

Distributed Coordination of Simulated Robots Based on Self-Organisation

Gianluca Baldassarre Domenico Parisi Stefano Nolfi

Institute of Cognitive Sciences and Technologies,

National Research Council (ISTC-CNR),

15 Viale Marx, 00137 Rome, Italy

baldassarre@ip.rm.cnr.it parisi@ip.rm.cnr.it nolfi@ip.rm.cnr.it

Abstract

Distributed coordination of groups of individuals accomplishing a common task without leaders, with little communication, and on the basis of self-organising principles, is an important research issue within the study of collective behaviour of animals, humans and robots. The paper shows how distributed coordination allows a group of evolved physically-linked simulated robots (inspired by a robot under construction) to display a variety of highly coordinated basic behaviours such as collective motion, collective obstacle avoidance, and collective light approaching, and to integrate them in a coherent fashion. In this way the group is capable of searching and approaching a light target in an environment scattered with obstacles, furrows, and holes, where robots acting individually fail. The paper shows how the emerged coordination of the group relies upon robust self-organising principles (e.g. positive feedback) based on a special sensor that allows the single robots to perceive the “average” group’s motion direction. The paper also presents a robust solution to a difficult coordination problem, that might also be encountered by some organisms, caused by the fact that the robots have to be capable of moving in any direction while being physically connected. Finally, the paper shows how the evolved distributed coordination mechanisms scale very well with respect to the number of robots, the way in which robots are assembled, the structure of the environment, and several other aspects.

1. Introduction

Consider the scenario shown in Figure 1 and Figure 8. A group of robots is placed in a maze with obstacles, furrows, and holes. The mission of the robots is to explore the maze and search for a light target. Some robots form a *swarm-bot* (i.e. a group of physically-linked robots with a particular topological structure) in order to pass over furrows and holes in which they would get stuck by moving alone. The challenges generated by this scenario are several: how can

the assembled robots move in the same direction, avoid obstacles, and approach the light once it is in sight, without centralised coordination and on the basis of poor implicit communication? To face these challenges, the robots should be able to display coordinated behaviours (e.g. coordinated movement, coordinated obstacle avoidance, and coordinated light approaching) and to integrate these behavioural capabilities into a single coherent behaviour.

This paper will show how these problems can be solved (a) by suitably designing the hardware of the robots, (b) by providing them with a special “traction sensor”, introduced for the first time in Baldassarre et al. [5], that allows the robots to detect the direction and intensity of the traction caused by the group, and (c) by developing the controllers of the robots through an evolutionary method that allows exploiting self-organising properties that might emerge from the fine-grained interactions among the robots and between the robots and the environment.

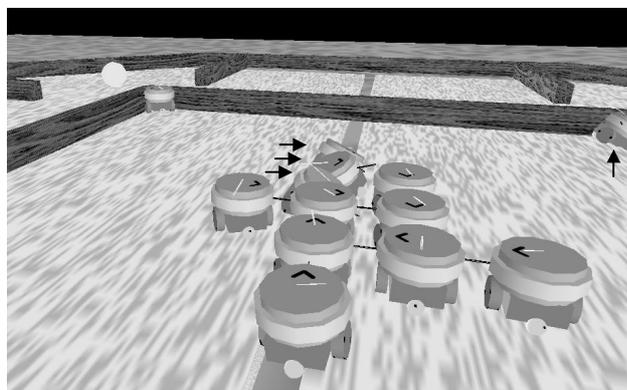


Figure 1: A group of eight assembled robots searches a light target (white sphere at the left top of the picture) in a maze with obstacles, furrows, and holes. Notice how single robots, like the four robots indicated by arrows, get stuck in them.

Collective robotics is a growing research area within the broader field of robotics [8, 14, 16, 19]. Part of this research area focuses on *distributed coordination* of groups of cooperating robots accomplishing common tasks [6, 13, 17, 23]. Distributed coordination implies that the characteristics of the group’s behaviour, such as the basic behaviours through which the task is accomplished, the roles played by the dif-

ferent robots, the synchronisation problems raised by their interactions, and so on, are not managed centrally by one or few “leaders” but are the result of self-organizing processes such as “positive feedback” (if each individual of a group follows a rule of the type “do what the majority does”, the individuals’ behaviours will tend to become homogeneous) or “consumption of building blocks” (e.g. if the number of individuals forming a group is limited, the process of convergence towards the same behaviour caused by a positive feedback mechanism will necessary slow down and then stop) [11, 24]. Distributed coordination is particularly interesting from an engineering point of view since, depending on the specific tasks to be accomplished, it might have some advantages compared to central coordination in terms of: (a) reduced communication needs, (b) resistance to failure of some robots, (c) simplicity of single robots in terms of sensors, actuators, and computational capabilities, (d) scalability to large numbers of individuals.

A research area of robotics that is related to the work presented here focuses on the problem of controlling the different parts of a single robot, for example the legs of an insect-like robot. Several models that tackle this problem assume that the control system of robot (e.g. controlling the legs) is composed of several neural modules (each controlling a leg) that are only scarcely connected between them (e.g. [12]). In these models some aspects that characterize the overall behaviour of the robot are not managed centrally but emerge from the interaction of the different control modules. Notwithstanding this, in these cases the whole system is controlled by a *single neural controller*. Differently from this approach, in the experiments reported here the robots are controlled by *independent neural controllers*: this condition requires distributed coordination.

Distributed coordination is also a central topic in the study of collective behaviour in animals. Bacteria use cell-to-cell communication to implement a self-organising mechanism named “quorum sensing”: with high population density there is enough signal to trigger a behaviour that is not performed at lower densities [15]. In social insects the interaction between simple behavioural rules followed by each individual might lead to rather complex collective behaviours such as chaining and formation of complex 3D structures [1], collective nest building [11], collective food retrieval [18]. These behaviours tend to rely upon self-organisation mechanisms [11]. Self-organization and behaviour coordination is also important in more complex animals such as non-human primates. For example baboons may use “move grunts” to coordinate the group’s transition from resting to moving [25]. Even humans, although endowed with sophisticated capacities such as language and planning to manage coordination, still heavily exploit distributed coordination. In these cases single individuals might even benefit from distributed coordination while being unaware of it (e.g. consider the coordination of people’s actions through market prices [27]). Research studies such as those presented in this paper should produce insights on the general mechanisms of distributed coordination that might underlie all these types of behaviours in different realms.

As mentioned above, in our research evolutionary techniques are used to evolve the behaviours of the robots in simulation. Unfortunately, the time required to run evolutionary experiments can be prohibitively long, especially when detailed simulations of robots and environments are used as it has been done here. To overcome this problem evolving robots were selected for their ability to solve critical aspects of the problem and then the final task was solved on the basis of the fact that they were able to generalize their abilities to new circumstances. This practical necessity, however, allowed shedding some light on what might be an important general property of nature, that is, searching for simple and robust “building blocks” that can be easily rearranged to generate new and more complex behaviours.



Figure 2: The hardware prototype of an individual robot.

The paper also presents a robust solution to tackle a difficult coordination problem generated by the hardware structure of the robots simulated, composed of a mobile chassis and a turret that cannot freely rotate over the chassis and that at the same time should allow the robot to move in any direction while being physically linked to other robots. This problem might affect a large number of possible architectures of robots that should be capable of moving while being physically linked between them, and might also furnish interesting hints on some strategies of movement of some animals that, for biological reasons, cannot have freely-rotating degrees of freedom between different parts of their bodies (e.g. consider some ants collectively transporting a food item).

Section 2 presents the experimental setup. Section 3 illustrates the results in terms of overall behaviour and shows how it generalises to several different conditions. Section 4 and 5 respectively analyse the functioning of the basic behaviours and their integration. Section 6 describes how the physical constraint on the turning capability of individual robots has been dealt with. Section 7 describes how the evolved behaviours scale to large numbers of robots. Finally, section 8 draws conclusions.

2. Experimental setup

The scenario described in the introduction includes some of the problems that are being faced within a research project

funded by the European Union (called “Swarm-bots”) which is developing *swarm-bots*, that is groups composed of fully autonomous robots able to *physically connect and disconnect* to form a single robotic system. These systems can assume different physical shapes and move in order to solve problems that cannot be solved by single robots. In this paper we focus on how a group of robots that are already assembled can produce coordinated movements. Other research carried out within the project is studying how robots can self-assemble and disassemble depending on the task to solve (e.g. [26]).

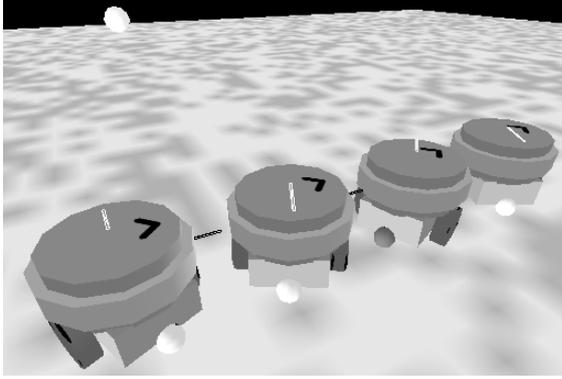


Figure 3: Four simulated robots connected to form a linear swarm-bot. Each robot is composed of a chassis (parallelepiped) to whom two motorised cylindrical wheels and two small passive spherical wheels are attached (the two passive wheels have different colours, dark and light grey, to allow distinguishing the two possible chassis’ fronts). The chassis underpins a cylindrical turret. The black segment between the turrets of two robots represent a physical link between them. The white line above each robot’s turret, that goes from the turret’s centre to a point on its perimeter, indicates the direction of traction and, with its size, the intensity of traction. The environment consists of a flat terrain and a light source (large white sphere at the top left corner of the picture).

