a better approximation of the real performance of the evolved controllers.

Direct observation of the behaviours evolved showed that all solutions rely on similar strategies. We observed the evolved behaviours placing the *swarm-bot* in the arena with holes, and starting with different orientations of the chassis of the *s-bots*⁸ (see Fig. 2(b)). At the beginning, the *s-bots* start to move in their initial direction, 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. Then, when one

⁷ The maximum value for F_e could be reached only if in the first sub-trial *s-bots* started with their chassis perfectly aligned, so that no coordination phase is required, allowing the *swarm-bot* to cover the maximum distance. This is normally not the case due to the random initialisation of the orientations of the chassis.

⁸ See <http://www.swarm-bots.org/hole-avoidance.html> for some movies of these behaviours.

s-bot perceives an edge with its ground sensors, it rotates the chassis and changes the direction of motion in order to avoid falling. This change in direction creates a traction force for the other *s-bots*, which they perceive by means of their traction sensors. At this point, a new coordination phase is triggered, which ends up in a new direction of motion leading 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 *swarm-bot* to a safer location.

4. Generalisation

The evolved strategy for hole avoidance is very robust, being able to work in a number of different situations. This is a result of the physical connections among the *s-bots* and, above all, of the use of the traction sensors. First of all, the evolved strategies are independent of the shape and position of the holes in the arena. We also tested the scalability of the evolved controllers varying the size and the shape of the *swarm-bot*. We observed that the evolved controllers perform well in many different conditions. For example, Fig. 3(a) shows the case of a *swarm-bot* comprising eight *s-bots* connected in a “star” shape. The *swarm-bot* is placed in a square arena without holes, but with open borders. The *swarm-bot* is still able to avoid falling out of the arena, notwithstanding the higher inertia of the star formation.

Another interesting feature of the evolved controllers is that they are able to perform collective obstacle avoidance. In fact, when an *s-bot* hits an obstacle, its turret exerts a force on the chassis in a direction opposite to the obstacle. This force is felt as a traction pulling the *s-bot* away from the obstacle. In response to this traction, the *s-bot* rotates its chassis to cancel it, as explained before. Moreover, the rigid connections between *s-bots* transmit the force resulting from the collision to the whole group, triggering a fast change in the direction of movement of the *swarm-bot*. As shown in Fig. 3(b), the *swarm-bot* is able to avoid both holes and obstacles, represented here by walls surrounding the arena. It is worth noting that the traction sensor works as an omni-directional bumper distributed on the whole body of the *swarm-bot*, allowing collective obstacle avoidance.

Finally, we tested the evolved controllers when the *s-bots* are linked using flexible, rather than rigid, connections. Flexible connections allow the rotation of the connecting *s-bots* around the turret of the connected *s-bot*. The use of this type of connection allows the shape of the *swarm-bot* to change during motion. Because of the flexibility of the connections, traction can be transmitted only in the radial direction, but not in the tangential one. Nevertheless, the evolved strategies still work. We performed tests with both a “star” and a chain formation composed of eight *s-bots* each. As shown in Fig. 3(c) and (d), the flexible formations are able to perform coordinated motion, obstacle and hole avoidance, changing




Fig. 3. Generalisation properties. The trajectories and the final position of the *swarm-bot* are shown. (a) Size and shape change. A “star” formation is tested in a square arena (grey area) without holes but with open borders. The trajectories indicate that the *swarm-bot* is able to avoid falling out, even if some *s-bots* are pushed out from the border. (b) Obstacle avoidance. The square arena with holes (grey area) is surrounded by walls (dark grey borders). The *swarm-bot* proves able to avoid both holes and obstacles. (c) Obstacle and hole avoidance of a “star” formation with flexible connections. Here the cylindrical obstacles (light grey objects) create a narrow passage with the edge of the arena (grey area), which is faced by the *swarm-bot* through reconfiguration of its shape. (d) Hole avoidance of a big linear formation with flexible connections. Here the *swarm-bot* completely deforms when it reaches the edge of the arena (grey area), therefore adapting its shape.

shape when passing through narrow passages. The flexible formation adapts more easily to the environment, and in some situations can avoid holes more efficiently than a rigid structure. In fact, the *s-bots* do not completely feel the inertia of the *swarm-bot*, because they can change their relative positions, therefore deforming the structure and adapting it to the edge of the hole. In order to assess to what extent flexible connections among *s-bots* make the system more efficient, we compared the performance of three different *swarm-bot* configurations using both rigid and flexible links. The first configuration is the standard linear formation with four *s-bots*, the second configuration is the “star” formation with eight *s-bots* shown in Fig. 3(a) and the third is a circular formation, again composed of eight *s-bots*. The performance was measured over 500 simulation cycles, and the two fitness components described in Section 2 were computed.⁹ Table 2 shows the exploration and the survival performance for these formations. The most interesting data is given by the survival factor, which indicates the ability of the *swarm-bot* to avoid falling out of the arena.¹⁰ We can notice that, as far as the linear and the circular formations are concerned, flexible formations are advantageous, leading to a higher survival

⁹ The arena used for the formations comprising eight *s-bots* is larger than before, having a side of 4 m, in order to ease the passage of the larger *swarm-bot*.

