

Self-Organized Coordinated Motion in Groups of Physically Connected Robots

Gianluca Baldassarre, Vito Trianni, Michael Bonani, Francesco Mondada, *Associate Member, IEEE*, Marco Dorigo, *Fellow, IEEE*, and Stefano Nolfi

Abstract—An important goal of collective robotics is the design of control systems that allow groups of robots to accomplish common tasks by coordinating without a centralized control. In this paper, we study how a group of physically assembled robots can display coherent behavior on the basis of a simple neural controller that has access only to local sensory information. This controller is synthesized through artificial evolution in a simulated environment in order to let the robots display coordinated-motion behaviors. The evolved controller proves to be robust enough to allow a smooth transfer from simulated to real robots. Additionally, it generalizes to new experimental conditions, such as different sizes/shapes of the group and/or different connection mechanisms. In all these conditions the performance of the neural controller in real robots is comparable to the one obtained in simulation.

Index Terms—Distributed control, evolutionary algorithms, intelligent mobile robots, neural networks, swarm intelligence, swarm robotics.

I. INTRODUCTION

SWARM ROBOTICS is an emergent field of collective robotics [1], [2] that studies systems composed of swarms of robots tightly interacting and cooperating to achieve common goals [3]. In a swarm robotic system, although each single robot

Manuscript received August 22, 2005; revised February 10, 2006 and May 22, 2006. This work was supported by the SWARM-BOTS project funded by the Future and Emerging Technologies Programme (IST-FET) of the European Commission under Grant IST-2000-31010 and by the ECagents project also funded by the Future and Emerging Technologies programme (IST-FET) of the European Commission under Grant IST-1940. The information provided is the sole responsibility of the authors and does not reflect the opinion of the community. The community is not responsible for any use that might be made of the data appearing in this paper. The Swiss participants to the project were supported under Grant 01.0012 by the Swiss Government. The work of M. Dorigo was supported by the Belgian Fonds National de la Recherche Scientifique (FNRS) and from the ANTS project, an Action de Recherche Concertée funded by the Scientific Research Directorate of the French Community of Belgium. This paper was recommended by Associate Editor H. Qiao.

G. Baldassarre and S. Nolfi are with the Laboratory of Autonomous Robotics and Artificial Life, Istituto di Scienze e Tecnologie della Cognizione, Consiglio Nazionale delle Ricerche (LARAL-ISTC-CNR), 00185 Rome, Italy (e-mail: gianluca.baldassarre@istc.cnr.it; stefano.nolfi@istc.cnr.it).

V. Trianni was with the Institut de Recherches Interdisciplinaires et de Développements en Intelligence Artificielle (IRIDIA), Université Libre de Bruxelles, 1050 Bruxelles, Belgium. He is now with the Laboratory of Autonomous Robotics and Artificial Life, Istituto di Scienze e Tecnologie della Cognizione, Consiglio Nazionale delle Ricerche (LARAL-ISTC-CNR), 00185 Rome, Italy (e-mail: vtrianni@ulb.ac.be).

M. Bonani and F. Mondada are with the Robotics Systems Laboratory, Ecole Polytechnique Fédérale de Lausanne (LSRO, EPFL-I2S-IPR), 1015 Lausanne, Switzerland (e-mail: michael.bonani@epfl.ch; francesco.mondada@epfl.ch).

M. Dorigo is with the Institut de Recherches Interdisciplinaires et de Développements en Intelligence Artificielle (IRIDIA), Université Libre de Bruxelles, 1050 Bruxelles, Belgium (e-mail: mdorigo@ulb.ac.be).

Digital Object Identifier 10.1109/TSMCB.2006.881299

is fully autonomous, the swarm as a whole can solve problems that a single robot cannot cope with because of physical constraints or limited behavioral capabilities. Swarm robotics emphasizes aspects such as decentralization of control, local and simple communication among robots, emergence of global behavior, and robustness [4]–[11]. Moreover, swarm robotics aims at exploiting self-organizing principles similar to those observed in social insects [12]–[14].

This paper focuses on a particular swarm robotic system (referred to as “swarm-bot”) that is composed of a number of individual robots (referred to as “s-bots”) that are assembled to each other through physical links [15], [16]. Each s-bot is provided with different types of sensors, motors, and connecting apparatuses that allow groups of s-bots to self-assemble and disassemble. A swarm-bot consisting of several connected s-bots should move as a whole and reconfigure its shape when needed. For example, it might have to change its shape in order to go through a narrow passage or overcome an obstacle [17]. Thus, swarm-bots combine the power of swarm intelligence, as they are based on the emergent collective intelligence of groups of robots, and the flexibility of self-reconfiguration as they might dynamically change their structure to match environmental variability.

There are different approaches that can be used to control such an artifact. In this paper, we aim at obtaining a completely decentralized system. Therefore, the behavior of the swarm-bot should not be defined by a central controller that establishes the actions to be performed by every single s-bot, nor should the s-bots act following a global template. The global behavior of the swarm-bot should rather be the result of a self-organizing process, that is, to emerge from the numerous interactions that take place among the s-bots and between the s-bots and the environment. Systems that feature self-organization are also characterized by other interesting properties, such as robustness, flexibility, and scalability [14]. Therefore, designing robotic systems that exploit self-organizing principles is highly desirable.

In this paper, we focus on a particular problem for the swarm-bot: coordinated motion. The s-bots are physically connected in a swarm-bot and have to coordinate their individual actions in order to move coherently. Coordinated motion is well studied in biology as it is present in many different animal species. Examples of this behavior can be seen in flocks of birds flying in a coordinated fashion or in 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-organizing behaviors. In Section V,

we review some important research work related to coordinated motion.

This paper shows how a coordinated motion of real physically linked robots can be achieved on the basis of simple and robust controllers that have access only to the local sensory information (similar results, obtained with simulated robots, are presented in [17] and [18]). Note that this paper focuses on the coordinated motion of swarm-bots in which s-bots are assembled since the beginning of the tests, while the complementary study on self-assembling has been reported elsewhere (see, for example, [19] and [20]). Swarm-bots' coordinated motion controllers are neural networks synthesized through artificial evolution [21]. This methodology proved to be very effective for the development of collective behaviors, but rarely were the obtained controllers tested on real robots (a noticeable exception is given in [11]). The main contribution of this paper consists in the demonstration that controllers evolved in simulation to coordinate physically assembled robots continue to exhibit high performance when downloaded and tested on real robots. The reason of such a successful transfer is mainly due to the properties of the evolved controllers, which were shaped by evolution in order to exploit the dynamical features of the system. This resulted in a simple and clever behavioral strategy at the individual level and in a robust self-organizing system at the collective level. To the best of our knowledge, this is the first work to date in which up to eight real physically assembled robots display coordinated behaviors clearly based on self-organizing principles (see Section V).

This paper is organized as follows. Section II presents the experimental setup, while Section III analyzes the functioning of the evolved controller. Section IV shows that the controller evolved in simulation produces a robust behavior when used to control real robots. Moreover, this section describes how the evolved neural controller generalizes its ability to produce coordinated motion in conditions that were never experienced during the evolutionary phase. In particular, this section shows that the controller evolved in simulation produces a robust behavior when used to control real robots. Finally, Section V reviews some literature related to this paper, and Section VI draws conclusions.

II. EXPERIMENTAL SETUP

This section describes the simulated and real s-bots' properties, the task, and the evolutionary method used to evolve the neural controller.

A. Robots and the Simulator

The s-bots used in this paper (shown in Fig. 1) have been developed within the "SWARM-BOTS" project [15], [16].¹ Each s-bot is composed of a turret and a chassis. The turret is a cylindrical body, with a diameter of 11.6 cm, equipped with a rigid gripper that allows the s-bot to connect to the perimeter of other s-bots. The chassis is a mobile base provided with two motors, each controlling a track and a teathed wheel. The

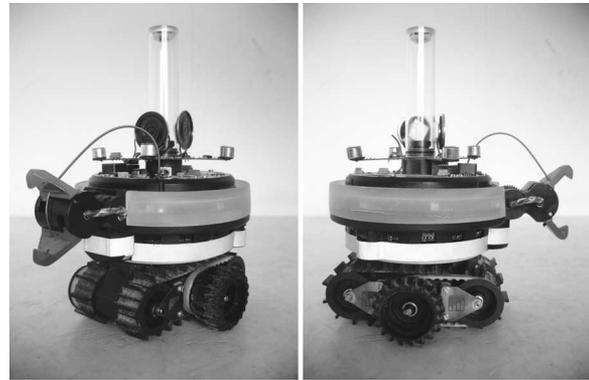


Fig. 1. S-bot. The bottom part (the chassis) includes the tracks and the teathed wheels and the four proximity sensors oriented toward the ground. The top part (the turret) includes the rigid gripper, one omnidirectional camera, four microphones, two speakers, 16 infrared proximity sensors, and a three-axis accelerometer. The traction sensor is placed between the turret and the chassis. The turret and the chassis can actively rotate with respect to each other. A flexible arm endowed with a gripper is also part of the s-bot, but it was neither used nor mounted on the s-bots used for the experiments presented in this paper (see [16] for more details).

turret and the chassis can actively rotate with respect to each other through an independent motor. Relative rotation is limited to $\pm 180^\circ$ due to power and control cables connecting the two parts.

S-bots are provided with several sensors, such as infrared proximity sensors, microphones, an omnidirectional camera, and many others (for more details, see [16]). However, in this paper, we used only the traction sensor, a sensor that detects the direction and the intensity of the pulling force that the turret exerts on the chassis. The sensor is composed of two portions: one connected to the turret and the other one to the chassis (see Fig. 2). The two parts can translate with respect to each other along two orthogonal horizontal axes and consequently can deform four thin iron plates that connect them. This deformation, which is proportional to the intensity of the traction force, is measured along the two axes by eight strain gauges placed on the plates. The two values so obtained are the x and y components of the traction force, which are measured with respect to a reference-frame integral with the chassis. The two orthogonal components are used to compute the intensity and direction of the traction force.²

It should be noted that in swarm-bots formed by two or more assembled s-bots, the body of each s-bot physically integrates the forces resulting from the traction and thrust that other s-bots exert on it. The traction sensor, by detecting the resultant of these forces, provides compact information on the mismatch between the s-bot's movement and the movement of the rest of the group. The perceived traction thus constitutes an implicit form of communication (cf. [14]) that, as we will see in Section III, can be exploited by s-bots to produce coordinated movements.

A simulator based on a 3-D rigid body dynamics simulation engine was developed to synthesize the robot controller through

¹For more information, see also the project website at <http://www.swarm-bots.org>.

²In this paper, the direction of the traction has been encoded from 0° to 360° , where 0° and 180° correspond to the backward and forward directions of motion of the chassis, respectively, while 90° and 270° correspond to the traction coming from the left- and right-hand side of the chassis, respectively.