Each robot [21] has a cylindrical body with a diameter of 11.6 cm and consists of a mobile base (“chassis”) with two motors each controlling a track and a toothed wheel (Figure 2), and a main body (“turret”). A third motor allows the turret and the chassis to actively rotate with respect to each other. Given that several cables connect different parts of the turret and the chassis, the turret-chassis degree of freedom is limited in the range of $\pm 180^\circ$. The problems due to this physical constraint and the way in which these problems has been solved will be discussed in section 6 (the experiment described in section 3-5 involve robots in which the joint between the turret and the chassis can freely rotate). The turret is provided with two grippers, one rigid and one flexible, that allow the robots to self-assemble and to grasp objects. The turret has a motor with which it can rotate with respect to the chassis. Each robot is provided with a number of different sensors [21], but only the traction sensor described below has been simulated and used in the experiments reported in this paper.

Given that the robots of the project are still under production, a simulator software based on the SDK Vortex™ toolkit, which allows programming realistic 3D simulations

of dynamics and collision of rigid bodies, was designed. Given the high computational costs of simulations, the evolutionary experiments were speeded up as follows: (a) only few relevant characteristics of the sensors, actuators and body of the robot were simulated; (b) the size of the robots, the gravitational acceleration coefficient, and the engines’ power were reduced, so as to have the possibility of increasing the simulation time step without having instabilities; (c) the controller of the robots was evolved in a simplified environment, see Figure 3, and then tested in the more complex environment shown in Figure 1 and Figure 8 (for a comparison of the behaviour of simulated and real robots, and for more details on the simulator, see [21]).

The motor system of a robot was modelled by four wheels: two lateral motorised wheels that modelled the external wheels of the hardware robot and two spherical passive wheels placed at the front and at the rear of the robot that assured its stability. The four wheels were connected to the chassis, which underpinned the rotating turret modelled as a cylinder (Figure 3). The turret was endowed with a simplified gripper which was modelled by creating a physical joint between two robots when needed (Figure 3). The active and passive wheels had a radius of 1.15 and 0.575 cm, respectively. The turret had a radius of 2.9 cm and a height of 2.3 cm.

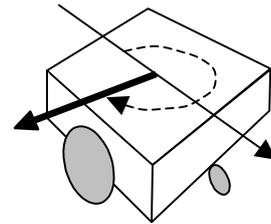


Figure 4: Traction force detected by the robot traction sensor. The parallelepiped represents the chassis. The turret has not been drawn for clarity. The large and small grey circles respectively represent the right motorised wheel and the front passive wheel. The thin arrow indicates the orientation of the chassis, the bold arrow indicates the vector of the traction force that the turret exerts on the chassis, the dotted arrow indicates the angle of the traction.

During evolution, spherical collision models were used for all the wheels and for the chassis, as they required less computational time, but equivalent results were obtained by testing evolved controllers with the collision models shown in Figure 3. The gravitational acceleration coefficient was set at 9.8 cm/s^2 . This low value, that caused a low friction of the wheels on the ground, was compensated for by setting the maximum torque of the motors at a low value, 70 dynes/cm. These decisions were made since the generation of low forces allows using longer integration time steps in the physical simulator without causing instabilities. Preliminary experiments that are currently being carried out with the prototype hardware robots show that those decisions do not preclude the possibility of successfully transferring the controllers evolved in simulation to the hardware robots [7]. The friction coefficient was set at 0.6 (Vortex simulates fric-

tion according to the Coulomb friction model). The desired speed of the wheels varied within ± 5 rad/s. The desired speed applied to the turret-chassis motor was equal to the difference between the desired speed of the left and right wheel times 0.26. The effect of this was that when the chassis turned the turret turned in the opposite direction, so that its orientation did not change with respect to the environment. This greatly aided the robots to turn the chassis when connected to other robots. The state of the sensors and the motors, and the differential equations used by Vortex to simulate the bodies' dynamics, were updated every 100 ms.

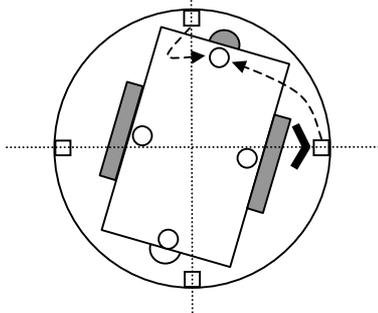


Figure 5: The four empty squares represent the light sensors placed on the turret. The four empty circles represent the virtual light sensors located on the chassis. As an example of how the virtual sensors were activated, the dotted arrows depart from the sensors that were used to compute the activation of the virtual sensor pointed by the arrows themselves.

Each robot was provided with a “traction sensor”, placed at the turret-chassis junction (Figure 4). This sensor returned the direction (angle with respect to the chassis' orientation) and intensity of the force of traction (henceforth called “traction”) that the turret exerted on the chassis. Traction was caused by the movements of both the connected robots and the robot's own chassis. Notice that, by being assembled, the turrets of the robots *physically integrated* the forces produced by all robots. As a consequence, the traction direction measured the mismatch between the directions toward which the robot and the group were trying to move. The traction intensity measured the size of this mismatch. As it will be shown below, given that robots are physically connected the traction sensor informs robots on what the other robots are doing and thus implements a sort of implicit communication [11, 16]. Traction, seen as a vector, was affected by a 2D noise of $\pm 5\%$ of its maximum length.

Each robot was also endowed with four light sensors, positioned on the perimeter of the turret, simulated by using a sampling procedure applied to a real sensor (see [20]). A noise of $\pm 5\%$ of the maximum intensity was added to the sensors. Shadows were simulated by computing the geometrical projections of obstacles (shadows are important for the success of the control strategies illustrated below since if the light target were visible from everywhere in the maze, e.g. if it were positioned in a position higher than the walls, the robots would run into these without being capable to perform obstacle avoidance). In order to provide information about the light gradient with respect to the orientation of the

chassis (this greatly eased control since the wheels were connected to the chassis) four virtual light sensors were simulated. These virtual sensors were activated on the basis of the weighted average of the activation of the two closest light sensors, with weights proportional to the angular distance of them from the virtual sensor (Figure 5).

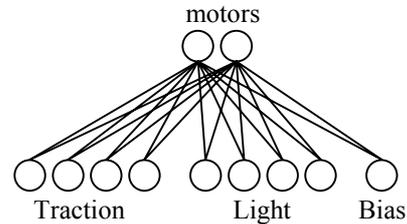


Figure 6: The neural controller of each robot consisted of a two-layer neural network with nine sensory neurons directly connected to two motor neurons.

Each robot's controller (Figure 6) consisted of a neural network with nine sensory neurons directly connected to two motor neurons. The first four sensory neurons encoded the intensity of the traction from four different preferential orientations with respect to the chassis (rear, left, front and right). Each sensory neuron had an activation proportional to the cosine of the angle between the sensor's preferential orientation and the traction direction when the angle was within ± 90 degrees, and zero otherwise. This activation was scaled by the traction intensity. The next four sensory neurons encoded the four virtual light sensors. The last sensor was the bias unit that was always activated with 1.

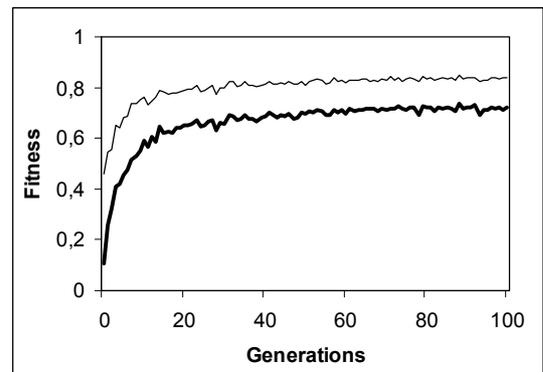


Figure 7: Fitness throughout 100 generations. Thin line: fitness of the best swarm-bot of each generation. Bold line: average fitness of the population. Each line indicates the average performance over 20 replications of evolution.

The activation state of the two motor neurons was normalized within ± 5 rad/s and was used to set the desired speed of the two corresponding wheels and of the turret-chassis motor. Notice that the chassis had two potentially equivalent fronts (this has been exploited by the solution of the turret-chassis constraint problem illustrated in section 6). The direction of the traction and the activation of the light sensors was computed with respect to the “first front”, that is the direction along which the robot moved when the two motors were fully activated with +5, in all the experiments

reported in sections 2 to 5 (experiments reported in sections 6 and 7 used also the “second front”, that is the direction along which the robot moved when the two motors were fully activated with -5).

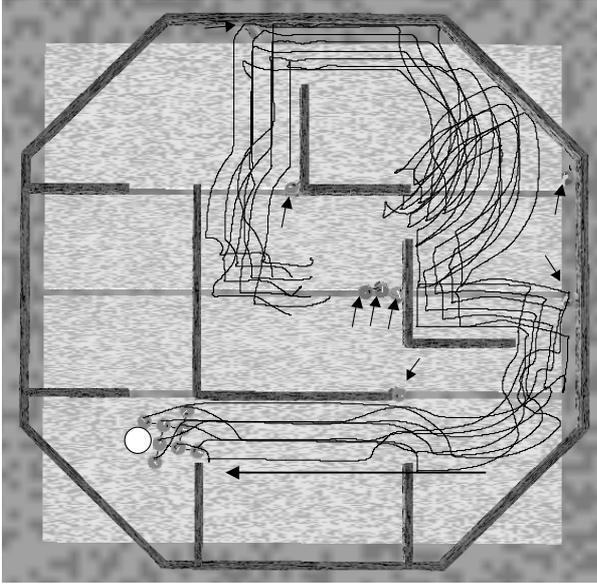


Figure 8: The trajectory produced by a star-shaped swarm-bot in the environment shown in Figure 1 (214 cm wide). The shape of the swarm-bot is depicted in its final position near the light target represented by the white sphere. The black irregular lines indicate the trajectories followed by the eight robots forming the swarm-bot. While the eight single robots (indicated by short arrows) gets stuck in furrows, the swarm-bot passes over them, succeeds to free its robots that fall in the holes near the walls, and searches and finds the light target that was not visible from the starting position (centre of graph). As soon as the target is in sight, the swarm-bot reaches it by following a quite direct path (long arrow).