¹⁰ The exploration factor is less relevant in this case. In fact, the longer the trial, the higher the probability that the *swarm-bot* retraces his steps visiting already covered zones. This justifies the observed drop in performance when using flexible links.

Table 2
Performance comparison of different *swarm-bot* formations using rigid and flexible connections among *s-bots*

| | Exploration | Survival |
|-----------------|-------------|----------|
| Rigid line | 5.82 | 0.46 |
| Flexible line | 6.52 | 0.66 |
| Rigid star | 5.43 | 0.30 |
| Flexible star | 4.4 | 0.33 |
| Rigid circle | 5.22 | 0.40 |
| Flexible circle | 5.04 | 0.60 |

factor. In contrast, when *s-bots* are connected in a “star” formation, the use of a flexible formation does not correspond to a significant improvement. This fact can be explained considering that the “star” formation is an intrinsically more rigid structure, and therefore it does not allow a drastic shape change. This means that the *swarm-bot* can adapt to the environment (i.e., the presence of holes, edges) only to a limited extent.

5. Conclusions

We presented a set of experiments for the evolution of hole avoidance behaviours in a group of simulated *s-bots* that are physically connected to form a *swarm-bot*. The solutions found by evolution are simple and in many cases they work in different environmental situations. The obtained results suggest that evolution is a suitable tool for synthesising controllers for a group of homogeneous robots. In this case, evolution was able to produce a self-organising system that relies on simple and general rules, a system that is consequently robust to environmental changes and to the number of *s-bots* involved in the experiment.

Our results demonstrate the traction sensor to be a powerful mechanism for achieving coordination in the *swarm-bot*. The traction sensor allows the *swarm-bot* to exploit the complex dynamics arising from the interactions among individual *s-bots* and between the *s-bots* and the environment. It provides robustness and adaptivity features with respect to environmental or structural changes of the *swarm-bot*. Besides, traction forces are used as a sort of communication of the presence of a hazard, allowing the group as a whole — and not only the *s-bots* that perceive the hole — to change direction of motion when heading toward a hole. Finally, the traction sensor can also work as a distributed bumper for the *swarm-bot*, allowing collective obstacle avoidance.

The hole avoidance task represents the first step toward the solution of more difficult problems. We will extend this work in order to obtain controllers that can pass over holes that are sufficiently small, while avoiding falling into holes that are too big to be traversed by the *swarm-bot*. Additionally, we plan to study problems that belong to the all-terrain navigation family, such as coping with uneven terrains. In these perspectives, physical connections among *s-bots* become an essential feature to be exploited. Finally, we intend to investigate functional self-assembly for all-terrain navigation, that is, the problem of forming and disbanding a *swarm-bot* with a functional

shape for the particular environmental conditions and task to be performed, in order to maximise the efficiency in the navigation [12].