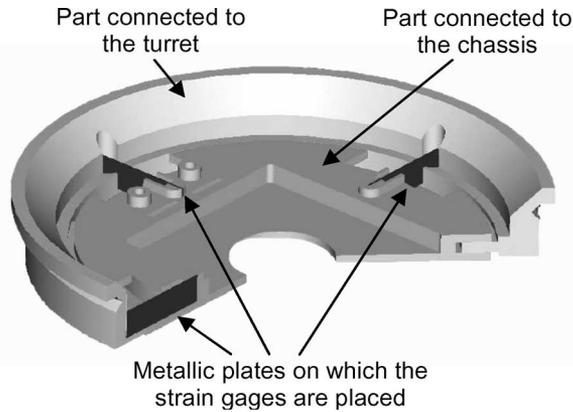
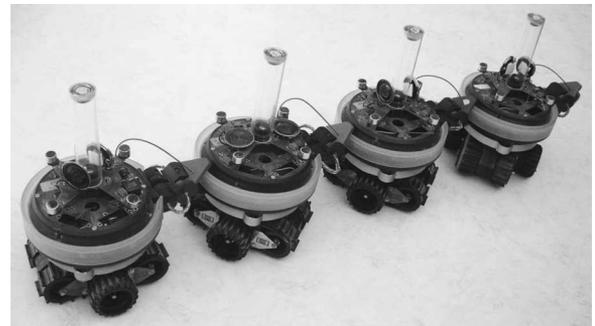


Fig. 2. Structure of the traction sensor (see text for details).

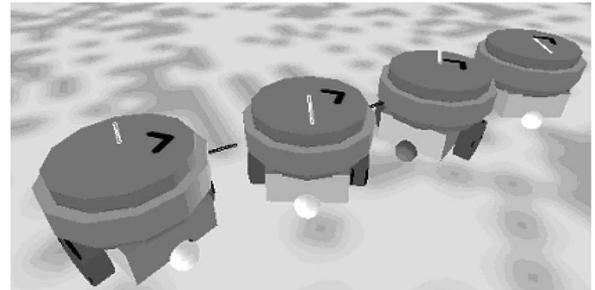
an evolutionary technique (see Section II-C). In fact, embedding the evolutionary process in the real robots would have been extremely time demanding: one evolutionary run would have taken about eight days if carried out with the real robots. Moreover, even running part of the evolution directly on the real robots (e.g., to “refine” the controller in real robots, cf. [22]) was not viable given the prototype stage of development of the robots. Also, many other factors, such as power issues and reinitialization difficulties, made evolution with real robots impractical in our particular case.

The simulator of the s-bots was based on a very simplified model in order to increase the speed of the simulations. This model preserves only the features of the real s-bots that were considered important for the experiments to be performed. The simulated s-bot consists of a cylindrical turret connected to a chassis by a motorized hinge joint. In the basic simulation model used for the evolution of the coordinated-motion behavior, the two bodies can rotate without limits. However, a second version of the model simulates the limit for which the turret can only rotate $\pm 180^\circ$ with respect to the chassis, as in the real s-bot. The latter model was used for comparing the results obtained in simulation with those obtained with the real s-bots (see Section IV). The chassis is modeled as a parallelepiped to which four spherical wheels are connected. The lateral wheels are connected to the chassis by motorized joints. Friction is modeled on the basis of the Coulomb friction model (the friction coefficient was set to 0.6). This setup implies that the s-bot’s wheels slip if motion is blocked by obstacles or by other connected s-bots. The front and back wheels are passive and can rotate in any direction. The gripper is not present in the model, and connections between two s-bots are simulated by creating a rigid joint between the two bodies.

The traction sensor is simulated measuring the horizontal components of the force acting on the hinge joint that connects the turret to the chassis. This force is computed at each cycle by the dynamic simulation engine and is therefore always available. The maximum force that the sensor can perceive was measured on the real s-bots and accordingly set in the simulation. Noise is added to the two horizontal components of the traction force by adding a value randomly selected with a uniform distribution within the range $[-5\%, +5\%]$. Note that, due to the high number of variables that influence the perception of



(a)



(b)

Fig. 3. (a) Four real s-bots forming a linear swarm-bot. (b) Four simulated s-bots forming the same linear structure. The cylinders represent the turret, while the chassis is shaped as a parallelepiped. The arrow on the cylinders indicates the orientation of the turret. The wheels are displayed as cylinders (motorized wheels) and spheres (passive wheels, which have different colors, dark and light gray, to allow distinguishing, respectively, the two chassis’ fronts). The black segment between the turrets of the two robots represents a physical link between them (gripper). The white line above each robot’s turret, which goes from the turret’s center toward its perimeter, indicates the direction of the traction force and, with its length, its intensity.

traction (i.e., number of robots involved, friction parameters, hysteresis, and interrobot variability), a precise characterization of the traction sensor was not feasible. Moreover, it would have been extremely difficult to use samples taken from the real s-bots in order to resort to a sampling technique [23]: the aforementioned procedure was the only viable option for simulating the traction sensor. For more details on the simulator and for a description of more detailed simulation models not used in the experiments reported in this paper, see [16].

B. Task

A swarm-bot can efficiently move only if the chassis of the assembled s-bots have the same orientation. As a consequence, the s-bots should be capable of negotiating a common direction of movement and then compensating possible misalignments that originate during motion.

The experiments presented in this paper study a group of s-bots that remain always connected in a swarm-bot formation (see Fig. 3). At the beginning of a trial, the s-bots have their chassis oriented in random directions. Their goal is to choose a common direction of motion on the basis of only the information provided by the traction sensor and then to move as far as possible from the starting position along such direction. Notice that this task is more difficult than it might appear at first sight.

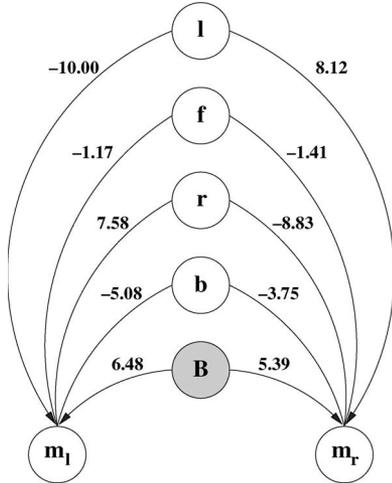


Fig. 4. Weights of the controller synthesized in the 30th run of the simulation. The sensory neurons associated with the left, front, right, and back traction-sensor readings are labeled as “*l*,” “*f*,” “*r*,” and “*b*,” respectively. “*B*” indicates the bias neuron, while m_l and m_r indicate, respectively, the left and right motor neuron.

First, the group is not driven by a centralized controller (i.e., the control is distributed) nor the s-bots can directly communicate or coordinate on the basis of synchronizing signals. Moreover, the s-bots cannot use any type of landmark in the environment, such as light sources, or exploit predefined hierarchies between them to coordinate (i.e., there are no “leader robots” that decide and communicate to the other robots the direction of motion of the whole group). Finally, the s-bots do not have a predefined trajectory to follow, nor do they have information about their relative positions or about the structure of the swarm-bot in which they are assembled. As a consequence, the common direction of motion of the group should emerge as the result of a self-organizing process based on local interactions perceived by the robots through the traction sensors. The problem of designing a controller capable of producing such a self-organized coordination was tackled using neural networks synthesized by artificial evolution, as illustrated in detail in the following section.

C. Neural Controllers and the Evolutionary Algorithm

In the experiments reported here, artificial evolution is used to synthesize the connection weights of simple neural controllers with fixed architecture (see Fig. 4). The controller of each s-bot consists of a neural network with four sensory neurons (plus a bias unit) directly connected to two motor neurons. The sensory neurons are simple relay units while the output neurons are sigmoid units whose activation is computed as follows:

$$y_j = \sigma \left(\sum_i w_{ji} x_i \right) \quad \sigma(z) = \frac{1}{1 + e^{-z}} \quad (1)$$

where x_i is the activation of the i th input unit, including the bias, y_j is the activation of the j th output unit, w_{ji} is the weight of the connection between the input neuron i and the output neuron j , and $\sigma(z)$ is the sigmoid function.

The sensory neurons encode the intensity of traction along the four directions corresponding to the direction of the semi-axes of the chassis’ reference frame (i.e., front f , back b , left l , and right r , see also Fig. 4). In particular, the sensory neurons are activated as follows:

$$\begin{aligned} r &= F_x & l &= 0, & \text{iff } F_x &\geq 0 \\ r &= 0 & l &= -F_x, & \text{iff } F_x < 0 \\ f &= F_y & b &= 0, & \text{iff } F_y &\geq 0 \\ f &= 0 & b &= -F_y, & \text{iff } F_y < 0 \end{aligned} \quad (2)$$

where F_x and F_y are the x and y components of the traction force. The bias neuron is clamped to one. The activation state of the two motor neurons is scaled onto the range $[-\omega_M, +\omega_M]$, where ω_M is the maximum angular speed of the wheels ($\omega_M \approx 3.375$ rad/s in simulated s-bots and $\omega_M \approx 3.5$ rad/s in the real s-bots: these settings allowed obtaining the same speed for simulated and real robots). The desired speed of the turret–chassis motor is set equal to the difference between the desired speed of the left and right wheels times a constant $k = r_w/d_w$, where r_w is the radius of the wheels, and d_w is the distance between the two wheels. This setting produces a movement of the turret with respect to the chassis that counterbalances the rotation produced by the wheels’ motion. In this way, the turret–chassis motor actively contributes to the rotation of the chassis by anchoring on the connected robots, especially in those situations in which one or both wheels partially or totally lose contact with the ground.

The s-bots are connected in a linear formation as shown in Fig. 3(b). The evolutionary algorithm is based on a population of 100 genotypes, which are randomly generated. This population of genotypes encodes the connection weights of 100 neural controllers. Each connection weight is represented with a ten-bit binary code mapped onto a real number in a range $[-10, +10]$. For each genotype, four identical copies of the resulting neural-network controllers are used: one for each s-bot (this implies that the s-bots forming the swarm-bot have homogenous controllers). The “fitness” of the genotype is computed as the average performance of the swarm-bot over five different trials. Each trial lasts $T = 150$ cycles, each corresponding to 100 ms of real time for a total of 15 simulated seconds. At the beginning of each trial, a random orientation of the chassis is assigned to each s-bot. The ability of a swarm-bot to display a coordinated motion is evaluated by computing the average over five trials of the distance D covered by the group. In particular, in each trial tr , the distance covered by the group is obtained by measuring the Euclidean distance between the position of the center of mass of the swarm-bot at the beginning and at the end of the test

$$D = \frac{1}{5} \sum_{tr=1}^5 \frac{\|\mathbf{c}_{tr}(T) - \mathbf{c}_{tr}(0)\|}{D_M(T)} \quad (3)$$

where $\mathbf{c}_{tr}(t)$ is the vector of coordinates of the group’s center of mass at time t , and $D_M(t)$ is the maximum distance that can be covered by an s-bot in t simulation cycles. Notice that this way of computing the “fitness” of the group is sufficient to obtain a coordinated-motion behavior. In fact, it rewards swarm-bots that maximize the covered distance and, therefore, their motion

speed. As a consequence, the s-bots should minimize the time required to align their chassis, move at maximum speed once coordinated, and reduce instabilities and noise disturbances that might impair the motion of the group while moving. This fitness measure promotes controllers that result in efficient coordination, as confirmed by the analysis of the evolved behavior performed in Section III.