The connection weights of the neural controllers were evolved [22]. The initial population consisted of 100 randomly generated genotypes that encoded the connection weights of 100 corresponding neural controllers. Each connection weight was represented in the genotype by eight bits that were transformed into a number within ± 10 .

Each genotype encoded the connection weights of four identical neural controllers that were used to control the four robots linked up to form the swarm-bot shown in Figure 3. The case of groups of robots with different neural controllers in which different robots might specialize by assuming different roles was not studied here. Although investigating this aspect would have been interesting, the fact that near-optimal performance was obtained without specialisation indicates that specialisation is not needed to solve the task discussed here (note that this circumstance is task dependent, see for example [3]).

Each swarm-bot was tested in six “epochs”, each lasting 150 time steps lasting 100 ms each. The light source was present only in three of the six epochs. At the beginning of each epoch the orientation of the chassis of the four robots was randomly assigned and, when present, the light source was placed at a random selected position at a distance of

100 cm from the swarm-bot. The 20 best genotypes of each generation were allowed to reproduce by generating five copies each, with 3% of their bits replaced by a new randomly selected value. The evolutionary process lasted 100 generations. The evolution was replicated 20 times by starting with different initial random genotype populations.

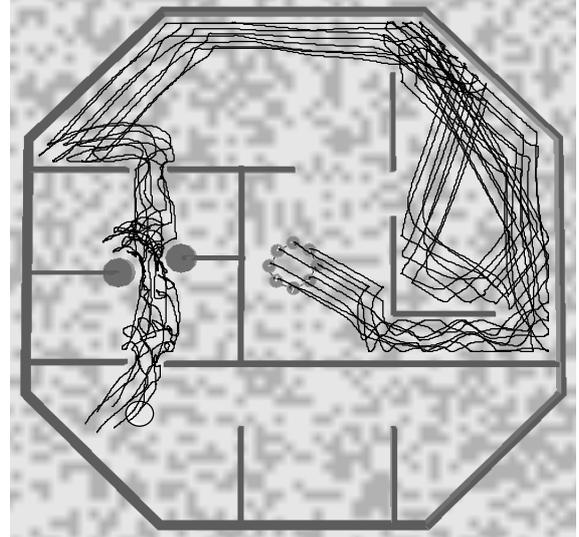


Figure 9: A swarm-bot with a circular shape formed by robots connected through flexible links. The irregular lines, that indicate the trajectories of the single robots, provide an indirect indication of how the shape of the swarm-bot changes while it is moving. With respect to the environment shown in Figure 8, this environment has the same size, does not have furrows and holes, and contains narrow passages such as that formed by the two close cylindrical obstacles.

To develop swarm-bots able to explore the environment and to approach the light target when it was in sight, the swarm-bots were selected for the ability to move as fast and as straight as possible when the light was off, and to move toward the light when the light was on. More specifically, the fitness of a swarm-bot was computed by summing the Euclidean distance between the centre of mass of the swarm-bot at the beginning and at the end of the epoch in the three epochs in which the light was off, and the Euclidean distance travelled by the swarm-bot toward the light in the three epochs in which the light was on. To normalize the value of the fitness in $[0.0, 1.0]$, the total fitness of one swarm-bot was computed by dividing the average distance travelled during one epoch by the maximum distance travelled by a single robot moving straight at maximum speed for 150 steps. Notice how the few and short *epochs focussed* only on the two critical aspects of the final task, namely coordinated motion and light approaching, since an evolution directly tackling the final task (see Figure 8) would have required a prohibitive amount of time.

3. The evolved behaviour

Figure 7 shows how the fitness of the best swarm-bot and the average fitness of the population increase throughout the

generations. By testing for 100 epochs the swarm-bots of the last generations of the 20 replications of the evolution, it was found that the performance of the best swarm-bot is 0.87 while the average performance is 0.78. This means that all evolved robots are able to coordinate so as to move straight when the light is off and to move toward the light when the light is on.

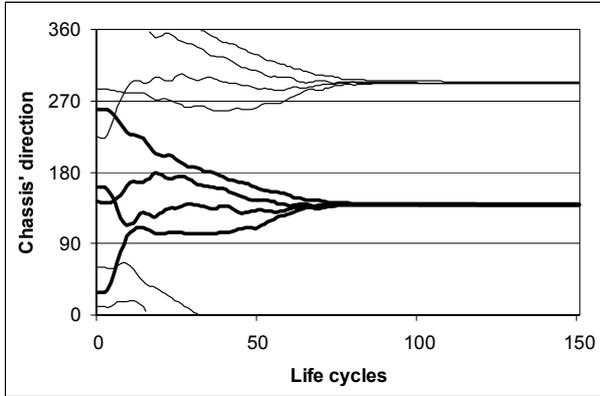


Figure 10: The orientation of the chassis of four robots assembled to form a linear structure placed in the simple environment shown in Figure 3 with the light turned off. The thick and thin lines represent data obtained in two independent tests in which the robots' chassis are assigned initial different random orientations.

By testing evolved swarm-bots in new conditions (the controllers used are those evolved with groups of four robots forming a linear structure) it can be seen that they are able to generalize their coordinated motion and coordinated light approaching ability in rather different circumstances and also to display new interesting behaviours (cf. also [5]).

By placing the evolved swarm-bots in the environment shown in Figure 1 and Figure 8, it can be seen that they are able to explore the environment avoiding walls and to display a coordinated light approaching behaviour as soon as the light is not shadowed by walls. Figure 8 shows how the neural controllers evolved in a team of four robots forming a linear structure are able to generalize their ability to swarm-bots made up by a *larger number* of robots (eight robots instead of the four used during the evolutionary process) and having a *different shape* (a star shape instead of a linear shape). This graph and direct observation of behaviour indicate that the robots forming the swarm-bot, whose chassis initially have different orientations, quickly negotiate a common direction of motion and move along such direction by compensating further mismatches arising during the movement. Later on, when one or more of the robots forming the swarm-bot hit a wall, the swarm-bot changes direction by displaying a very effective and coordinated obstacle avoidance behaviour. After avoiding one wall, the swarm-bot keeps moving straight in the arena until it hits another obstacle. The combination of coordinated motion and collective obstacle avoidance behaviour allows the swarm-bot to explore the environment. When finally the swarm-bot sees the light, it reaches it following a straight path.

Figure 9 shows how the same neural controllers (i.e. the controllers evolved with a group of four robots forming a linear structure through rigid links) are able to generalize their ability to swarm-bots having a circular shape and in which robots are connected through *flexible links*, that is links that have at their centre a hinge joint with a passive degree of freedom rotating around a vertical axis. Interestingly, swarm-bots formed by robots connected through flexible links can modify their shape while moving. Indeed, as shown in Figure 9, the evolved neural controllers display an ability to dynamically adapt the shape of the swarm-bot to the local environmental configuration in order, for example, to pass through narrow passages. As explained in detail below, this result can be explained by considering that the robots that see the light move toward it and, by doing so, produce a traction in that direction that is perceived and followed by the other robots. The generalization ability with respect to the shape of the swarm-bot, can be explained by considering that the control system receives only sensory information with respect to the orientation of the chassis (this is the result of the careful design of the sensors, such as the traction sensor and the light virtual sensors).

Figure 8 shows that single robots get stuck in furrows and holes while the swarm-bot succeeds to pass over them thanks to its larger size. Of course one might have built robots with larger sizes capable of passing over the furrows and holes individually. However, robots with a larger size might have had difficulties in passing through narrow passages, such as those shown in Figure 9. The general idea is that if the robots are capable of self-assembling on the basis of the task to solve and the properties of the environment, they can succeed in a number of different scenarios where single robots with a given size would fail [21].

It has been mentioned that, for efficiency reasons, the evolutionary process was focussed on some critical aspects of the final task, namely coordinated motion and light approaching. When the controllers obtained in *different replications* of the evolution are tested in the final task (see Figure 8 and Figure 9) they display different behaviours. For example, some swarm-bots are more stubborn than others in passing through narrow corridors, some explore the arena very efficiently by doing a sort of collective wall following, some sometimes get stuck in narrow passages when the light is in sight, some other rotate on the spot with some particular initial chassis' orientations (i.e. they reach a stable local minimum in terms of the minimisation of the traction intensity: this dynamical rotating equilibrium is quite important for the performance of the swarm-bots, see section 8). Such a variety of behaviours offers the experimenter the possibility of choosing the controllers of those replications that better satisfy the requirements of the final task.

In order to analyse more in detail how the evolved swarm-bots are able to display the described complex behaviours and to generalize their ability to different circumstances the next two sections will describe how different basic behaviours are produced and how they are integrated.

4. Analysis of the basic behaviours

An analysis of how evolved swarm-bots perform coordinated motion when the light is off or in shadow is now presented. In this condition the robots (a) start to pull/push in different directions, (b) orient their chassis in the direction where the majority of the other robots are pulling/pushing,

and (c) move straight along the direction that emerges from the initial negotiation by compensating successive mismatches that arise while moving. As shown in Figure 10, the absolute direction that emerges from the robots' negotiation changes in different tests depending on the initial orientation of the robots, but the robots always converge toward a single direction.