Acknowledgements

This work was supported by the SWARM-BOTS project, funded by the Future and Emerging Technologies programme of the European Commission, under grant IST-2000-31010. 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. The authors thank Rehan O’Grady, Roderich Groß, Thomas H. Labella, Shervin Nouyan and Elio Tuci for stimulating discussions and for their feedback during the preparation of this paper. Marco Dorigo acknowledges support from the Belgian FNRS, of which he is a Research Director, 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] G. Baldassarre, S. Nolfi, D. Parisi, Evolution of collective behavior in a team of physically linked robots, in: R. Gunther, A. Guillot, J.-A. Meyer (Eds.), Applications of Evolutionary Computing - Proceedings of the Second European Workshop on Evolutionary Robotics, EvoWorkshops2003: EvoROB, Springer-Verlag, Berlin, Germany, 2003, pp. 581–592.
- [2] S. Camazine, J.-L. Deneubourg, N. Franks, J. Sneyd, G. Theraulaz, E. Bonabeau, Self-Organization in Biological Systems, Princeton University Press, Princeton, NJ, 2001.
- [3] A. Castano, W. Shen, P. Will, CONRO: Towards deployable robots with inter-robot metamorphic capabilities, Autonomous Robots 8 (2000) 309–324.
- [4] M. Dorigo, V. Trianni, E. Şahin, R. Groß, T.H. Labella, G. Baldassarre, S. Nolfi, J.-L. Deneubourg, F. Mondada, D. Floreano, L.M. Gambardella, Evolving self-organizing behaviors for a *swarm-bot*, Autonomous Robots 17 (2–3) (2004) 223–245.
- [5] S. Hirose, T. Shirasu, E.F. Fukushima, Proposal for cooperative robot Gunryu composed of autonomous segments, Robotics and Autonomous Systems 17 (1996) 107–118.
- [6] A. Lioni, C. Sauwens, G. Theraulaz, J.-L. Deneubourg, Chain formation in *Aecophylla longinoda*, Journal of Insect Behaviour 15 (2001) 679–696.
- [7] L. Matthies, E. Gat, R. Harrison, B. Wilcox, R. Volpe, T. Litwin, Mars microrover navigation: Performance evaluation and enhancement, Autonomous Robots 2 (4) (1995) 291–312.
- [8] F. Mondada, G.C. Pettinaro, A. Guignard, I.V. Kwee, D. Floreano, J.-L. Deneubourg, S. Nolfi, L.M. Gambardella, M. Dorigo, SWARM-BOT: A new distributed robotic concept, Autonomous Robots 17 (2–3) (2004) 193–221.
- [9] S. Nolfi, D. Floreano, Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines, MIT Press/Bradford Books, Cambridge, MA, 2000.
- [10] M. Quinn, L. Smith, G. Mayley, P. Husband, Evolving teamwork and role allocation with real robots, in: R.K. Standish, M.A. Bedau, H.A. Abbass (Eds.), Proceedings of the 8th International Conference on Artificial Life, MIT Press, Cambridge, MA, 2002, pp. 302–311.
- [11] K. Støy, W.-M. Shen, P. Will, Global locomotion from local interaction in self-reconfigurable robots, in: W.M. Shen, C. Torras, H. Yuasa (Eds.), Proc. of the 7th International Conference on Intelligent Autonomous

Systems, IAS-7, IOS Press, Amsterdam, The Netherlands, 2002, pp. 309–316.

- [12] V. Trianni, E. Tuci, M. Dorigo, Evolving functional self-assembling in a swarm of autonomous robots, in: S. Schaal et al. (Eds.), *From Animals to Animats VIII. Proceedings of the 8th International Conference on Simulation of Adaptive Behavior*, MIT Press, Cambridge, MA, 2004, pp. 405–414.
- [13] C.R. Ward, F. Gobet, G. Kendall, Evolving collective behavior in an artificial ecology, *Artificial Life* 7 (2) (2001) 191–209.
- [14] M. Yim, D.G. Duff, K.D. Roufas, PolyBot: A modular reconfigurable robot, in: *Proceedings of the 2000 IEEE/RAS International Conference on Robotics and Automation*, vol. 1, IEEE Robotics and Automation Society, Piscataway, NJ, 2000, pp. 514–520.



Vito Trianni received the Laurea (Master of Technology) degree in computer science engineering in 2000 from Politecnico di Milano, Italy, a Master in Information Technology in 2001, from CEFRIEL, Milan, Italy and the Diplôme d'Études Approfondies in 2003 from the Université Libre de Bruxelles, Belgium. He is a Ph.D. student at IRIDIA, Université Libre de Bruxelles. His research interests are in the field of evolutionary robotics, swarm intelligence and self-organisation.



Stefano Nolfi is a senior researcher at the Institute of Cognitive Sciences and Technologies of the National Research Council (CNR) in Rome, Italy, where he leads the Laboratory of Artificial Life and Robotics. His research interests are in the field of neuro-ethological studies of adaptive behaviour in natural and artificial agents and include: evolutionary robotics, artificial life, complex systems, neural networks, and genetic algorithms. He has published more than 70 peer-reviewed articles and a book: *Evolutionary Robotics*, MIT Press, 2000.



Marco Dorigo is a Research Director of the FNRS, the Belgian National Funds for Scientific Research, and co-director of IRIDIA, the artificial intelligence laboratory of the Université Libre de Bruxelles. He is the inventor of the ant colony optimisation metaheuristic and one of the founders of the swarm intelligence research field. Its current research interests include metaheuristics for discrete optimisation, swarm intelligence and swarm robotics. He is the author of three books: *Robot Shaping*, MIT Press, 1998; *Swarm Intelligence*, Oxford University Press, 1999; and *Ant Colony Optimization*, MIT Press, 2004. In 1996 he was awarded the Italian Prize for Artificial Intelligence and in 2003 the Marie Curie Excellence Award.