Once the fitness of every genotype of the population has been computed, the 20 best individuals are selected for reproduction. Each genotype is reproduced five times, applying a mutation with 3% probability of replacing a bit with a new randomly generated value (crossover was not used due to the simplicity of the controller). The evolutionary process, which is a run in simulation, lasts 100 generations and is replicated 30 times starting with different initial randomly generated genotypes.

III. RESULTS

All the 30 evolutionary runs successfully synthesized controllers that produced coordinated motion in the linear swarm-bot. The obtained results are described in detail in Section III-A. Section III-B describes how the problem related to the rotational limit of the turret/chassis degree of freedom was solved. The solution to this problem was important in testing the evolved controllers on the real robots, as described in Section IV.

A. Results in Simulation

The controllers evolved in simulation allow the s-bots to coordinate by negotiating a common direction of movement and to keep moving along such direction by compensating small misalignments arising during the movement (see Fig. 5). Direct observation of the evolved behavioral strategies shows that at the beginning of each trial, the s-bots try to pull or push the rest of the group in the direction of motion that they initially have. This disordered motion results in traction forces that are exploited for coordination as the s-bots tend to orient their chassis in the direction of the perceived traction, which roughly corresponds to the “average” direction of motion of the group. This allows the s-bots to rapidly converge toward a common direction and to maintain it.

All the 30 controllers evolved in the different replications of the evolutionary process present similar dynamics: In all trials, the s-bots converge to a common direction of motion in a very fast and effective way. As shown in Fig. 5, this common direction of motion varies across trials. In fact, the direction of motion of the group is not *a priori* defined but rather emerges as a result of the coordination phase and depends on the initial random orientations of the s-bots' chassis.

By testing the best neural controller of the last generation of each evolutionary run for 100 trials, it was observed that the performance varies in the range [0.81, 0.91], which is not far from the theoretical maximum (corresponding to 1.0) that can be achieved only by a single s-bot moving at full speed in a fixed direction. Notice that the maximum performance cannot be reached in practice by a swarm-bot, since assembled s-bots can move at maximum speed only once they have achieved

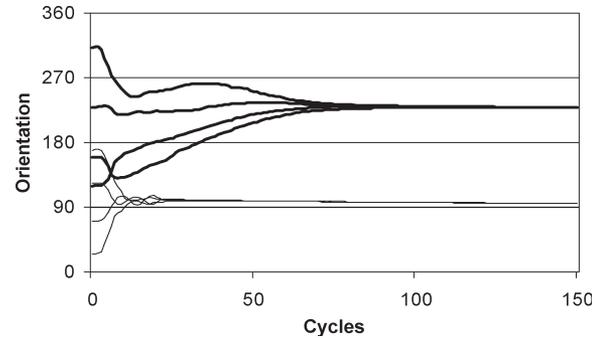


Fig. 5. Absolute orientation of the chassis of four s-bots forming a linear structure in two trials lasting 150 cycles each (thick and thin lines, respectively). At the beginning of each trial, the s-bots start moving with randomly assigned orientations, as can be seen by the different starting points of the curves. As time elapses, the robots achieve coordination and converge to the same direction of motion, as shown by the curves' overlap at the end of the graph. Notice how the final emergent direction of motion of the swarm-bot is different in the two trials.

coordination. In the rest of this paper, the controller synthesized by the 30th evolutionary run is used because it resulted to have the best performance. Fig. 4 shows both the architecture of this controller and the weights of each connection between the input and the output neurons, as generated by the evolutionary process.

In order to understand the functioning of the controller at the individual level, the activation of the motor units of an s-bot was measured in correspondence to a traction force whose angle and intensity were systematically varied. The results reported in Fig. 6 indicate the following.

- 1) Whenever the traction intensity is low or when the traction comes from the front (i.e. around 180°), the s-bot moves forward at maximum speed (see the portions of Fig. 6 indicated by number 1). These conditions take place, respectively, when the s-bot's chassis is oriented toward the same direction in which the other s-bots are pulling/pushing it or when all s-bots' chassis are aligned.
- 2) When traction comes from the left- or the right-hand side (i.e., around 90° or 270° , respectively), the s-bot turns toward the direction of traction (see the portions of Fig. 6 indicated by number 2). This condition takes place when there is a significant mismatch between the motion's direction of the s-bot and the average direction of motion of the group.
- 3) When traction comes from the rear (i.e., around 0°), the s-bot moves forward at maximum speed independently of the traction intensity (see the portions of Fig. 6 indicated by number 3). Notice that this is an unstable condition: As soon as the angle of traction differs from 0° , for example due to noise, the s-bot rotates its chassis following the rules specified in point 2. This type of condition is normally caused by the movement of the s-bot itself, whenever the resultant of the forces produced by the other s-bots in the group tends to be null.

In other words, at the individual level, each s-bot exhibits two tendencies. One consists in following the rest of the group (e.g., when the perceived traction comes from the left- or right-hand side) and the other consists in persevering in moving straight

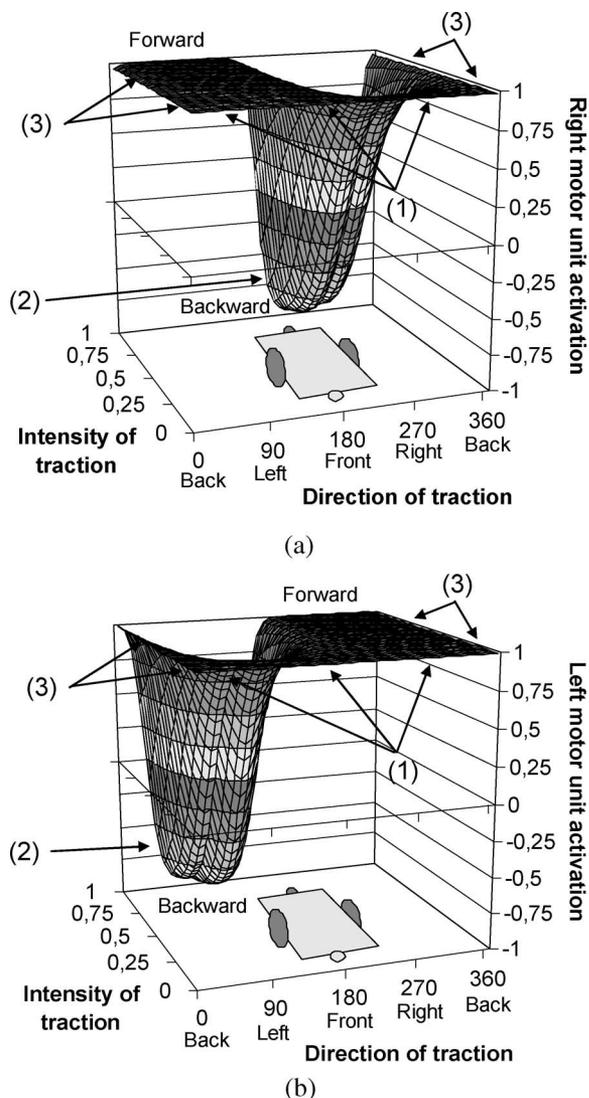


Fig. 6. Motor commands issued by (a) the left and (b) right motor units mapped onto a $[-1, 1]$ interval (-1 and $+1$ correspond to the maximum backward and forward speeds, respectively) of one of the best evolved neural controllers in correspondence to the traction forces having different directions and intensities (see text for the explanation of the numbers in round brackets).

(e.g., when the perceived traction comes from the rear or from the front, or has a low intensity). The effects of the individual behavior at the group level can be described as follows. At the beginning of each test, all s-bots perceive traction forces with low intensity, and so they move forward at maximum speed (according to point 1). The different traction forces generated by these movements are physically summed up by the turret of each robot. This causes a unique force to emerge at the group level, which has a direction that characterizes the movement of the whole group. The s-bots that have small misalignments with respect to the average group's motion direction perceive traction forces from the rear, and so they tend to persevere in their motion (according to point 3). In so doing—and this has a very important role for coordination—they continue to generate a traction signal in the same direction, which is perceived by the rest of the group. In contrast, the s-bots that have large misalignments with respect to the average group's direction of motion perceive traction from the left- or right-hand side,

and so they tend to turn so as to follow the rest of the group (according to point 2). Overall, these behaviors quickly lead the whole group of s-bots to converge toward the same direction of motion (see [24] for a more detailed quantitative analysis of the self-organizing principles at work in these processes).

As will be shown in the rest of this paper, this simple behavioral strategy is very effective and robust. In some cases, however, the same strategy does not lead the s-bots to converge toward a common direction of motion but rather to a rotational dynamic equilibrium in which all s-bots move around the center of mass of the swarm-bot. This rotational equilibrium is stable since, while turning in circle, the s-bots perceive a traction force toward the group's center that keeps them moving by slightly turning toward it. This rotational equilibrium is never observed in the experimental conditions used to evolve the controller, involving four simulated s-bots forming a linear structure, but only in the generalization tests performed with the real robots in different situations (see Section IV).

B. Coping With the Limits of the Turret–Chassis Degree of Freedom

As previously mentioned, the chassis of the s-bots can rotate only 180° clockwise or counterclockwise with respect to the turret, due to the cables connecting the two parts. This implies that in order to coordinate with the other s-bots, an individual s-bot cannot simply turn its chassis toward the direction of traction. In fact, if the rotational limit is located between the current orientation of the s-bot's chassis and the direction of traction, the s-bot should turn in the opposite direction (up to 360°) in order to reach the desired orientation.

Rather than introducing the limit in the simulation model and asking evolution to solve the problem, we designed a solution that consists in inverting the front of motion when the limit on the turret–chassis degree of freedom is reached (this solution was proposed for the first time, and tested in simulation, in [17]). This solution exploits the fact that s-bots have two equivalent fronts of motion. In fact, the chassis is symmetric with respect to the wheel's axis; the motorized wheels can turn in both directions; and the sensors are homogeneously distributed. As a consequence, the same behavior described in the previous section continues to work properly when the two fronts of motion are “swapped.” Specifically, the direction of motion of the s-bot can be easily inverted (forward with backward and vice versa), provided that the encoding of the sensor and motor neurons is properly modified. More in particular, a front inversion can be implemented as follows: 1) The motor commands are swapped (left with right and right with left), and their signs are inverted and 2) the encoding of the sensory neurons that determines which are the front and rear input units and which are the left and right input units, is rotated 180° along the perimeter of the robot.

The solution to the rotational limit consists in triggering a front inversion each time the turret exceeds the rotational limit both while turning clockwise or counterclockwise. The effect of the front inversion at the level of the single robot is illustrated in Fig. 7. In the example shown in the figure, the robot is initially moving by using the first front. Since it perceives

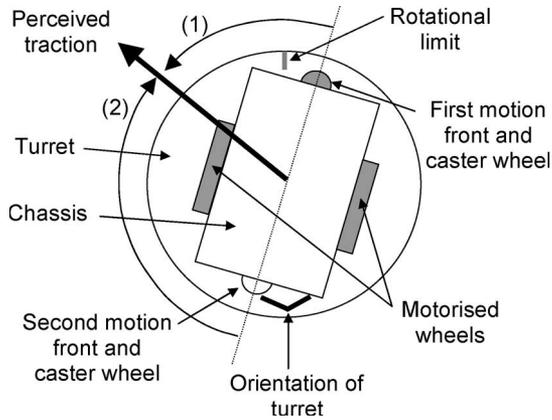


Fig. 7. Schematic representation of the effect of the front inversion from the point of view of a single robot. The bold arrow indicates the direction of the traction perceived by the s-bot. The gray caster wheel cannot pass the rotational limit. The arrows “1” and “2” indicate the direction in which the chassis turns before and after the front inversion, respectively. In this case, the inversion was from the first to the second front (see text).