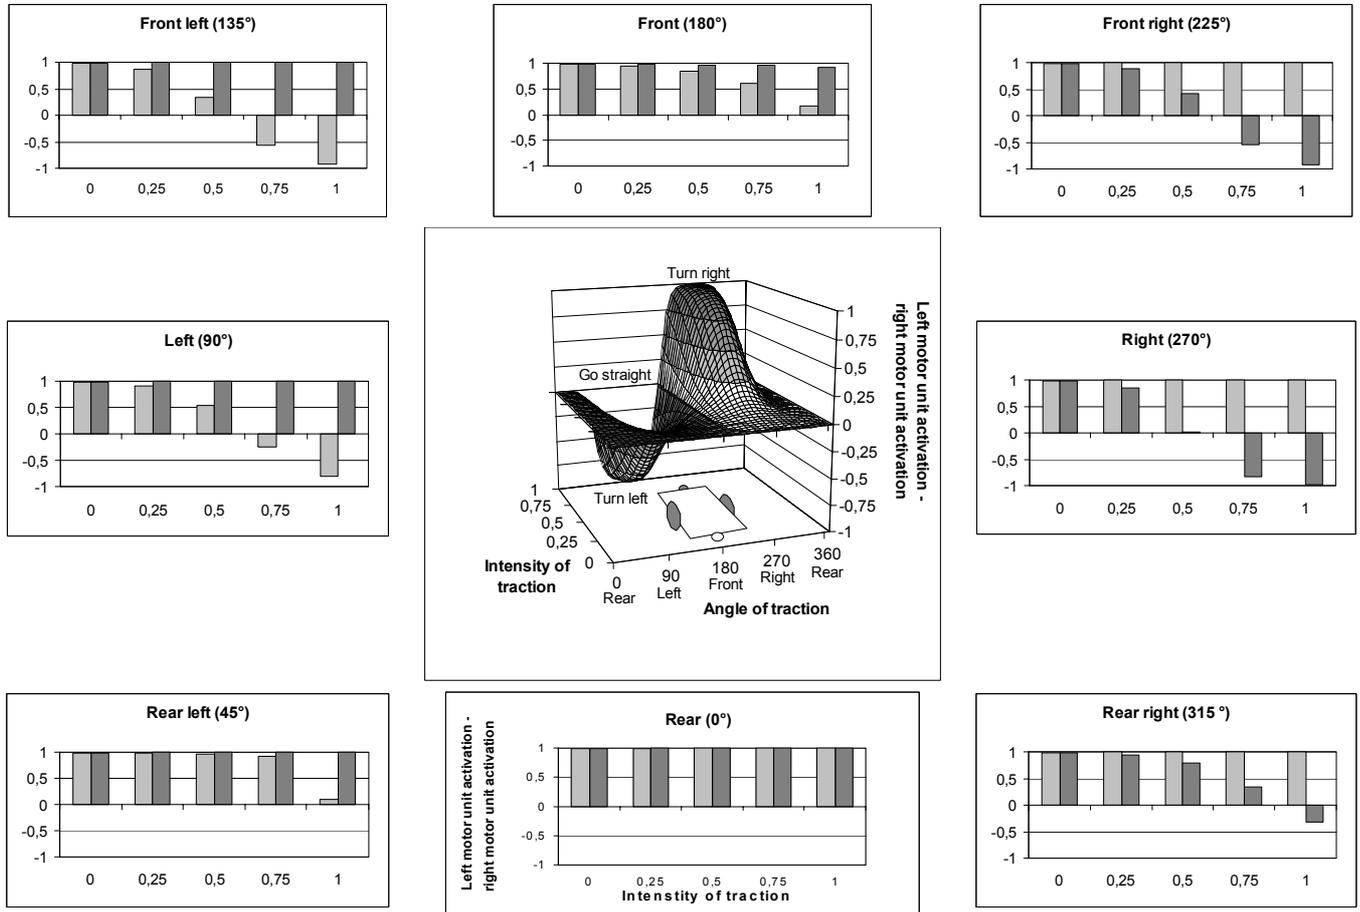


Figure 11: The central graph shows how a robot reacts to a traction with different angles and intensities. The vertical axis indicates the difference between the left and the right motor neurons that, in turn, indicates whether the robot goes straight, turns left or turns right, with different speed. The schematic little picture represents a chassis and should aid the “visualization” of the direction of traction with respect to the chassis itself: the white little wheel represents the rear of the chassis and corresponds to an angle of traction of 0°, measured anticlockwise. The histograms show the desired speed of the left and right wheel as indicated by the activation state of the left and right motor neurons (respectively light grey and dark grey columns of the histograms, measured on the y-axis) for a traction with different angles (see the title of each histogram) and intensities (histograms' x-axis shows four different intensities of traction normalized in [0.0, 1.0]). Data of histograms were obtained with a typical evolved neural controller by manually setting the activation state of the sensory traction neurons and by measuring the corresponding activation state of the motor neurons (the activation state of the light sensors was always set to zero).

By analysing how evolved robots react to traction with different angles and intensities (Figure 11) and by observing the behaviour of the corresponding swarm-bots, it is possible to describe the control's functioning as follows:

- 1) When the traction comes from the front (about 180°, Figure 11, top histogram), the robot is oriented toward a direction that is close to the mean direction of the other robots. In this situation, when the intensity of the traction is low, the robot moves straight (when the intensity is above 0.25, see Figure 11, the robot tends to turn left; this condition, a product of an evolutionary drift, does

not impair performance: in fact traction from front with a high intensity is rare since when the robots move in the same direction the intensity of the traction is low).

- 2) When the traction comes from the left or the right hand side (Figure 11, histograms on the left and the right respectively) there is a significant mismatch between the orientation of the robot and the mean orientation of the other robots. In this condition the robot turns toward the direction of traction, that is towards the mean direction of the other robots, by turning left when the traction

comes from the left hand side and right otherwise. The speed of turning is proportional to the traction intensity.

- 3) When the traction comes from the rear (about 0° , Figure 11, bottom histogram) the robot goes straight at maximum speed independently of the traction intensity. This might be due to the fact that, since when the traction comes from the left or the right the robot has to respectively turn left or right, the 0° point (rear) represents the discontinuous separation between the two different turning behaviours. However, this behavioural trait also produces a “stubborn” tendency of the robots that might have a functional role for coordination (see below).

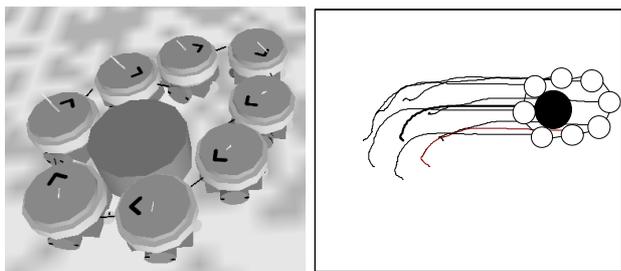


Figure 12: *Left*: eight robots drag an object around which they are connected. *Right*: the irregular thin and bold lines indicate the trajectories left respectively by the robots and the object. The white and black circles represent the final position of respectively the robots and the object.

Summarising, the ability to display coordinated motion is the result of a conformist tendency, that is a tendency to follow the direction of the traction that provides an indication of the average direction of motion of the other robots (see previous points 1 and 2). As mentioned in the introduction, this ability is a specific instance of a general self-organising principle, named *positive feedback*, at the basis of which there is usually a conformist tendency (“conform your behaviour to that of the others”). This principle allows groups of individuals to converge on the same selection among multiple possible alternatives. The principle is important for distributed coordination since it allows groups, such as school of fish and flocks of birds, to behave as if their members were parts of a single organism or system without the need for a centralised decision maker [11, 24]. With this respect, the novelty of the results presented here is having shown that the same principle can work on the basis of the physical interactions between the robots, mediated by the their bodies and perceived through the traction sensors.

Likely, coordinated motion is also a result of a *stubborn tendency*, that is a tendency to persevere in one direction of motion independently of the intensity of the traction. The stubborn tendency is due to the following factors: (a) turning toward the traction direction takes time (i.e. the mismatch cannot be compensated for instantaneously), (b) the tendency to go straight when the traction intensity is low, independently of its direction, and (c) the tendency to go straight when the traction comes from the rear (i.e. from around 0° , see point 3 above). The stubborn tendency is likely to play an important role in the ability of the robots to keep the

equilibrium state once a common direction of movement has emerged from the negotiation and to avoid never-ending negotiation phases in which robots keep changing orientation in order to eliminate small mismatches (in this case the swarm-bot would enter into a limit cycle dynamics).

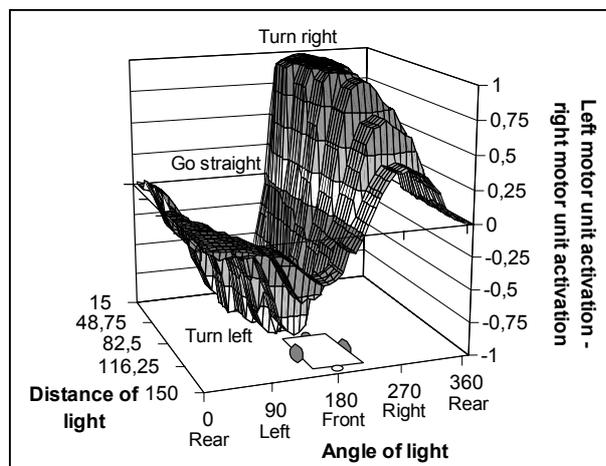


Figure 13: This graph shows how a robot reacts to different orientations and distances with respect to the light. The vertical axis indicates the difference between the left and the right motor neurons that, in turn, indicates whether the robot goes straight, turns left or turns right with different speed. Data have been obtained for a typical evolved neural controller by placing a single robot at different orientations and distances with respect to the light and by recording the corresponding activation state of the motor neurons. The activation state of the traction sensors was always set to zero.

