Evolution of Collective Behavior in a Team of Physically Linked Robots

Gianluca Baldassarre, Stefano Nolfi, and Domenico Parisi

Institute of Cognitive Sciences and Technologies,
National Research Council of Italy (ISTC-CNR),
Viale Marx 15, 00137 Rome, Italy
{baldassarre, nolfi, parisi}@ip.rm.cnr.it

Abstract. In this paper we address the problem of how a group of four assembled simulated robots forming a linear structure can co-ordinate and move as straight and as fast as possible. This problem is solved in a rather simple and effective way by providing the robots with a sensor that detects the direction and intensity of the traction that the turret exerts on the chassis of each robot and by evolving their neural controllers. We also show how the evolved robots are able to generalize their ability in rather different circumstance by: (a) producing co-ordinated movements in teams with varying size, topology, and type of links; (b) displaying individual or collective obstacle avoidance behaviors when placed in an environment with obstacles; (c) displaying object pushing/pulling behavior when connected to or around a given object.

1 Introduction

How can a group of robots that are physically connected to form a single physical structure display coordinated movements? In this paper we consider a group of physically connected robots forming a linear structure that have to move as straight and as fast as possible in the environment. Given that the initial orientation of the tracks of each robot is randomly chosen, the robots should first negotiate a common direction and then move along that direction in a coordinated fashion.

This is one of the research problem we are facing within a project founded by CEC in which we are developing “swarm-bots” [2, 3, 5], i.e. a groups of robots (each called “s-bot”) that are able to self-assemble so as to form different physical structures and to cooperate in order to solve problems that cannot be solved individually.

Each s-bot has its own neural-network controller that generates motor outputs in response to sensory inputs. Since the s-bots are physically connected they need to coordinate in order to move together. As we will see, by providing the individual s-bots with traction sensors that detect the direction and intensity of the force that the turret exerts on the chassis and by utilizing an evolutionary technique [4], we were able to find a simple and effective solution: evolved s-bots display the ability to coordinate toward a unique direction that emerges from the negotiation between the individuals.

Moreover, we will show how neural controllers evolved for the ability to produce coordinated movements in a swarm-bot which includes four assembled s-bots forming a linear structure generalize to rather different circumstances. In particular, the evolved s-bots are able to: (a) produce coordinated movements in swarm-bots with varying size, topology, and type of links; (b) display individual and collective obstacle avoidance behaviors when placed in an environment with obstacles; (c) spontaneously produce object pushing/pulling behavior when connected to or around a given object.

2 Experiments and Results with a Simple Linear Formation

In order to investigate the problem described, we evolved the control systems of four physically linked s-bots forming a linear structure that were asked to move as fast and as straight as possible (Fig. 1). Experiments were run in simulation by constructing a software based on the rigid body dynamics simulator SDK Vortex^{TM}. This simulator reproduces the dynamics, friction and collisions between physical bodies.

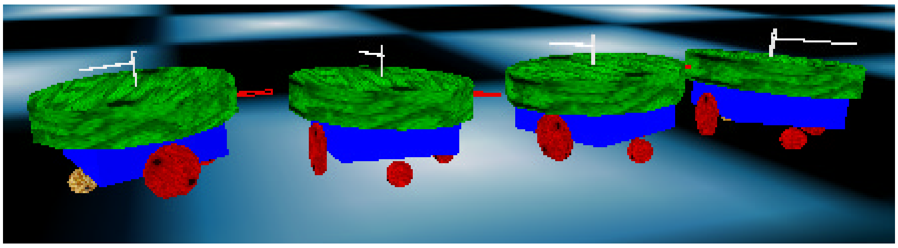

Fig. 1. Four physically linked s-bots forming a linear structure. For each s-bot, the cylinder and the parallelepiped respectively represent the turret and the chassis. The large disks and small spheres respectively represent the motorized and passive wheels. The line between two s-bots represents the link between them. The white line above each s-bot indicates the direction and intensity of the traction (see below)

Each s-bot (Fig. 1) consists of a rectangular chassis of size 3.5×3.5×1.0 cm provided with two motorised and two passive wheels with a width of 0.2 cm and a radius of 0.75 and 0.375 cm, respectively. Each s-bot is also provided with a cylindrical turret with a radius of 2.75 cm and a height of 1.0 cm that is connected to the chassis through a motorised "hinge joint" that can rotate around a vertical axis. Each s-bot has a physical link through which it can be attached to another s-bot along the perimeter of its turret. The link consists of another "hinge joint" that has a rotation axis parallel to the ground plane and is perpendicular to the line formed by the four s-bots. The density of the turret was set at 0.5 and the density of the other components at 1.0.