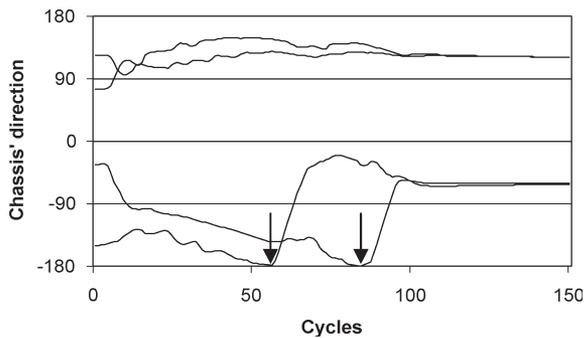


Fig. 8. Absolute orientations of the chassis of four s-bots (y axis) during a trial lasting 150 cycles (x axis). The arrows indicate the cycles in which two s-bots reach the rotational limit and invert their front of motion. During the last phase, the two s-bots that never changed their front still move by using their first front, while the other two s-bots use the second front.

the traction from its left-hand side, the robot starts turning its chassis counterclockwise (along the direction indicated by the arrow “1” in the figure). While turning, the chassis reaches the rotational limit, and the front inversion is triggered. At this point, the controller perceives the traction from the right-hand side, and therefore, the chassis starts turning clockwise (along the direction indicated by the arrow “2” in the figure). Consequently, the robot can successfully align its current front (the second front in this case) to the direction of traction without exceeding the rotational limit.

The effect of the front inversion at the level of the swarm-bot is shown in Fig. 8, which indicates the absolute orientation (with respect to the first front) of the chassis of the four s-bots forming a linear structure provided with the rotational limit and the front-inversion mechanism. Initially, the s-bots, all having random orientations, use the first front. Between cycles 50 and 100, two s-bots reach the rotational limit and invert their front. Finally, from about cycle 100 onward, the four s-bots converge to the same direction of movement. Notice how two robots, after converging, use the first front and have an absolute orientation of the chassis of about 120° , while two robots use the second front and have an orientation of about

-60° . The result is that all s-bots move in the same absolute direction in the last phase of the trial.

The front-inversion mechanism actually solves the problem introduced by the rotational limit, but it could also affect the performance of the swarm-bot in the coordinated-motion task. We measured the effects of this solution measuring the average distance covered by a swarm-bot over 20 trials lasting 25 s each. We noticed only a slight decrease with respect to the baseline performance, that is 8% of the covered distance (see Fig. 10). This indicates that the front-inversion mechanism is a viable solution to cope with the rotational limit. This is an important result in view of testing the evolved controllers with real robots because in this condition the constraint imposed by the rotational limit cannot be neglected.

IV. TESTING WITH REAL ROBOTS

The introduction of the front-inversion mechanism provides the controller evolved in simulation with all the required characteristics to be directly transferred to the real s-bots. We therefore tested the functionality of the evolved behavior in reality comparing the obtained performance with the results of the simulations.

In all the tests performed in this section, s-bots are provided with the rotational limit of the turret–chassis motor and with the front-inversion mechanism. The s-bots always start connected to each other with randomly assigned chassis’ orientations. Each experimental condition is tested for 20 trials, each lasting 25 s (250 cycles).

We initially tested the functionality of the evolved neural controller in the experimental conditions identical to those used during evolution (see Section IV-A). Afterward, we studied the ability of the controller to generalize to different situations that were never met during the evolutionary process: rough terrain and varying size and shapes of the swarm-bots. Then, we tested the coordination capabilities of the controller when using semirigid connections between s-bots (implemented by a slightly loose gripping) or indirect connections between them (that is, robots attached to an object to be transported). The good performance recorded in all these new conditions suggests that the evolved controller is very robust and flexible.

A. Testing the Controller Evolved in Simulation on Real S-Bots

We tested the best controller evolved in simulation using four real s-bots forming a linear structure. The results show that the controller allows the real s-bots to coordinate without the need of any adjustment and despite significant differences from the simplified simulation model previously described. Indeed, as shown in Fig. 9, the simulated and real s-bots display a qualitatively similar behavior.

Quantitatively, on average, the performance of the best controller evolved in simulation decreases 23% when tested with the real s-bots (see the second and third histogram bars of Fig. 10 and the first two columns of Table I). The data shown in Table I also indicate that the swarm-bot never fell into the rotational equilibrium, either in tests with simulations nor

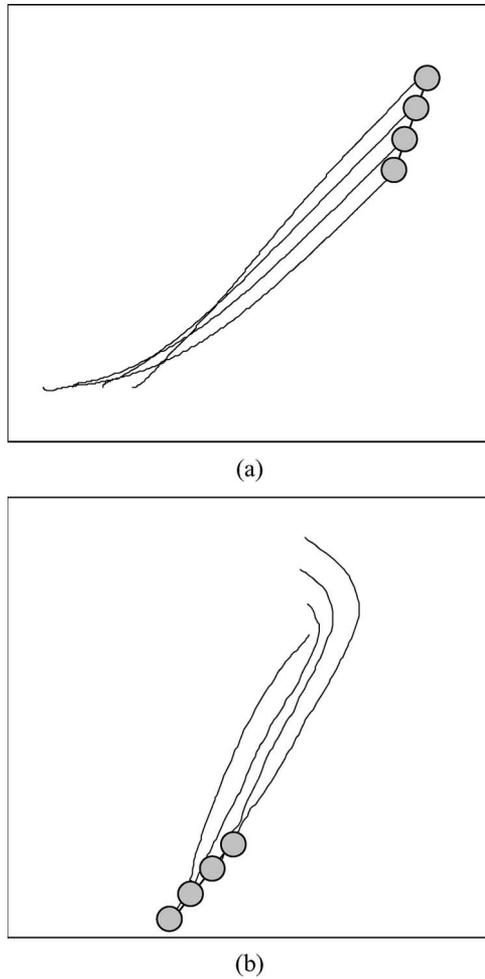


Fig. 9. Trajectory of (a) four simulated and (b) four real s-bots forming a linear swarm-bot in a coordinated-motion test lasting 15 s. The gray circles indicate the final position of the s-bots. In the case of real s-bots, trajectories have been automatically extracted from a video obtained by recording the behavior of the real s-bots from a camera mounted on the ceiling.

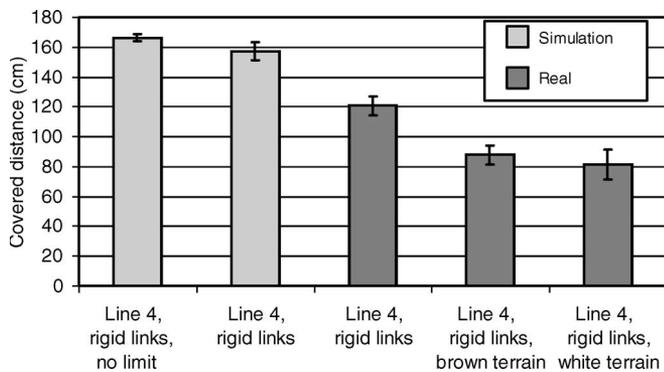


Fig. 10. Performance of the best evolved controller in simulation and reality (average and standard error of the distance covered in 20 trials, each lasting 25 s). Light and dark gray bars represent the tests carried out with the simulated and real s-bots, respectively. Labels indicate the experimental conditions. "Line 4" indicates tests involving four s-bots forming a linear structure; "rigid links" indicate rigid connections between s-bots; "no limit" indicates tests performed without the introduction of the rotational limit and of the front-inversion mechanism; "brown terrain" and "white terrain" indicate two different rough terrain conditions (see text).

TABLE I
PERFORMANCE OF THE BEST EVOLVED CONTROLLER TESTED IN SIMULATION AND REALITY. TESTS INVOLVE FOUR S-BOTS FORMING A LINEAR STRUCTURE. THE FIRST TWO COLUMNS INDICATE THE PERFORMANCE ON FLAT TERRAIN IN THE CASE OF SIMULATED AND REAL S-BOTS, RESPECTIVELY. THE LAST TWO COLUMNS INDICATE THE PERFORMANCE OF REAL S-BOTS ON BROWN AND WHITE ROUGH TERRAIN, RESPECTIVELY (SEE TEXT). THE SIX ROWS INDICATE IN ORDER: THE AVERAGE PERFORMANCE OVER 20 TRIALS, THE STANDARD DEVIATION, THE STANDARD ERROR, THE RATIO OF THE PERFORMANCE WITH RESPECT TO THE THEORETICAL MAXIMUM, THE RATIO OF THE PERFORMANCE WITH RESPECT TO THE CORRESPONDING SIMULATED TEST, AND THE NUMBER OF TRIALS (OUT OF 20) IN WHICH THE SWARM-BOTS DID NOT MANAGE TO PERFECTLY COORDINATE

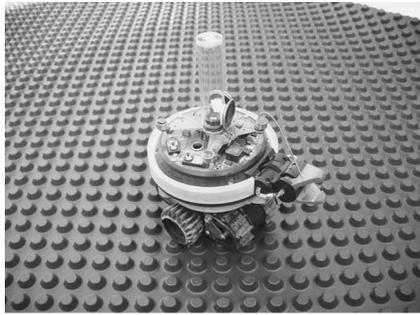
	Line 4, rigid links, flat terrain		Line 4, rigid links, rough terrain	
	Simulation	Real	Brown	White
Avg. perf.	156.96	120.85	87.75	81.25
Std. dev.	28.39	29.53	43.95	39.45
Std. err.	6.35	6.60	9.82	8.82
Ratio with th. max.	0.85	0.65	0.47	0.44
Ratio with sim.	1.00	0.77	0.56	0.52
Partial coord.	0	0	4	6

in those with real robots. The lower performance of the real swarm-bot with respect to the simulated swarm-bot is due to the longer time required by the real s-bots to coordinate. This is caused by many factors, among which the fact that tracks and teathed wheels of the real s-bots sometimes get stuck during the initial coordination phase, which is due to a slight bending of the structure that caused an excessive thrust on the tracks. This leads to a suboptimal motion of the s-bots, for example, while turning on the spot. However, coordination is always achieved, and the s-bots always move away from the initial position. This result proves that the controller evolved in simulation can effectively produce coordinated motion when tested in real s-bots, notwithstanding the fact that the whole process takes some more time with respect to simulation.

B. Testing the Controller on Rough Terrain

The evolved controller is also able to produce coordinated movements on rough terrain. Fig. 10 and Table I show the performance obtained by the real s-bots placed on two types of terrain. The brown rough terrain is a very regular surface made of brown plastic isolation foils. This terrain remains mostly flat, but it is impossible to access for most standard wheeled robots. Only robots with tracks like the s-bot can move on it. The plastic is composed of a grid of cones, which are spaced 2.1 cm apart. Each cone is 1.2 cm large and 0.7 cm high [see Fig. 11(a)]. The white rough terrain is an irregular surface made of stonelike plaster bricks. The bricks measure 13 × 28 cm, and their heights range from 0.9 to 2.1 cm [see Fig. 11(b)].