Another effect of the stubborn tendency, with particular reference to the tendency to move straight when the traction come from the rear, is that evolved robots spontaneously display coordinated object pushing/pulling behaviours when linked to or around an external object. Figure 12 shows how robots connected around an object coordinate their motion toward a single direction and succeed to move the object notwithstanding the fact that the object’s friction with the ground produces a traction in the opposite direction of the group’s motion (of course the robots fail to coordinate if the object is heavy and generates a large traction intensity).

Now a description of how evolved single robots and swarm-bots consisting of several assembled robots spontaneously display obstacle avoidance behaviour will be presented. When a single robot hits an obstacle, the collision produces a traction with a direction that points away from the obstacle. The robot follows this traction and so avoids the obstacle. When a swarm-bot consisting of several assembled robots hits an obstacle, only a single or few robots collide with it. However, the resistance of the obstacle will propagate through the turrets of these robots to the turrets of the other robots via the links. As a consequence, all the robots will start to turn away from the obstacle, eventually tending to select slightly different directions. Traction will immediately average between this tendencies and will guarantee a synchronous well coordinated turning of the whole swarm-bot. Indeed, direct observation of behaviour shows

how the single robots perform a surprisingly highly synchronised turning when the swarm-bot hits an obstacle.

Finally, how evolved swarm-bots display a collective light approaching behaviour is now explained. The fact that individual robots display such behaviour can be explained by analysing how evolved neural controllers react to light sources located at different angles and distances with respect to the robot. As shown in Figure 13, when the traction is null, evolved robots follow the light gradient by turning left or right when the light is respectively located at the left or at the right side of the robot's chassis.

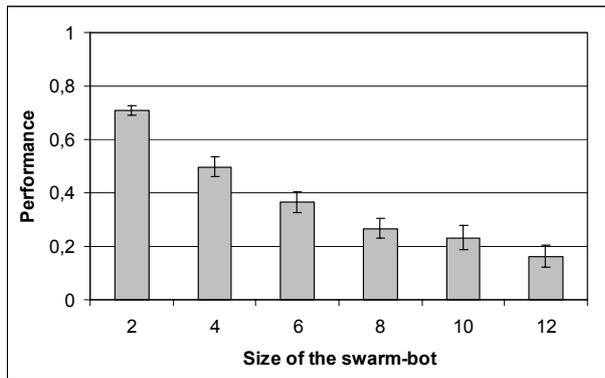


Figure 14: Performance of swarm-bots composed of two to twelve robots forming a linear structure in the simple environment shown in Figure 3 with the light always on. The performance is measured as the fraction of distance covered by the swarm-bots with respect to the space covered by a single robot moving at maximum speed for the same amount of time (150 cycles). Only one robot forming the swarm-bots is provided with the light sensors information (the light sensors of the other robots are always set to zero). Data have been obtained by using the usual neural controller evolved with four robots forming a linear swarm-bot. Columns and bars respectively indicate averages and standard errors over 60 epochs.

The ability to display a coordinated light approaching behaviour despite the fact that individual robots forming a swarm-bot have different sensory information (because of their different relative positions and because of shadows) can be explained by considering, once more, the ability to coordinate through the traction sensor and the effects of the motor behaviour of each individual robot on the other robots. Indeed, as soon as a single robot starts to perceive and follow a light gradient, it creates a traction force that is felt by the other robots which, as a consequence, will also move in the direction of the light. To confirm this, Figure 14 shows an experiment where a single robot provided with light sensors is able to drive a swarm-bot composed of up to twelve robots deprived of the signal from the light sensors towards a light target. Although performance decreases with the size of the swarm-bot, the experiment shows the power of traction in supporting coordination (notice that while this test with robots endowed with different sensors was directed to show the power of coordination based on traction, section 7 will systematically investigate how the whole set-up scales to larger swarm-bots composed of many robots endowed with the same sensors).

5. Integration of behaviours

The ability to integrate the different behaviours described in the previous section into a single coherent behaviour results from the combination of three mechanisms.

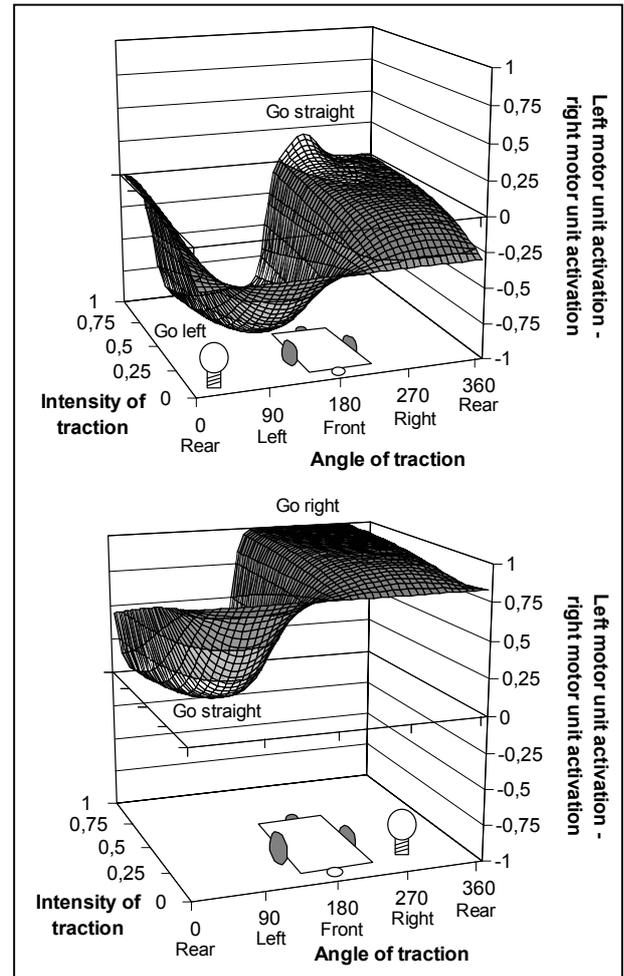


Figure 15: The pictures show how a robot reacts to traction with different direction and intensities when it also perceives a light at a distance of 100 cm coming from the left (90°) or from the right (270°) side (top and bottom pictures, respectively). The y-axis indicates the difference between the left and the right motor neuron that, in turn, indicates whether the robot goes straight or turns left or right with different speed.

As the reader might have already realised from the analyses presented above, one important mechanism is the ability to coordinate through the traction sensor and through the effects that arise from the fact that the robots are physically linked. In fact, as it has been shown, the ability to coordinate through the traction sensor and the fact that robots are physically assembled not only allow swarm-bots to display coordinated motion but also play a crucial role in their ability to display coordinated obstacle avoidance and coordinated light approaching behaviours. In fact, the tendency to coordinate through the traction sensors and links between the robots assures that the swarm-bot produces a coherent and coordinated behaviour even when the individual robots

have different, incomplete or noisy perceptions. In other words, the ability to coordinate through traction constitutes an important behavioural building block also for the ability to display other more complex behaviours.

The second mechanism, already mentioned above, consists in the fact that the control strategies responsible for the ability to display the different basic behaviours described in the previous section are independent of the particular shape of the swarm-bot. This is due to the fact that the direction of the traction returned by the traction sensor corresponds to the direction of the force that the turret exerts on the chassis *independently* of the particular *turret's orientation*. This characteristic is important not only because it allows the swarm-bots to behave robustly independently of the number of robots that form them and the way in which they are assembled, but also because it allows the swarm-bots to behave robustly even when their shape changes dynamically adapting to the environment's properties.

A third important aspect consists in the fact that the ability to coordinate through the traction sensor and the ability to approach the light and to avoid obstacles integrate without interfering with each other. This can be seen by examining, for example, how the evolved robots react when the information coming from the traction and the light sensors are consistent or inconsistent. With this regards, Figure 15 shows how an evolved robot reacts to traction with different directions and intensities when it also perceives a light at a fixed distance of 100 cm coming from a fixed position. When the light comes from the left side (Figure 15, top graph) the robot turns mostly left by varying the speed of turning on the basis of the intensity of the traction. On the contrary, when the light comes from the right side (Figure 15, bottom graph) the robot turns mostly right by varying the speed of turning on the basis of the intensity of the traction. The combination of these two tendencies allows the robots to approach the light and to maintain the group's coherence at the same time. Figure 15, however, only shows how the control mechanisms responsible for the two behaviours interact from a static point of view. The interaction from a dynamical perspective is more informative.

One first important aspect to notice is that the three basic behaviours described above are different forms of *taxis* or *anti-taxis* (i.e. behaviours through which a robot should approach or avoid a stimulation source: notice how the surface graphs of Figure 11 and Figure 13 have similar shapes), and as a consequence in some conditions they can suitably integrate and sum up. A second important aspect is that the behaviour based on traction is a *special type of taxis* in that it is *not* related to a stimulus *anchored* to the environment, as the light pursuing behaviour, but it is related to the average motion direction of the group. This direction tends to change dynamically on the horizontal plane (due to the mechanisms described above) so as to eventually coincide with the direction of the light. When the two directions match completely, the tendency to move toward the average direction of the group and toward the light gradient tend to sum up and to amplify each other as shown in Figure 16.

Referring to the single robot, the possibility of integrating the coordination behaviour based on traction with other behaviours eliminates the need to use behavioural selection mechanisms such as the "subsumption architectures" [10] in which only one control mechanism among many is activated at each time. The possibility of avoiding behaviour selection mechanisms, in turn, eliminates the need of evolving arbitration mechanisms able to select the right behaviour at the right time (e.g. coordinated motion or obstacle avoidance or light approaching). Moreover, it allows evolution to exploit the synergies that might emerge from the interplay between different basic behaviours. With these respects, the controller emerged during evolution has a functionality that has some similarities with that of the "motor-schema based architectures" [2]. These architectures generate the overall robots' behaviour as a weighted summation of building-block behaviours (see [4] for some hand-coded controllers based on this intuition, and for a comparison of them with evolved controllers in tasks similar to those tackled in this paper).