To speed up the simulations, we used a low gravitational acceleration coefficient (9.8 cm/s^2). This low value, that causes a low friction of the wheels on the ground, was compensated for by setting the maximum torque of the motors at a low value (see below). This allowed us to set the parameter that determines the granularity at which

Vortex approximates the differential equations used to simulate the bodies' dynamics at a rather high value, 0.1 (this parameter is interpreted as 100 ms).

Vortex uses a Coulomb friction model. Coulomb friction causes a force, opposite to the forces that cause the relative motion of objects, with an intensity μ N. μ is the friction coefficient and N is the total force, normal to the plane of contact, that depends on the forces applied to the bodies. μ is set at 0.6 and 0.0 respectively for the motorised and passive wheels. The activation of the motor neurons is used to control the motorised wheels by setting their desired velocity within the range [-10, +10] radians per second, and the maximum torque of the motors at 20 dynes-centimetre. The desired velocity applied to the turret-chassis motor is set on the basis of the difference between the activation of the left and right output neuron. If this difference is positive the chassis tends to rotate rightward (with respect to the turret and when seen from above), otherwise it tends to rotate leftward.

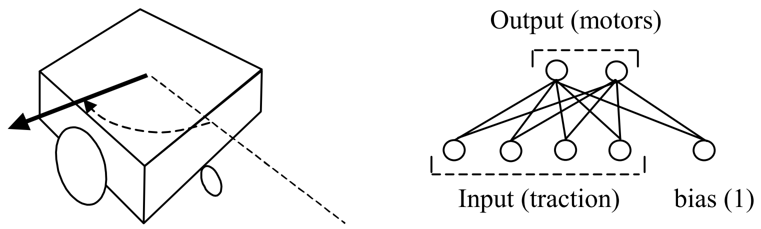

Fig. 2. *Left:* Traction force detected by an s-bot's traction sensor. The large and small circles respectively represent the right active wheel and front passive wheel. The dashed line and the full arrow respectively indicate the s-bot's orientation and the direction and intensity of the traction. The dashed arrow indicates the angle between the chassis' orientation and the traction. *Right:* The neural controller of the s-bots

Each s-bot is provided with a "traction sensor", placed at the turret-chassis junction, that returns the direction (i.e. the angle with respect to the chassis' orientation) and the intensity of the force of traction (henceforth called "traction") that the turret exerts on the chassis (Fig. 2). Traction is caused by the movements of both the connected s-bots and the s-bot's own chassis. Notice that the turret of each s-bot physically integrates the forces that are applied to the s-bot by the other s-bots. As a consequence the traction sensor provides the s-bot with an indication of the average direction toward which the team is trying to move as a whole. More precisely, it measures the mismatch between the directions toward which the entire team and the s-bot's chassis are trying to move. The intensity of the traction measures the size of this mismatch. From the point of view of each s-bot, this type of information is very easy to use for changing the direction of its own chassis in order to follow the rest of the team or to push the team to move toward a different desired direction.

Each s-bot's controller (Fig. 2, right) is a neural network with 4 sensory neurons that encode the traction. These neurons are directly connected with 2 motor neurons that control the two motorized wheels and the turret-chassis motorized joint. The 4 sensory neurons encode the intensity of the traction (normalized in [0.0, 1.0] on the basis of its maximum value registered in an experiment where the s-bots moved randomly) from four different preferential orientations with respect to the chassis (front, right, back and left). For each sensor, this intensity decreases linearly with respect to the absolute value of the difference between the sensor's preferential orientation and the traction's direction, and is 0 when this value is bigger than 90 degrees. The activation state of the motor units is normalized between [–10, +10] and is used to set the desired speed of the two corresponding wheels and the turret-chassis motor.