With the exception of a few cases in which coordination is only partially achieved, the performance of the swarm-bot on the rough terrains is comparable with what is achieved on the flat terrain. However, in these experimental conditions, we observed a decrease of the performance, which is mainly due to a more difficult gripping of the tracks and teathed wheels on the irregular surface. In fact, the roughness leads to very noisy signals perceived by the traction sensors. As a consequence, the swarm-bots in some cases do not reach a complete



(a)



(b)

Fig. 11. Two types of rough terrain used to test the robustness of the controller. (a) Very regular rough terrain made of brown plastic isolation foils. (b) Irregular rough terrain made of white plaster bricks that look like rough stones.

coordination since the s-bots have similar but different orientations. In these situations, the swarm-bots move in large circles, sometimes returning to the initial position, therefore scoring a low performance.

C. Testing With Swarm-Bots Consisting of a Larger Number of Assembled S-Bots

The best evolved controller was tested with the linear swarm-bots composed of six s-bots. The results showed that larger swarm-bots preserve their ability to produce coordinated movements both in simulation and in reality. As shown in Fig. 12 and Table II, the performance in the new experimental condition is only 10% and 8% lower than the one measured with the swarm-bots formed by four s-bots, respectively, in tests with the simulated and real s-bots. The performance of the experiments performed with the six real s-bots is 21% lower than the corresponding simulated experiments in line with the results presented in Section IV-A. Moreover, in all cases, swarm-bots never fall into the rotational equilibrium. This test suggests that the evolved controller produces a behavior that scales very well with the number of individuals forming the group both in simulated and real robots.

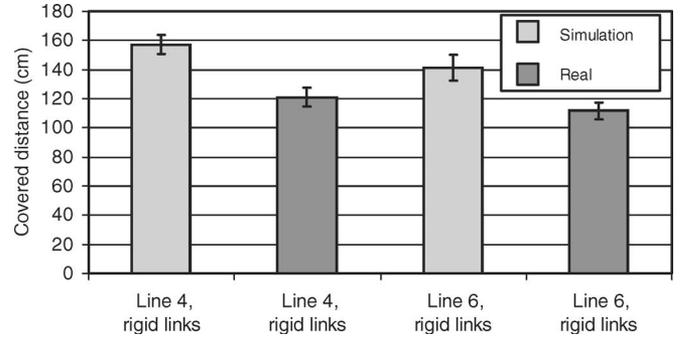


Fig. 12. Performance of the best evolved controller in simulated and real swarm-bots formed by a different number of s-bots (average and standard error of the distance covered in 20 trials, each lasting 25 s) (see the caption of Fig. 10 for an explanation of the figure). Additionally, “line 6” indicates tests involving six s-bots forming a linear structure.

TABLE II
PERFORMANCE OF THE BEST EVOLVED CONTROLLER TESTED IN SIMULATION AND REALITY. COMPARISON BETWEEN LINEAR STRUCTURES INVOLVING FOUR AND SIX S-BOTS, RESPECTIVELY (SEE CAPTION OF TABLE I FOR MORE DETAILS)

	Line 4, rigid links, flat terrain		Line 6, rigid links, flat terrain	
	Simulation	Real	Simulation	Real
Avg. perf.	156.96	120.85	141.03	111.65
Std. dev.	28.39	29.53	39.36	26.05
Std. err.	6.35	6.60	8.80	5.82
Ratio with th. max.	0.85	0.65	0.76	0.60
Ratio with sim.	1.00	0.77	1.00	0.79
Rot. equil.	0	0	0	0

D. Testing With Swarm-Bots Having Different Shapes

The best controller evolved in simulation was tested varying the shape and the size of the swarm-bot. In particular, we tested the swarm-bots composed of four s-bots forming a square structure and swarm-bots composed of eight s-bots forming a “star” shape (see Fig. 13). The results show that the controller displays an ability to produce coordinated movements independently of the swarm-bot’s shape, although the tests that use real s-bots show a higher drop in performance. As shown in Fig. 14, in simulation the performance of square and “star” swarm-bots is not very different from the performance of the linear swarm-bot composed of four s-bots. Comparing the data reported in Tables I and III, the performance of the simulated swarm-bots in square and “star” formations is, respectively, 13% and 17% lower than for a linear swarm-bot. The corresponding experiments performed with the real swarm-bots present a performance drop of 18% and 35% with respect to the real swarm-bots having a linear structure. These higher decrements of performance of the real robots are due to a higher chance of falling in the rotational equilibrium (up to seven times in the case of the “star” formation) and, to a minor extent, to an increased difficulty in converging toward a common direction of motion and in maintaining it (see also Section IV-G). With respect to the rotational equilibrium, we observed that the chance of falling in it is higher in swarm-bots having shapes that tend to be central symmetrical. Additionally, increasing the size of the swarm-bots leads to a slower coordination. This not only lowers the performance, but also likely increments the probability that the group falls in the rotational equilibrium. As a



(a)



(b)

Fig. 13. (a) Swarm-bot composed of four s-bots forming a square shape. (b) Swarm-bot composed of eight s-bots forming a “star” shape.

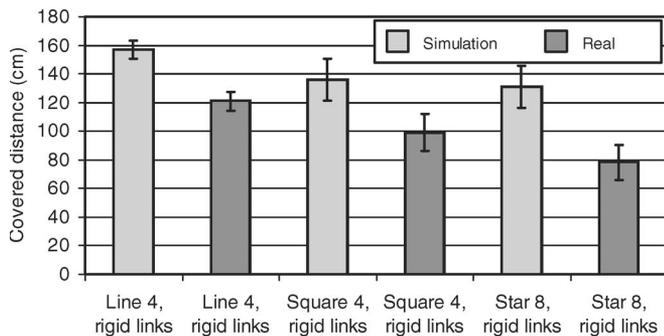


Fig. 14. Performance of the best evolved controller in simulation and reality (average and standard error of the distance covered in 20 trials, each lasting 25 s) (see the caption of Fig. 10 for a detailed explanation of the figure). Additionally, “square 4” indicates tests involving four s-bots forming a square shape; “star 8” indicates tests involving eight s-bots forming a “star” shape.

consequence, the performance of the square and “star” formations in reality is 27% and 40% lower than the corresponding simulated ones (see Table III).

E. Testing With Swarm-Bots Assembled Through Semirigid Links

The experiments presented in this section are conceived to test the generalization capability with respect to different types of links between s-bots. The neural controllers have been evolved with a linear swarm-bot composed of four s-bots connected through rigid links. Here, we test the same controller

TABLE III
PERFORMANCE OF THE BEST EVOLVED CONTROLLER TESTED IN SIMULATION AND REALITY. COMPARISON BETWEEN A SQUARE SWARM-BOT INVOLVING FOUR S-BOTS AND A “STAR” SWARM-BOT INVOLVING EIGHT S-BOTS (SEE CAPTION OF TABLE I FOR MORE DETAILS)

	Square 4, rigid links		Star 8, rigid links	
	Simulation	Real	Simulation	Real
Avg. perf.	136.02	99.00	131.05	78.10
Std. dev.	65.44	57.22	64.96	55.15
Std. err.	14.63	12.79	14.53	12.33
Ratio with th. max.	0.74	0.53	0.71	0.42
Ratio with sim.	1.00	0.73	1.00	0.60
Rot. equil.	4	5	4	7

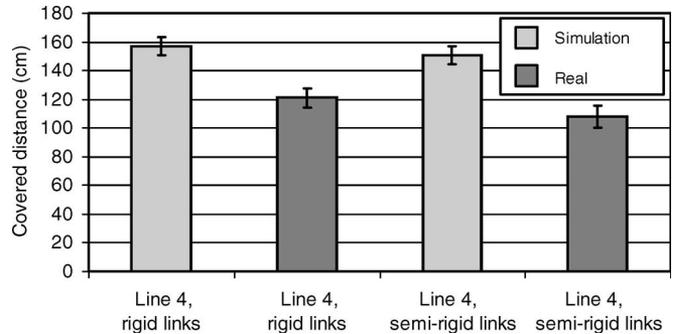


Fig. 15. Performance of the best evolved controller in simulation and reality (average and standard error of the distance covered in 20 trials, each lasting 25 s) (see the caption of Fig. 10 for a detailed explanation of the figure). Additionally, “semirigid links” indicate tests involving s-bots connected through slightly opened grippers.

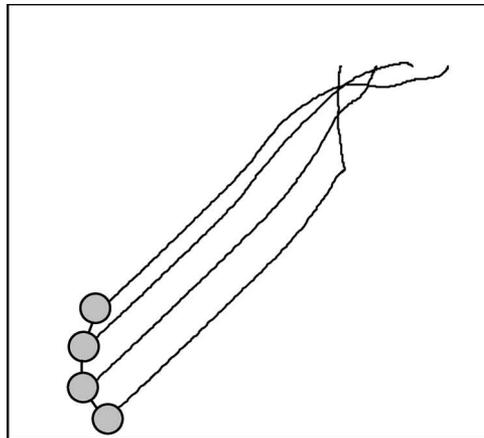
with s-bots connected through “semirigid” links. Contrary to the other experiments illustrated in this paper, in the case of semirigid links, the gripper is not completely closed, and the assembled s-bots are partially free to move with respect to each other. In fact, a partially open gripper can slide around the turret perimeter and can partially rotate by pivoting on the gripping point.

One interesting aspect of semirigid links is that they potentially allow swarm-bots to dynamically rearrange their shape in order to better adapt to the environment. Indeed, experiments conducted in simulation show how the swarm-bots assembled through semirigid links are able to dynamically rearrange their shape in order to pass through narrow passages and avoid falling into holes [17], [25]. The way in which the torque produced by the motors controlling the wheels and the turret of each individual s-bot affect the traction perceived by other s-bots, however, significantly differs in the case of rigid and semirigid links. Whereas in the case of rigid links the forces produced by the motors and the collisions directly affect the traction perceived by other s-bots, in the case of semirigid links these forces might also affect the shape of the swarm-bot. As a consequence, the traction forces are transmitted only in part when using semirigid links.