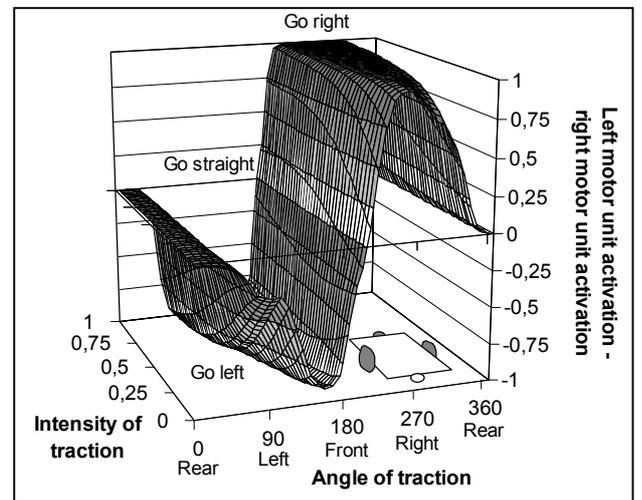


Figure 16: The graph shows the behaviour of a robot when it perceives a traction with different angles and intensities and a light with the same angle as the traction and a fixed distance of one metre.

It should be noted that the fact that the control mechanisms responsible for the three different basic behaviours cooperate and never interfere with each other is not only due to the characteristics of the basic behaviours themselves but also to the way in which sensory information is encoded in the sensors, to the morphology of the robots' body, and to the interaction of these characteristics with the environment. For instance, the possibility of integrating obstacle avoidance and coordinated motion behaviours is due to the fact that robots collide with obstacles with the turret (since the turret is bigger than the chassis) and this generates a traction with a direction opposite with respect to the collision point. A second example refers to the obstacle avoidance and the light approaching behaviours. The fact that obstacles shadow the light has the effect that the anti-taxis behaviour related to obstacles never interferes with the taxis behaviour

related to light. In conclusion, this means that the possibility of finding simple and robust solutions, from the point of view of the control system of the robots, crucially depends on the way in which sensory and motor information are encoded, on the structure of the robot's body, and on the relation of these with the environment's properties.

6. Solution of the problem caused by the constraint of the turret-chassis joint

As previously mentioned, the chassis of the robots has a limited rotational degree of freedom with respect to the turret due to the control and energy cables that connect them. In particular, starting from the position aligned with the turret front, the chassis can rotate only 180° , either clockwise or anticlockwise with respect to the turret itself before encountering physical stops. This constraint makes the coordination problems significantly more difficult. For example, in order to coordinate with the other robots, an individual robot cannot simply turn toward the direction of the traction. In fact if the physical constraint is located between the robot current orientation and the traction direction, it should turn (up to 360°) in the opposite direction. The constraint thus introduces a severe difficulty in the type of sensory-motor mapping that should be produced by the neural controller.

It should be noted that a similar problem might affect self-assembled natural organisms that produce collective coordinated movements (e.g. ants carrying a food item). In fact, due to the physiological continuity of biological bodies, organisms cannot freely rotate the body components used for locomotion with respect to the body components linked to the rest of the group or to a transported object. As a consequence, variations in the direction of movement of the group might force single individuals to suddenly change their motion behaviour.

By exploiting the fact that robots have two potentially equivalent orientations, it was possible to find a very robust and effective solution to the problem. This solution allows the control strategies described in the previous section to continue to work even in the presence of the physical constraint. The solution is now explained in detail. The circular shape of the robots, the fact that the motorised wheels can turn in both directions, and the fact that sensors are homogeneously distributed on the perimeter of the robot imply that the two fronts of the chassis are potentially equivalent both from the motor and the sensory point of view. Henceforth, the two orientations will be respectively referred to as "first front" and "second front". The first and second front correspond to the direction of movement of the chassis when the motors are both fully activated respectively with $+5$ or -5 rad/s (in the figures shown in the paper the two fronts respectively correspond to the position of the dark grey and light grey passive wheels). In the remaining part of this section, first it is shown how the same controllers evolved in the experiments illustrated in the previous sections can be used to guide the robots' chassis using either one of the fronts. Then it is shown how the inversion of the

front used, triggered in particular conditions, allows robots with the constraint to display the same behavioural capabilities of the robots described in the previous sections.

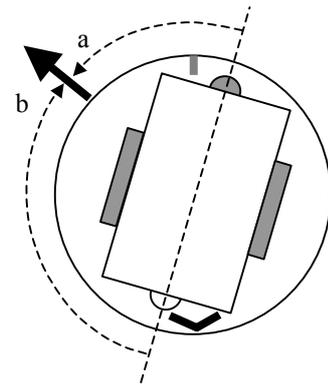


Figure 17: Schematic representation of the effect of a front inversion from the point of view of a robot that is turning toward the direction of traction. The circle and the bold open "V" within the circle indicate respectively the turret and its orientation. The empty rectangle, the two small grey rectangles, and the two semicircles indicate respectively the chassis, the active wheels, and the passive wheels. The bold arrow outside the turret indicates the traction direction. The small grey segment indicates the physical constraint on the turret that cannot be overcome by the chassis first front (grey passive wheel) while turning in either one of the two directions. The dotted arrows a and b indicate the direction in which the chassis would turn while respectively using the first and second front.

In order to invert the front used by an evolved robot, one has to invert the encoding of the sensory and motor neurons. To invert the encoding of the sensory neurons one has to compute the activation of the sensory neurons encoding traction and the sensory neurons encoding the light with respect to the other front (so one has to exchange the activations of the left and right neurons, and of the front and rear neurons). To invert the direction of movement and turning, one should invert the sign of the motor commands (ranging from -5 to $+5$ rad/s) and the encoding of the motor neurons (i.e. the left motor neuron should control the right motor and the right motor neuron should control the left motor). By using the same controllers described in the previous section and by inverting the front according to the procedure just described, it was observed that: (a) when a front inversion takes place while robots are navigating in an open arena, all robots synchronously invert their front of motion so that the entire swarm-bot suddenly inverts its direction of motion; (b) when a front inversion takes place while robots are navigating towards a light target, they invert their direction of movement, turn 180° in the direction of the light target, and then continue to navigate toward it.

The problems introduced by the constraint on the turret-chassis degree of freedom can be solved by switching each robot's front from the current front to the other front each time the robot's chassis reaches the constraint, either while turning anticlockwise or clockwise. The detailed functioning of a front inversion solution from the point of view of a single robot is illustrated by the scheme in Figure 17. The bold

arrow indicates the direction of the traction. Assuming that the robot starts to move by using the first front, the controller first turns the chassis anticlockwise along the direction indicated by the dotted arrow “a”. Given that during this

turning the chassis reaches the joint constraint, the front is changed. This causes the controller to turn the chassis clockwise along the dotted arrow “b”, hence moving away from the joint constraint.

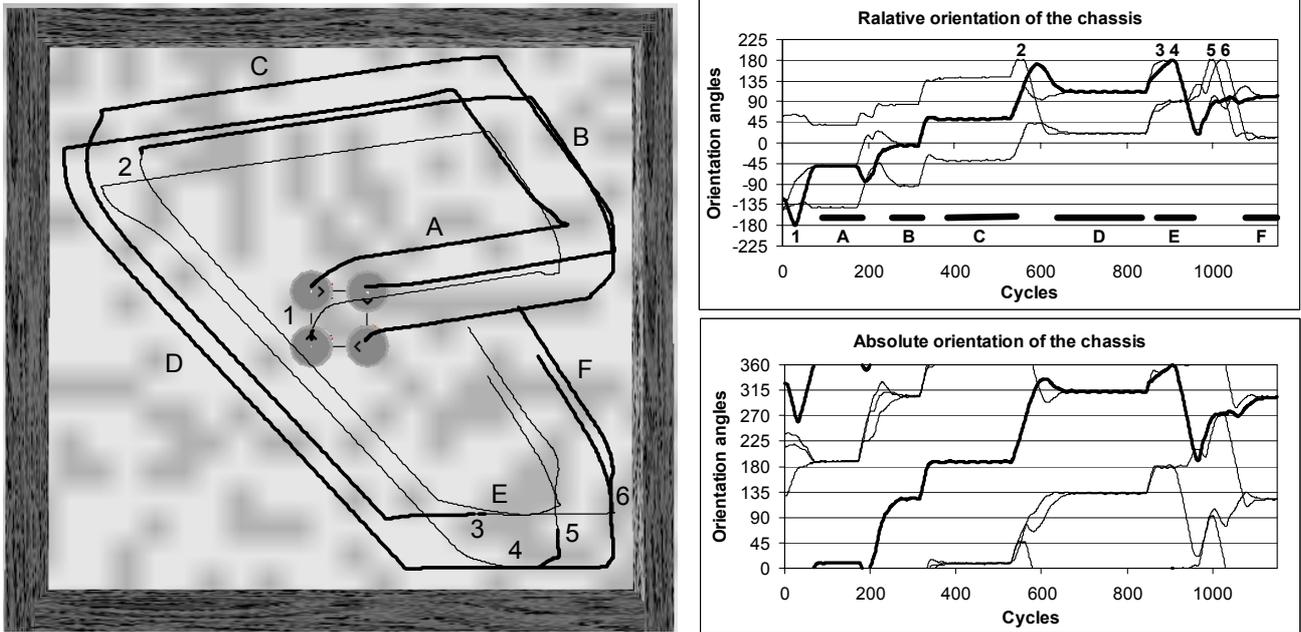


Figure 18: *Left*: A swarm-bot consisting of four robots with the turret-chassis degree freedom constrained, tested in an environment surrounded by walls. Bold and thin lines indicate the trajectories produced by the robots during phases in which respectively the first or the second front are used (this should illustrate how the robots use the two fronts). *Right top*: Orientations of the chassis with respect to the turret of the four robots during the behaviour displayed in the left part of the figure. Orientations are indicated with angles ranging in $\pm 180^\circ$ measured anticlockwise. Bold segments and letters indicate the phases during which the swarm-bot moves straight (cf. figure on the left). Numbers indicate the cycles in which a robot reaches the constraint of the turret-chassis joint and inverts its front. In both graphs on the right, the orientations of the robot at the bottom left of the swarm-bot shown in the map on the left are highlighted by a bold line to allow the reader to follow the behaviour of a particular robot. *Right bottom*: Absolute orientations of the chassis of the four robots during the behaviour displayed in the map on the left. The angle is measured anticlockwise and 0° corresponds to the west on the map.