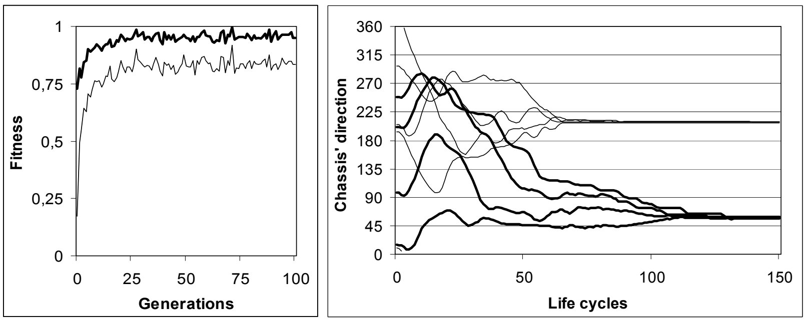

Fig. 3. *Left:* Performance across 100 generations. The bold line and thin line respectively represent the performance of the best team of each generation and the average performance of the population of teams, averaged over the 10 replications. *Right:* The graph shows the direction (angle) of the chassis of the four s-bots in two trials of 150 steps each, starting with two different initial random orientations (the bold and thin lines refer to the two trials)

The connection weights of the neural controller of the s-bots are synthesized by using an evolutionary technique [4]. The initial population consists of 100 randomly generated genotypes that encode the connection weights of 100 corresponding neural controllers. Each connection weight is represented in the genotype by 8 bits that are transformed in a number in the interval [–10, +10]. Therefore, the total length of the genotype is 10 (8 connection weights and two biases) × 8 = 80 bits.

As in the experiments reported in [1] and [6], we used a single-pool-single-genotype selection schema, i.e. we evolved a single population of genotypes each of which encodes the connection weights of a team of identical neural controllers. Each genotype is translated into four identical neural controllers ("clones") that in the experiments reported here correspond to a team of four s-bots forming a linear structure. As shown in [6], neural controllers might differentiate during lifetime as a result of their internal dynamics. However, in the experiments of this paper, as in the case of [1], neural controllers do not adapt online and do not have any internal state that can allow them to respond differently depending on their previous sensory experiences. Therefore, all s-bots always react in the same way to the same sensory states.

The team is allowed to "live" for 5 epochs each including 150 cycles. In each cycle, for each s-bot: (1) the activation state of the sensors is set according to the procedure described above; (2) the activation state of the two motor neurons is computed according to the standard logistic function; (3) the desired speed of the two wheels and of the motor controlling the motorized joint between the chassis and the turret are set according to the activation states of the motor units. At the beginning of each epoch the chassis of the four s-bots are assigned random orientations. However, to have a fair comparisons between different teams, all the teams of the same generation started with the same 5 randomly selected orientations in the 5 epochs. The 20 best genotypes of each generation were allowed to reproduce by generating 5 copies of their genotype with 3% of their bits replaced with a new randomly selected value. The evolutionary process lasted 100 generations. The experiment was replicated 10 times by starting with different randomly generated initial populations.