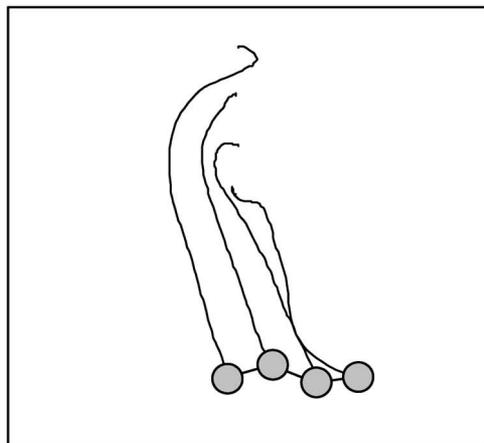
Despite the increased complexity, the obtained results show that the evolved controller preserves its capability of producing coordinated movements both in simulation and in reality (see Fig. 15 and Table IV). Moreover, the performance drops only of 4% and 11% passing from the rigid to the semirigid links, respectively, in the tests with the simulated and real swarm-bots. The performance of the experiments performed with the real

TABLE IV
PERFORMANCE OF THE BEST EVOLVED CONTROLLER TESTED IN
SIMULATION AND REALITY. COMPARISON BETWEEN THE
SWARM-BOTS WITH RIGID AND SEMIRIGID LINKS
(SEE CAPTION OF TABLE I FOR MORE DETAILS)

	Line 4, rigid links		Line 4, semi-rigid links	
	Simulation	Real	Simulation	Real
Avg. perf.	156.96	120.85	150.57	108.00
Std. dev.	28.39	29.53	27.87	34.14
Std. err.	6.35	6.60	6.23	7.63
Ratio with th. max.	0.85	0.65	0.81	0.58
Ratio with sim.	1.00	0.77	1.00	0.72
Rot. equil.	0	0	0	2



(a)



(b)

Fig. 16. Trajectories produced by (a) four simulated and (b) four real s-bots forming a linear swarm-bot with semirigid links during a test lasting 15 s. The gray circles indicate the final position of the s-bots. In the case of real s-bots, the trajectories have been automatically extracted from a video recording of the real s-bots.

s-bots with semirigid links is 28% lower than the corresponding simulation experiments in line with the results presented in Section IV-A. Fig. 16 shows an example of the behavior of the simulated and real swarm-bots assembled through semirigid links. Notice how the swarm-bots modify their shape while moving without losing their ability to coordinate.

F. Coordinated Object Pushing/Pulling Behavior

Fig. 17 shows the case of four s-bots connected to an object rather than between them. In this situation, the s-bots

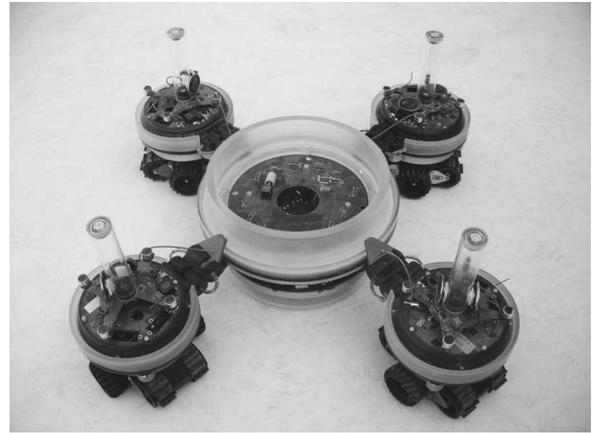
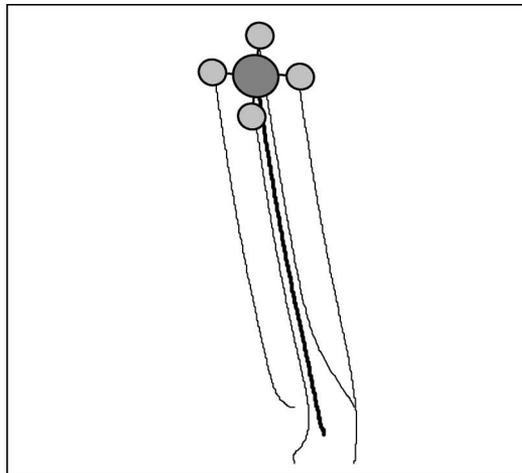


Fig. 17. Four s-bots connected to a cylindrical passive object.

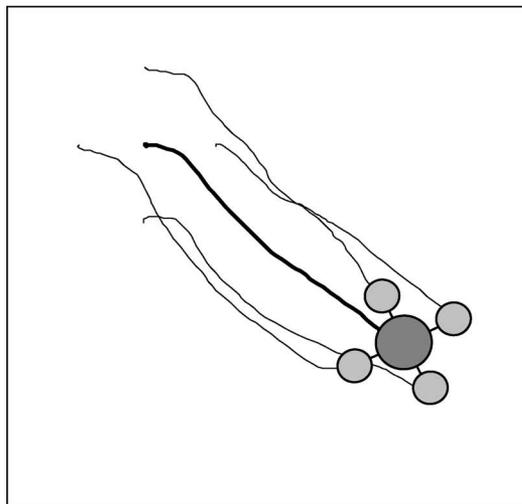
continue to coordinate moving in a common direction while pushing/pulling the object. Notice that the four s-bots and the cylindrical object form a single physical system. In such a situation, as soon as the resistance given by static friction is overcome, the pushing/pulling forces are transmitted through the rigid links of the structure, and coordination can take place. Moreover, a slight resistance produced by the dynamic friction of the passive object does not disturb the coordinated motion because, as shown in Section III-A, the evolved controller keeps moving despite the small traction that comes from the rear. However, as the s-bots are only able to coordinate if the friction of the object with the ground is not too high, the tests in simulation and in reality used a lightweight object. Note that this test was not carried out to study the problem of collective transport, which is not within the scope of this paper (see Section V for a review of the corresponding literature). Its aim was rather to study the robustness of the evolved behavior. In particular, we verified whether the coordination mechanisms underlying such behavior were capable of exploiting the “indirect” traction signals perceived by the s-bots through a passive object to which they were connected.

The tests performed in this experimental condition show that the s-bots preserve their ability to coordinate and to move in a coherent fashion both in simulation and in reality. Consequently, also, the object is transported by the coordinated action of the s-bots. Fig. 18 shows two examples of the trajectories traced by a simulated and a real swarm-bot. The figure shows that after an initial coordination phase, the robots succeed to move in the same direction while transporting the object.

A quantitative comparison between this experimental condition and the case of the four s-bots assembled in a square formation (i.e., the most similar shape) showed a slight performance drop (see Fig. 19 and Table V). In particular, the performance drops of 23% and 29% in the tests run in simulation and in reality, respectively. The decrement of performance is mainly due to a higher probability of falling in the rotational equilibrium. The resistance to motion of the passive object is probably the main cause of this. Indeed, the performance of the experiments performed with real s-bots is 33% lower than the corresponding simulation experiments in line with the case of square formations (27% lower).



(a)



(b)

Fig. 18. Trajectories followed by four (a) simulated and (b) real s-bots connected to a cylindrical object during a test lasting 15 s. The light and dark gray circles indicate the final position of the s-bots and the object, respectively.

G. Analysis of Scalability

To have a general idea of how the performance scales with the number of robots, we measured the time the real s-bots take to converge to a single direction of motion in swarm-bots composed of different numbers of individuals. The time needed by the s-bots to convergence was estimated on the basis of graphs analogous to the one reported in Fig. 8. The results of the tests are reported in Fig. 20 and Table VI. These results indicate that the swarm-bots formed by a higher number of assembled s-bots take longer to coordinate. These data confirm similar results obtained in simulation, for which it was found that the coordinated-motion behavior scales well with the number of robots (see [15] and [17]).

V. RELATED WORK

Coordinated motion is a task that attracted the interest of many researchers and has been commonly studied in the literature. Also referred to as “formation control,” it requires that a

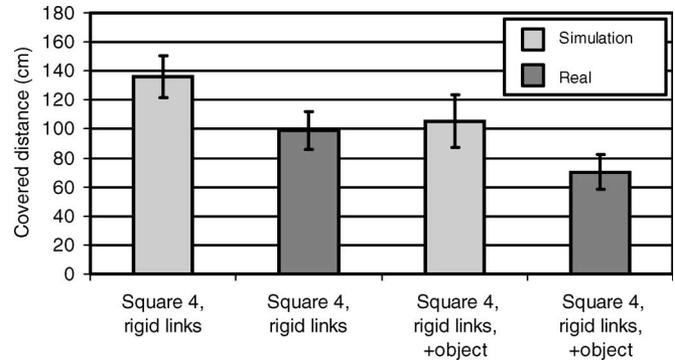


Fig. 19. Performance of the best evolved controller in simulation and reality (average and standard error of the distance covered in 20 trials, each lasting 25 s) (see the caption of Fig. 10 for a detailed explanation of the figure). Additionally, “+object” indicates tests involving s-bots connected through a passive cylindrical object.

TABLE V
PERFORMANCE OF THE BEST EVOLVED CONTROLLER TESTED IN SIMULATION AND REALITY. COMPARISON BETWEEN A SQUARE SWARM-BOT AND S-BOTS CONNECTED TO A CYLINDRICAL OBJECT IN A SQUARE-LIKE FORMATION

	Square 4, rigid links		Square 4, + object	
	Simulation	Real	Simulation	Real
Avg. perf.	136.02	99.00	105.34	70.4
Std. dev.	65.44	57.22	80.72	53.28
Std. err.	14.63	12.79	18.05	11.91
Ratio with th. max.	0.74	0.53	0.57	0.38
Ratio with sim.	1.00	0.73	1.00	0.67
Rot. equil.	4	5	8	9

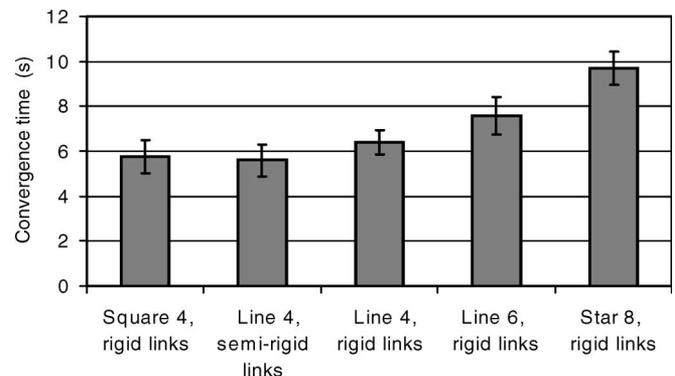


Fig. 20. Time that real s-bots take to converge to a single direction of motion in swarm-bots formed by a different number of robots (average and standard error of the distance covered in 20 trials, each lasting 25 s).

TABLE VI
AVERAGE TIME THAT THE REAL S-BOTS COMPOSING SWARM-BOTS FORMED BY A DIFFERENT NUMBER OF MEMBERS TAKE TO CONVERGE TO THE SAME DIRECTION OF MOTION. AVERAGE, STANDARD DEVIATION, AND STANDARD ERROR FOR 20 REPLICATIONS

	Square 4 rigid links	Line 4 semi-rigid links	Line 4 rigid links	Line 6 rigid links	Star 8 rigid links
Conv. time	5.75	5.60	6.40	7.57	9.70
Stand. dev.	3.38	3.17	2.37	3.66	3.28
Stand. error	0.76	0.71	0.53	0.82	0.73

number of independent entities coordinate their actions in order to move coherently. One of the first works on this topic dates back to 1991 when Wang proved how a simple leader–follower

mechanism could produce a coordinated motion in a group of simulated robots [26]. This is a common strategy to perform a decentralized control of a group of robots, as it reduces coordination to the *a priori* definition of a hierarchy among the robots. The leader–follower paradigm has many different instantiations, in which either the leader role is fixed [27], or it varies according to some arbitration rule [28] or it emerges from the interaction among the robots or between the robots and the environment [11]. In some cases, the leader role is taken by a centralized controller, which plans a trajectory that the robots follow keeping a certain group formation [27], [29], [30]. Finally, a kind of leader–follower paradigm is accomplished defining a neighbor-based hierarchy according to which robots maintain the relative position with respect to a given neighbor [27], [31]. On the contrary, the work presented in this paper does not define any leader that drives the group coordination, because the latter is the emergent result of a self-organizing process.