Figure 18 (left) shows the behaviour of four robots assembled to form a square swarm-bot in an environment surrounded by walls (notice that the robots turrets have four different relative orientations). The turret-chassis degree of freedom is limited within $\pm 180^\circ$. Robots are provided with the same evolved neural controllers described in the previous sections but the front is inverted every time the turret and-chassis joint reaches the constraint. In the test the swarm-bot is set at the centre of the environment. Initially all robots use the first front and have chassis with random orientations. The robots quickly converge to a common direction of movement. During this phase, at about cycle 50, one of the robots reaches its constraint and changes front (see point indicated with number 1 in Figure 18). Then the swarm-bot moves straight by avoiding two walls. During this phase the joint constraints are never reached and consequently no further change of fronts occurs. The second change of front occurs in one of the robots while the swarm-bot avoids the next wall (see point indicated with number 2 in Figure 18). During the next two wall avoidance behaviours some of the robots alternate between the two fronts before stabilizing on one of the two (see points indicated with number 3-4 and 5-6 in Figure 18). Summarising, the

front inversions allow each robot to keep the turret-chassis joint within the $\pm 180^\circ$ constraint (Figure 18, top-right graph) and, at the same time, to contribute to the movement of the swarm-bot in any direction (Figure 18, bottom-right graph).

7. Scalability of the controllers

This section describes the results of a set of tests run to evaluate how evolved neural controllers generalize their abilities to larger swarm-bots. More specifically, the section shows how the best controller evolved for the ability to control a swarm-bot of four robots with a linear shape described in section 3 behaves in swarm-bots with a square shape and composed of 4, 9, 16, 25 and 36 robots respectively (see Figure 19 for an example of a swarm-bot composed of 36 robots). More specifically, swarm-bots of different size, with and without the turret-chassis joint constraint, were tested with respect the ability to: (1) coordinate as quickly as possible; (2) move as fast as possible after the initial coordination; (3) avoid obstacles; (4) approach a light target.

Figure 21 reports the performance of swarm-bots of different size, with and without the joint constraint, with respect to the four abilities mentioned above. Each histogram

indicates the average distance travelled by the swarm-bots, normalised by the maximum distance that a swarm-bot could travel in the conditions set up to test the aforementioned four abilities (the conditions are described below). Each histogram bar is the average of up to 100 trials in which robots' chassis had different initial random orientations. As mentioned in section 3, sometimes the swarm-bots fall into a dynamical equilibrium where they rotate on themselves instead of moving along a straight line. As shown in Figure 20 (coordinated motion test), this is less likely for swarm-bots with a larger size. To deparate the results reported in Figure 21 from this effect, the graphs of this figure show averages referring only to trials where the performance was above 0.3 (the swarm-bots usually do not achieve this performance threshold when they fall into the rotating dynamical equilibrium).

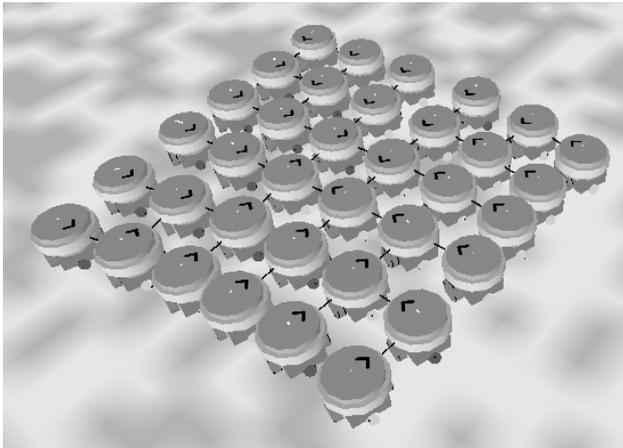


Figure 19: A snapshot of a swarm-bot composed of 36 robots with the turret-chassis joint constraint that is moving after a coordination phase. Notice how the robots have their turrets oriented in four different directions (as shown by the arrows placed on the turrets) and how during the negotiation phase different robots have selected different fronts (indicated by the different orientations of the grey and white passive wheels).

The top-left graph of Figure 21 shows the distance travelled by the swarm-bots in 150 cycles in an open arena. The distance travelled by the swarm-bots provides an indication of how quickly the robots start to move, and this in turn provides an indication of how quickly they negotiate a common direction of movement. As the graph shows, robots evolved for the ability to produce coordinated movement in a swarm-bot composed of four robots assembled to form a linear structure generalize their ability very well in swarm-bots with a larger size. However, the results also show that swarm-bots with larger size tend to take more time to coordinate.

The top-right graph of Figure 21 shows the distance travelled by the swarm-bots between cycle 150 and cycle 300 in an open arena. Since during the first 150 cycles the swarm-bots coordinate and start to move in most of the trials, the distance travelled during the second 150 cycles provides an indication of how fast and how straight swarm-bots are able to move once the robots are coordinated. The

results indicate that the swarm-bots are able to move close to maximum speed independently of their size.

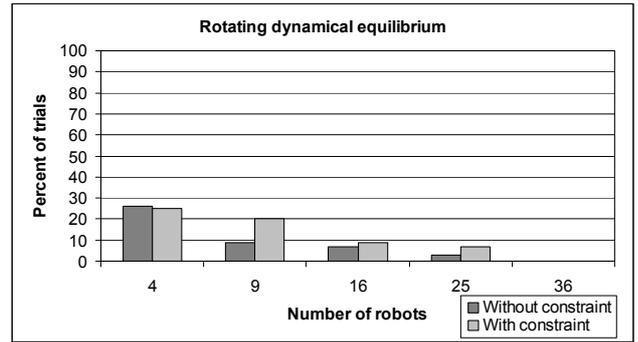


Figure 20: Percent of times in 100 coordinated motion tests in which some swarm-bots having different sizes fall into the rotating dynamical equilibrium.

The bottom-left graph of Figure 21 shows the distance travelled by swarm-bots between cycle 50 and 250 placed in the central part of an arena surrounded by walls (see Figure 18). Since after 50 cycles swarm-bots are usually coordinated, and since during the following 200 cycles they typically bump and avoid one or two walls, the distance travelled provides an indication of the ability of the swarm-bots to avoid obstacles. The results indicate how also this ability generalizes very well with respect to the size of the swarm-bots, even if a decrease of performance can be observed since the robots tend to take longer to negotiate a new direction of movement after avoiding the obstacles if they are more numerous (cf. the top-left graph of Figure 21).

Finally, the bottom-right graph of Figure 21 shows the distance travelled by the swarm-bots toward a light during 150 cycles when placed in an open arena one meter away from the light target. The distance travelled in this conditions provides an indication of the ability of the robots to coordinate and to move toward the light. The results indicate that there is a good level of generalization but performance tends to decrease with the size of the swarm-bots. This decrease in performance can be explained by considering that the proportion of robots that perceive the light decreases with the increase of the size of the swarm-bots due to shadows. This result confirms what was said in section 4: performance tends to decrease when only a small subset of the robots forming the swarm-bot can see the light target.

A last important result shown in all graphs of Figure 21 is that robots provided with the ability to change their front generalize very well with respect to the presence of the turret-chassis joint constraint. In fact, as shown in all the graphs of the figure, the performance of swarm-bots with the constraint is only slightly lower than the performance of swarm-bots without the constraint in tasks in which the robots should coordinate. In the case of the task reported in the bottom-right graph of Figure 21, robots with the constraint sometimes even outperform those without (an explanation of this result has not been found yet).

8. Conclusions and future work

The paper has described how an evolved group of physically linked robots can solve problems that could not be solved by single robots, on the basis of simple control mechanisms, a suitable integration of them, and robust coordination mechanisms. In particular, the analysis of the evolved robots indicates that the ability to coordinate through a novel “traction sensor”, that in turn allows exploiting the physical interaction between assembled robots, not only is at the basis of the ability of the robots to display coordinated motion but

it also constitutes an important building block for the ability to display other more complex behavioural capabilities. Moreover, the evolved controller generalises and scales quite well with respect to the number of robots, the way in which robots are assembled, the structure of the environment, and several other aspects. Given the success of the controller based on the traction sensor, this sensor has been implemented in the hardware robots. Preliminary experiments with these robots are showing that the controllers evolved in simulation have a quite good performance in coordinated motion tests [7].