A coordinated motion can also be performed without keeping the team in a precise formation. In this case, the resulting behavior is closer to what can be observed in many different animal species, such as flocks of birds or schools of fish. Many researchers have provided models for schooling behaviors and replicated them in artificial-life simulations [14]. As an example, it is worth mentioning the seminal work of Reynolds, who defines the behavior of virtual creatures, called “boids,” making use of only local rules [32]. The work of Reynolds has stimulated many other studies on coordinated motion, which are all based on some biological inspiration [33], [34]. These works have self-organization as a common “feature” with the experiments presented in this paper. However, the obtained results are usually limited to simulation, and the experimental setup does not consider the possibility of testing the controllers in real robots.

Among the related works, it is worth mentioning a class of robotic systems developed for collective transport/manipulation. This task is slightly different from the coordinated-motion task studied in this paper, since particular attention is given to the displacement of an object toward a given location or along a given trajectory. In this task, tight coordination among the robots is needed especially in the cases in which the object to be transported must be first lifted and then moved. In such situations, force sensors are often used that provide a feedback mechanism to control the stability of the transported object. Force sensors are not exploited for achieving coordination in the group, as in the experiments presented in this paper. They are rather used to keep under control the planned force to be applied on the transported object [35]–[37] or for correctly distributing the payload in the group [38]. In some cases, collective manipulation has been achieved through centralized approaches [35], [36], a distributed leader–follower approach [28], [38], [39], or a distributed approach based on *a priori* planned trajectories [37].

A different approach characterizes other works that are devoted to minimalism: Collective transport/manipulation is distributed, and individual complexity is minimized [40]. The work of Kube and Zhang [4] is an interesting example of this approach. They start from the assumption that cooperation does

not necessarily require intention, but it can be easily achieved exploiting perceptual cues freely offered by the environment and positive feedback loops that reinforce the collective response. A similar approach is taken in the work presented in this paper. The main difference, which is apart from the experimental details, lies in the coordination mechanism exploited by the robots. In fact, in the former case, the environment contains landmarks (i.e., light bulbs) that guide the robots in locating the object and in moving toward the goal location. On the contrary, in the experiments presented in this paper, no such environmental cue is exploited by the group, but coordination is based solely on a self-organizing process.

Self-organization is also at the basis of some experiments in clustering and sorting of objects [6], [7]. In these works, a number of objects are scattered in a closed arena. The objects can be of different types, and the robots are programmed to collect them in one cluster or to segregate them in concentric rings. The individual behavior can be summarized as follows: Pick up an item, and drop it where the local density of the same type items is higher. This simple rule makes no reference to the formation of a single cluster, which instead emerges through a self-organizing process [6], [7]. Differently from the work presented in this paper, no real coordination within the group is necessary for clustering and sorting. The collective action, instead, enhances the self-organization aspects and speeds up the accomplishment of the task.

VI. CONCLUSION

This paper showed how a group of several robots physically assembled in a swarm-bot can display a coherent behavior on the basis of a simple distributed control system in which individual robots have access only to local sensory information. More specifically, the paper showed how it is possible to evolve a behavior that allows the robots to coordinate their movements on the basis of self-organization principles. The robots start by negotiating a common direction of motion, and then, once coordinated, they continuously compensate possible misalignments caused by noise or other environmental factors. This solution is based on a traction sensor that is able to detect the intensity, and the orientation of the traction that the top part of the robot (which is physically connected with the other robots) exerts on the bottom part (which is in contact with the ground).

The most significant achievement presented in this paper concerns the successful transfer of controllers evolved in simulation to the real robots. The results illustrated show that the neural controller can generalize to conditions that are very different from those in which it was evolved. In particular, the evolved behavior was successfully tested in the following conditions: 1) swarm-bots composed of a larger number of assembled robots (up to eight real robots, but similar results have been obtained in simulation using up to 36 robots [15], [17]); 2) swarm-bots with varying shape; 3) swarm-bots assembled through semirigid links that allow relative motion of the connected robots; 4) swarm-bots that navigate on rough terrains, which produce high noise and disturbances; and 5) robots indirectly connected through a passive object.

Very few works in the literature present collective behaviors tested with physical robots, which have an effectiveness comparable to the system presented in this paper. Such effectiveness is the result of the design methodology that allowed obtaining self-organization in the robotic system along with its characteristic properties. Among these characteristics, we observed the high flexibility of the evolved behavior, both with respect to modifications in the environment and to the structure of the robotic system itself. Another fundamental property of the presented robotic system is the high complexity of the behavior exhibited at the collective level, notwithstanding the simplicity of the mechanisms characterizing the individual level. For instance, the sensory-motor apparatus of the robots involves only one sensor and few motors. Also, the neural controller is the simplest possible, that is, a feedforward single-layer neural network with very few input and output neurons. Therefore, all the complexity of the observed collective behavior resides in the interactions that take place among the robots and between the robots and the environment. These interactions are shaped as traction forces, which are captured by the traction sensor despite the variety of configurations of the robotic system and the number of robots forming it. The analysis of the individual behavior reveals that the interactions through the traction forces can be exploited resorting to two opposing tendencies: the first consists in complying with the motion of the rest of the group. This behavior corresponds to the “positive feedback” mechanisms which is at the basis of the self-organization of the group [14], [24]. The second tendency consists in persevering in the current direction of motion, and it has the important role of favoring the emergence of a common direction of motion and stabilizing the system against temporary disturbances.

It is worth noting that this behavior was obtained through an automatic design methodology, that is, artificial evolution, which is particularly tailored for the synthesis of self-organizing behaviors [15], [21]. In fact, evolutionary methods work in the bottom-up direction, as they define the controller at the individual level and evaluate the performance of the system as a whole. They also tend to produce robust behaviors because unstable solutions and solutions easily affected by disturbances are rapidly eliminated, as they have a poor performance.

It is also relevant to stress that the evolved behavior constitutes an important building block for swarm-bots that have to perform more complex tasks such as coordinately moving toward a light target [17], and coordinately exploring an environment by avoiding walls and holes [17], [25].

In future works, we will continue studying the coordinated motion with the aim of reducing or completely removing those stagnation conditions in which all robots keep moving around their center of mass. This rotational equilibrium may be avoided in different ways, such as providing the robots with additional information (e.g., additional sensors detecting the speed of the two wheels), or by providing the controller with recurrent connections, or both. With these modifications, the robots should be able to detect that the system is in a stagnation condition and, therefore, to trigger a behavior that could break the equilibrium.

REFERENCES

- [1] Y. U. Cao, A. S. Fukunaga, and A. B. Kahng, “Cooperative mobile robotics: Antecedents and directions,” *Auton. Robots*, vol. 4, no. 1, pp. 1–23, Mar. 1997.
- [2] G. Dudek, M. Jenkin, and E. Milios, “A taxonomy of multirobot systems,” in *Robot Teams: From Diversity to Polymorphism*, T. Balch and L. E. Parker, Eds. Wellesley, MA: A K Peters Ltd., 2002.
- [3] M. Dorigo and E. Şahin, “Swarm robotics—Special issue editorial,” *Auton. Robots*, vol. 17, no. 2/3, pp. 111–113, Sep. 2004.
- [4] C. R. Kube and H. Zhang, “Collective robotics: From social insects to robots,” *Adapt. Behav.*, vol. 2, no. 2, pp. 189–219, 1993.
- [5] C. R. Kube and E. Bonabeau, “Cooperative transport by ants and robots,” *Robot. Auton. Syst.*, vol. 30, no. 1/2, pp. 85–101, Jan. 2000.
- [6] R. Beckers, O. Holland, and J.-L. Deneubourg, “From local actions to global tasks: Stigmergy and collective robotics,” in *Proc. 4th Int. Workshop Synthesis and Simulation Living Syst. (Artificial Life IV)*, R. A. Brooks and P. Maes, Eds. Cambridge, MA: MIT Press, 1994, pp. 181–189.
- [7] O. Holland and C. Melhuish, “Stigmergy, selforganization, and sorting in collective robotics,” *Artif. Life*, vol. 5, no. 2, pp. 173–202, 1999.
- [8] A. J. Ijspeert, A. Martinoli, A. Billard, and L. M. Gambardella, “Collaboration through the exploitation of local interactions in autonomous collective robotics: The stick pulling experiment,” *Auton. Robots*, vol. 11, no. 2, pp. 149–171, Sep. 2001.
- [9] A. Martinoli, “Swarm intelligence in autonomous collective robotics: From tools to the analysis and synthesis of distributed control strategies,” Ph.D. dissertation, Comput. Sci. Dept., Ecole Polytechnique Fédérale de Lausanne, Lausanne, Switzerland, 1999.
- [10] M. J. B. Krieger, J.-B. Billeter, and L. Keller, “Ant-like task allocation and recruitment in cooperative robots,” *Nature*, vol. 406, no. 6799, pp. 992–995, Aug. 2000.
- [11] M. Quinn, L. Smith, G. Mayley, and P. Husbands, “Evolving controllers for a homogeneous system of physical robots: Structured cooperation with minimal sensors,” *Philo. Trans. R. Soc. London Ser. A: Math. Phys. Eng. Sci.*, vol. 361, no. 1811, pp. 2321–2344, Oct. 2003.
- [12] C. Anderson, G. Theraulaz, and J.-L. Deneubourg, “Selfassemblage in insects societies,” *Insectes Soc.*, vol. 49, no. 2, pp. 99–110, May 2002.
- [13] E. Bonabeau, M. Dorigo, and G. Theraulaz, *Swarm Intelligence: From Natural to Artificial Systems*. New York: Oxford Univ. Press, 1999.
- [14] S. Camazine, J.-L. Deneubourg, N. Franks, J. Sneyd, G. Theraulaz, and E. Bonabeau, *Self-Organization in Biological Systems*. Princeton, NJ: Princeton Univ. Press, 2001.
- [15] 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 selforganizing behaviors for a *swarm-bot*,” *Auton. Robots*, vol. 17, no. 2/3, pp. 223–245, 2004.
- [16] 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,” *Auton. Robots*, vol. 17, no. 2/3, pp. 193–221, Sep. 2004.
- [17] G. Baldassarre, D. Parisi, and S. Nolfi, “Distributed coordination of simulated robots based on self-organization,” *Artif. Life*, vol. 12, no. 3, pp. 289–311, 2006.
- [18] G. Baldassarre, S. Nolfi, and D. Parisi, “Evolution of collective behavior in a team of physically linked robots,” in *Proc. 2nd Eur. EvoWorkshops2003: EvoROB—Applications of Evolutionary Computing*, R. Gunther, A. Guillot, and J.-A. Meyer, Eds., 2003, pp. 581–592.
- [19] R. Groß, M. Bonani, F. Mondada, and M. Dorigo, “Autonomous self-assembly in swarm-bots,” *IEEE Trans. Robot.*, vol. 22, no. 6, pp. 1115–1130, Dec. 2006.
- [20] V. Trianni, E. Tuci, and M. Dorigo, “Evolving functional self-assembling in a swarm of autonomous robots,” in *Proc. 8th Int. Conf. SAB—From Animals to Animats 8*, S. Schaal, A. Ijspeert, A. Billard, S. Vijayakumar, J. Hallam, J.-A. Meyer, Eds., 2004, pp. 405–414.
- [21] S. Nolfi and D. Floreano, *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*. Cambridge, MA: MIT Press, 2000.
- [22] S. Nolfi, D. Floreano, O. Miglino, and F. Mondada, “How to evolve autonomous robots: different approaches in evolutionary robotics,” in *Proc. 4th Int. Workshop Synthesis and Simulation Living Syst. (ArtificialLifeIV)*, R. A. Brooks and P. Maes, Eds., 1994, pp. 190–197.
- [23] O. Miglino, H. H. Lund, and S. Nolfi, “Evolving mobile robots in simulated and real environments,” *Artif. Life*, vol. 2, no. 4, pp. 417–434, 1995.
- [24] G. Baldassarre, D. Parisi, and S. Nolfi, “Measuring coordination as entropy decrease in groups of linked simulated robots,” in *Proc. ICCS 2004*, 2006, to be published.