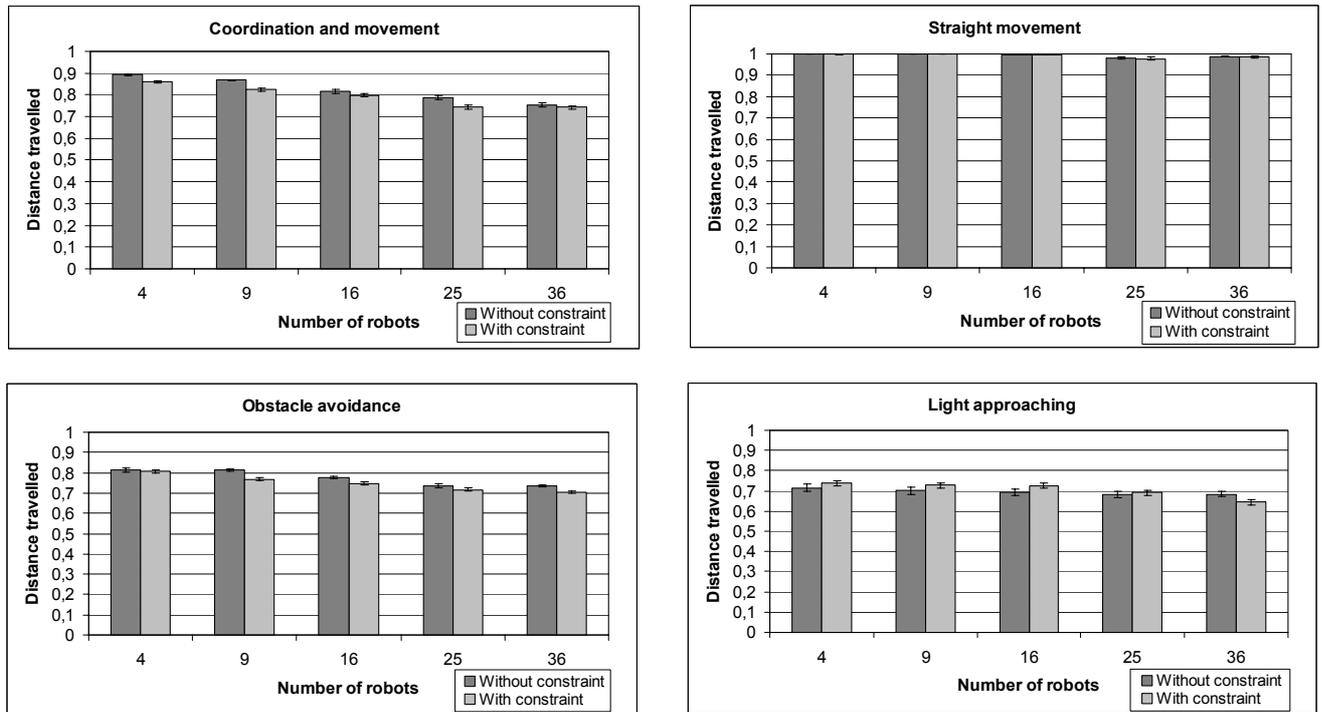


Figure 21: Performance with respect to the ability to: coordinate and move as quick as possible (top-left), move as fast and as straight as possible (top-right), avoid obstacles (bottom-left), and approach a light target (bottom-right). Histograms indicates the distance travelled by swarm-bots with different size, in four different conditions (see histograms’ titles and text). Dark and light grey histogram bars respectively indicate the performance of the robots without and with the turret-chassis joint constraint, and show the standard errors over the different trials (data refer to up to 100 trials per each condition since trials with performance below 0.3 were not consider).

The possibility of achieving these results can be ascribed to (a) the use of an evolutionary technique, and (b) a careful design of the robots’ hardware structure. The importance of the robots’ hardware structure can be explained by considering that the particular shape of the robots’ body, the type and positions of the sensors, and the way in which sensory information is encoded in the neural controllers might crucially affect the complexity of sensory-motor mapping that should be produced by the control systems of the robots and the possibility of integrating different behavioural abilities. These characteristics have been carefully designed and often re-designed on the basis of the results obtained by exploratory experiments performed in simulation. The importance of the latter point can be explained by considering that, as discussed elsewhere [6, 22], artificial evolution allows evolved robots to find solutions that exploit useful behavioural characteristics emerging from the interaction between

the control systems of the robots, the structure of their bodies, and the external environment, including the *social* environment made up by other robots. While in the work presented in this paper these designing activities have been carried out mainly on the basis of trial-and-error and engineering intuition, an alternative sounder method would be allowing the evolutionary process to co-evolve both the controller and the relevant aspects of the hardware [9].

The paper also presented a novel robust solution to tackle a difficult coordination problem generated by the hardware structure of the robots used, composed of a mobile chassis and a turret linked to other robots. This solution is interesting from an engineering perspective and might also furnish hints on the control solutions adopted by animals engaged in transporting objects alone or in groups.

This work has focussed on the coordination problems faced by some robots that are manually assembled by the

experimenter before being tested. Future work should focus on the interesting issue of how robots might self-assemble on the basis of the challenges of tasks, for example to overcome furrows and obstacles of different sizes and transporting objects with different sizes and weights [26].

Acknowledgements

This research has been supported by the “SWARM-BOTS” project funded by the Future and Emerging Technologies program (IST-FET) of the European Commission under grant IST-2000-31010.

References

1. Anderson, C., Theraulaz, G., & Deneubourg, J.-L. (2002). Self-assemblages in insect societies. *Insectes Sociaux*, 49, 1-12.
2. Arkin, R. C. (1989). Motor schema based mobile robot navigation. *International Journal of Robotics Research*, 8, 92-112.
3. Balch, T. (1998). *Behavioral Diversity in Learning Robot Teams*. Ph.D. Thesis. Atlanta, GE: College of Computing, Georgia Institute of Technology.
4. Baldassarre, G., & Nolfi, S. (in preparation). Synergies between evolved and hand-coded control in collective robotics.
5. Baldassarre, G., Nolfi, S., & Parisi, D. (2003). Evolution of collective behaviour in a group of physically linked robots. In G. Raidl, A. Guillot, & J.-A. Meyer (Eds.), *Applications of Evolutionary Computing - Proceedings of the Second European Workshop on Evolutionary Robotics* (581-592). Berlin: Springer Verlag.
6. Baldassarre, G., Nolfi, S., & Parisi, D. (2003). Evolving mobile robots able to display collective behaviour. *Artificial Life*, 9, 255-267.
7. Baldassarre, G., Trianni, V., Dorigo, M., & Nolfi, S. (in preparation). Coordinated motion and generalisation in groups of physically linked robots.
8. Bonabeau, E., Dorigo, M., & Theraulaz, G. (1999). *Swarm intelligence*. Oxford: Oxford University Press.
9. Bongard, J. C., & Pfeifer, R. (2002). A method for Isolating morphological effects on evolved behaviour. In Hallam, B., Floreano, D. (Eds.), *From Animals to Animats 7: Proceedings of the Seventh International Conference on the Simulation of Adaptive Behaviour* (305-311). Cambridge, MA: The MIT Press.
10. Brooks, R. A. (1986). A robust layered control system for a mobile robot. *Journal of Robotics and Automation*, 2, 14-23.
11. Camazine, S., Deneubourg, J. L., Franks, N. R., Sneyd J., Theraulaz, G., & Bonabeau, E. (2001). *Self-organization in biological systems*. Princeton, NJ.: Princeton University Press.
12. Cruse, H., Kindermann, T., Schumm, M., Dean, J., & Schmitz, J. (1998). Walknet - A biologically inspired network to control six-legged walking. *Neural Networks*, 11, 1435-1447.
13. Dorigo, M., Trianni, V., Sahin, E., Grob, R., Labella, T. H., Baldassarre, G., Nolfi, S., Deneubourg, J-L, Floreano, D., & Gambardella, L. M. (in press). Evolving self-organizing behavior for a swarm-bot. *Autonomous Robots*.
14. Grabowski, R., Navarro-Serment, L. E., & Khosla, P. K. (2003). An army of small robots. *Scientific American*, November, 42-47.
15. Greenberg, P. E. (2003). Tiny teamwork. *Nature*, 242, 134.
16. Holland O., Melhuish C. (1999). Stimergy, self-organization, and sorting in collective robotics. *Artificial Life*, 5, 173-202.
17. Ijspeert, A. J., Martinoli, A., Billard, A., & Gambardella, L. M. (2001). Collaboration through the exploitation of local interactions in autonomous collective robotics: The stick pulling experiment. *Autonomous Robots*, 11, 149-171.
18. Kube, R. C., & Bonabeau, E. (1998). Cooperative transport by ants and robots. *Robotics and autonomous systems*, 30, 85-101.
19. Martinoli, A. (1999). *Swarm intelligence in autonomous collective robotics: From tools to the analysis and synthesis of distributed control strategies*. PhD Thesis. Lausanne: Computer Science Department, Ecole Polytechnique Fédérale de Lusanne.
20. Miglino, O., Lund, H. H., & Nolfi, S. (1995). Evolving mobile robots in simulated and real environments. *Artificial Life*, 4, 417-434.
21. Mondada, F., Pettinaro, G., Guignard, A., Kwee, I., Floreano, D., Deneubourg, J-L, Nolfi, S., Gambardella, L. M., & Dorigo, M. (in press). Swarm-bot: A new distributed robotic concept. *Autonomous Robots*.
22. Nolfi, S., Floreano, D. (2001). *Evolutionary Robotics*. Cambridge, MA: The MIT Press.
23. Quinn, M., Smith, L., Mayley, G., & Husband, P. (2002). 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* (302-311). Cambridge, MA: MIT Press.
24. Reynolds, C. W. (1987). Flocks, herds, and schools: A distributed behavioral model. *Computer Graphics*, 21, 25-34.
25. Seyfarth, R. M., & Cheney, D. L. (2003). Signallers and receivers in animal communication. *Annual Review of Psychology*, 54, 145-173.
26. Trianni, V., Tuci, E., & Dorigo, M. (in press). Evolving Functional Self-Assembling in a Swarm of Autonomous Robots. *From Animals to Animats 8 - Proceedings of the 8th International Conference on the Simulation of Adaptive Behavior* (405-414). Cambridge, MA: MIT Press.
27. Von Hayek, F. A. (1945). The use of knowledge. *American Economic Review*, 35, 519-530.