- [25] V. Trianni, S. Nolfi, and M. Dorigo, "Cooperative hole avoidance in a swarm-bot," *Robot. Auton. Syst.*, vol. 54, no. 2, pp. 97–103, 2006.
- [26] P. K. C. Wang, "Navigation strategies for multiple autonomous mobile robots moving in formation," *J. Robot. Syst.*, vol. 8, no. 2, pp. 177–195, 1991.
- [27] T. Balch and R. C. Arkin, "Behavior-based formation control for multi-robot teams," *IEEE Trans. Robot. Autom.*, vol. 14, no. 6, pp. 926–939, Dec. 1998.
- [28] Z. D. Wang, E. Nakano, and T. Takahashi, "Solving function distribution and behavior design problem for cooperative object handling by multiple mobile robots," *IEEE Trans. Syst., Man, Cybern. A., Syst., Humans*, vol. 33, no. 5, pp. 537–549, Sep. 2003.
- [29] T. D. Barfoot and C. M. Clark, "Motion planning for formations of mobile robots," *J. Robot. Auton. Syst.*, vol. 46, no. 2, pp. 65–78, Feb. 2004.
- [30] J. P. Desai, J. P. Ostrowski, and V. Kumar, "Modeling and control of formations of nonholonomic mobile robots," *IEEE Trans. Robot. Autom.*, vol. 17, no. 6, pp. 905–908, Dec. 2001.
- [31] J. Fredslund and M. J. Matarić, "A general algorithm for robot formations using local sensing and minimal communication," *IEEE Trans. Robot. Autom.*, vol. 18, no. 5, pp. 837–846, Oct. 2002.
- [32] C. W. Reynolds, "Flocks, herds, and schools: A distributed behavioral model," *Comput. Graph.*, vol. 21, no. 4, pp. 25–34, Jul. 1987.
- [33] C. R. Ward, F. Gobet, and G. Kendall, "Evolving collective behavior in an artificial ecology," *Artif. Life*, vol. 7, no. 1, pp. 191–209, 2001.
- [34] L. Spector, J. Klein, C. Perry, and M. Feinstein, "Emergence of collective behavior in evolving populations of flying agents," *Genetic Program. Evolvable Mach.*, vol. 6, no. 1, pp. 111–125, Mar. 2005.
- [35] T. G. Sugar and V. Kumar, "Control of cooperating mobile manipulators," *IEEE Trans. Robot. Autom.*, vol. 18, no. 1, pp. 94–103, Feb. 2002.
- [36] W.-H. Zhu and J. De Schutter, "Control of two industrial manipulators rigidly holding an egg," *IEEE Control Syst. Mag.*, vol. 19, no. 2, pp. 24–30, Apr. 1999.
- [37] O. Khatib, K. Yokoi, K. Chang, D. Ruspini, R. Holmberg, and A. Casal, "Coordination and decentralized cooperation of multiple mobile manipulators," *J. Robot. Syst.*, vol. 13, no. 11, pp. 755–764, 1996.
- [38] T. Huntsberger, P. Pirjanian, A. Trebi-Ollennu, H. D. Nayar, H. Aghazarian, A. J. Ganino, M. Garrett, S. S. Joshi, and P. S. Schenker, "CAMPOUT: A control architecture for tightly coupled coordination of multirobot systems for planetary surface exploration," *IEEE Trans. Syst., Man, Cybern. A, Syst., Humans*, vol. 33, no. 5, pp. 550–559, Sep. 2003.
- [39] X. Yang, K. Watanabe, K. Izumi, and K. Kiguchi, "A decentralized control system for cooperative transportation by multiple non-holonomic mobile robots," *Int. J. Control*, vol. 77, no. 10, pp. 949–963, Jul. 2004.
- [40] K. Böhringer, R. Brown, B. Donald, J. Jennings, and D. Rus, "Distributed robotic manipulation: Experiments in minimalism," in *Proc. 4th ISER—Experimental Robotics IV*, O. Khatib *et al.*, Eds. Berlin, Germany: Springer-Verlag, 1997, vol. 223, pp. 11–25.



Gianluca Baldassarre was born in Rome, Italy, on February 21, 1969. He received the Laurea degree (B.A. and M.A.) in economics from the Università degli Studi di Roma "La Sapienza," Rome, in 1997 and the Ph.D. degree in computer science from the University of Essex, Colchester, U.K., in 2001.

From 2001 to 2005, he was a Research Fellow with the Laboratory of Autonomous Robotics and Artificial Life, Istituto di Scienze e Tecnologie della Cognizione, Consiglio Nazionale delle Ricerche (LARAL-ISTC-CNR), Rome, where he worked within the "SWARM-BOTS" EU-funded project. Since 2006, he has been a Researcher with the same Laboratory and Team Leader within the EU-funded Integrated Project "ICEA—Integrating Cognition, Emotion, and Autonomy." He has published more than 25 peer-reviewed papers. His research interests are complexity, neural networks, genetic algorithms, mobile robots and robotic manipulators, computational-embodied-neuroscience models of sensorimotor behavior, agent-based models of economic systems and subjective well being.



Vito Trianni received the Laurea (Master of technology) degree in computer science engineering from the Politecnico di Milano, Milan, Italy, in 2000, the Master in information technology from Società Consortile a Responsabilità Limitata, ICT Center of Excellence for Research, Innovation, Education and industrial Labs partnership (CEFRIEL), Milan, in 2001, and the Diplôme d'Études Approfondies and Ph.D. degree in applied sciences from the Université Libre de Bruxelles, Brussels, Belgium, in 2003 and 2006, respectively.

Currently, he is a Research Fellow with the Laboratory of Autonomous Robotics and Artificial Life, Istituto di Scienze e Tecnologie della Cognizione, Consiglio Nazionale delle Ricerche (LARAL-ISTC-CNR), Rome, Italy. His research interests are in the field of evolutionary robotics, swarm intelligence, and self-organization.



Michael Bonani received the M.Sc. degree in micro-engineering from the Ecole Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland, in 2003. Since 2005, he has been developing an educational robot called e-puck at EPFL, where he is working toward the Ph.D. degree in microengineering.

In 2003 and 2004, he was a Research Associate within the SWARM-BOTS project. His research interests include swarm robotics and educational robotics.



Francesco Mondada (S'93–A'95) received the M.Sc. degree in microengineering and the Ph.D. degree from the Ecole Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland, in 1991 and 1997, respectively.

During his thesis, he Cofounded the company K-Team, being both CEO and President of the company for about five years. He is one of the three main developers of the Khepera robot, which is considered as a standard in bio-inspired robotics and used by more than 1000 universities and research centers worldwide. Fully back in research in 2000 and after a short period at CalTech, he has been the main developer of the s-bot platform within the SWARM-BOTS project. He published more than 50 papers in the field of bio-inspired robotics and system-level robot design. He is Coeditor of several international conference proceedings. He is currently a Senior Researcher with EPFL. His research interests include the development of innovative mechatronic solutions for mobile and modular robots, the creation of know-how for future embedded applications, and making robot platforms more accessible for education, research, and industrial development.

Dr. Mondada was awarded the Swiss Latsis University prize for his contributions to bio-inspired robotics, in 2005.



Marco Dorigo (S'92–M'92–SM'96–F'06) was born in Milan, Italy, on August 26, 1961. He received the Laurea (Master of technology) degree in industrial technologies engineering and the doctoral degree in information and systems electronic engineering from Politecnico di Milano, Milan, in 1986 and 1992, respectively, and the title of Agrégé de l'Enseignement Supérieur from the Université Libre de Bruxelles, Belgium, in 1995.

From 1992 to 1993, he was a Research Fellow with the International Computer Science Institute of Berkeley, CA. In 1993, he was a North Atlantic Treaty Organisation–Consiglio Nazionale delle Ricerche (NATO–CNR) Fellow, and from 1994 to 1996, a Marie Curie Fellow. Since 1996, he has been a tenured Researcher with the Belgian National Fund for Scientific Research, and a Research Director of Institut de Recherches Interdisciplinaires et de Développement en Intelligence Artificielle (IRIDIA), the artificial intelligence laboratory of the Université Libre de Bruxelles. He is the author of three books: *Robot Shaping: An Experiment in Behavior Engineering* (MIT Press, 1998), *Swarm Intelligence: From Natural to Artificial Systems* (Oxford Univ. Press, 1999), and *Ant Colony Optimization* (MIT Press, 2004). He is the Inventor of the ant colony optimization metaheuristic. His current research interests include metaheuristics for discrete optimization, swarm intelligence, and swarm robotics.

In 1996, Dr. Dorigo was awarded the Italian Prize for Artificial Intelligence, in 2003, the Marie Curie Excellence Award, and in 2005, the Dr. A. De Leeuw–Damry–Bourlart award in applied sciences.



Stefano Nolfi was born in Rome, Italy, on September 23, 1963. He received the Laurea degree (B.A. and M.A.) in literature and philosophy from the University of Rome “La Sapienza,” in 1987.

Since 1988, he has been a Tenured Researcher with the Italian National Research Council (CNR), Rome. From 1994 to 2001, he coordinated the Laboratory of Neural Systems and Artificial Life of the Istituto di Psicologia, CNR. Since 2002, he has been coordinating the Laboratory of Autonomos Robotics and Artificial Life of the Istituto di Scienze e

Tecnologie della Cognizione, CNR (LARAL-ISTC-CNR). He has been an Adjunct Professor with the University of l'Aquila, L'Aquila, Italy; Lumsa University, Rome, Italy; and University of Calabria, Rende, Italy. He has been a Fellow of the University of California, San Diego; SONY Computer Science Laboratory, Tokyo, Japan; Institute for Advanced Studies of Berlin, Berlin, Germany; and University of New South Wales, Canberra, Australia. He is the author of the book S. Nolfi and D. Floreano, *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines* (MIT Press/Bradford Books, 2000) and of more than 100 peer-reviewed articles. He is one of the inventors of the evolutionary robotic method for the synthesis of autonomous robots. His research interests include autonomous robotics, complex systems, and cognitive science.