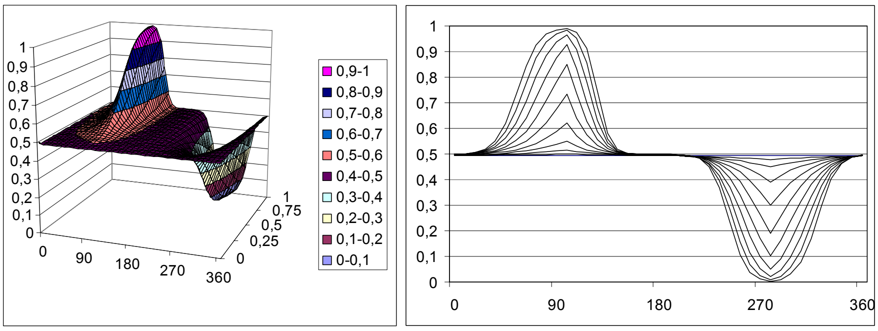

Fig. 4. *Left:* A typical evolved strategy (seed 4). Horizontal axis: angle of traction (0 is front, 90 right, 180 back, and 270 left). Depth axis: intensity of traction, between 0 and 1. Vertical axis: difference of the activation of the two neurons controlling the left and right wheels and the turret-chassis motor. When this value is around 0.5 the two wheels have the same speed (the s-bot is going straight), when it is above 0.5 the left wheel has a speed higher than the right wheel and the chassis turns right, and when it is below 0.5 the chassis turns left. *Right:* The graph shows the sections of the surface reported on the left graph, for different intensity levels

To force the teams of s-bots to move as fast and as straight as possible, we devised a fitness function based on the Euclidean distance between the barycentre of the team at the beginning and at the end of each epoch:

$$\text{Fitness} = ((x_0\text{-}x_{150})^2) * (y_0\text{-}y_{150})^2)^{1/2} / \text{MaxiDist} \qquad \mathbf{(1)}$$

where x_t and y_t are the coordinates, on the ground plane, of the average position occupied by the four s-bots at cycle t and MaxiDist is the maximum distance that a single s-bot can cover in 150 cycles by moving straight at maximum speed.

Fig. 3 shows how the fitness of the population, averaged over the 10 replications, changes across 100 generations. At the end of the evolution, the best team of each replication was tested for 100 epochs, and their average performance was: .947 .943 .931 .923 .839 .934 .765 .860 .946 .945.

Direct observation of behaviour shows that s-bots start to pull in different directions, orient their chassis in the direction where the majority of the other s-bots are pulling, move straight along this direction that emerges from this negotiation, and compensate successive mismatches in orientation that arise while moving. Fig. 4 shows a compact representation of the typical strategy of an evolved team. In particular, the figure shows the difference between the activation of the two motor neurons corresponding to the left and right wheels of an s-bot (normalized in the range [0.0, 1.0]) plotted against the angle and the intensity of the traction. The analysis of the figure and of the corresponding behaviour displayed by s-bots, indicates that evolved individuals adopt a simple strategy that can be described in the following way:

1) When the chassis of the s-bots are oriented toward the same direction, the intensity of the traction is null and the s-bots move straight with maximum speed.

2) When the traction intensity is low, the chassis of the s-bots are oriented toward similar but non-identical directions. In this case s-bots tend to turn toward the average direction in which the whole group is moving, i.e., they tend to turn left when the traction comes from the left side and right when the traction comes from the right side.

3) When the intensity of the traction is high and it comes from the rear direction, the chassis of the s-bots are oriented in rather different directions. For instance, three s-bots might be oriented toward north and one s-bot might be oriented toward south. In this case the s-bots tend to change their direction. The s-bots that have the higher mismatch with respect to the rest of the group feel a stronger traction than others, and this assures that a unique direction finally emerges for the whole team. In particular, in the example just described, the s-bot facing south will change its direction more quickly than the other three s-bots facing north. Notice that in this case all s-bots would feel a traction from the rear. The only difference between the s-bots is that the individual oriented toward south feels a traction intensity stronger than the other individuals. Aside from this schematic description, note that the non-linearities in how s-bots react to traction coming from different angles and of different intensities seem to play an important functional role that we are still trying to understand.

The right graph of Fig. 3 shows that the team direction that emerges from the s-bots' negotiation can be any possible one. This demonstrates that the strategy does not rely upon any type of alignment between the turret and the chassis. This is the key aspect of the traction sensor that allows the strategy to generalize to the wide variety of situations illustrated in the following sections.

3 Generalization

The strategy evolved with teams of four aligned s-bots illustrated in the previous section, generalises under many different conditions and exhibits a number of interesting emergent properties when placed in different environments. This suggests that this simple control strategy might be useful in a large number of cases.

3.1 Coordination in Swarm-bots with Varying Size, Topologies, and Links

To verify how this form of behaviour generalizes in new circumstances, we tested the control strategies evolved in the experiments described above on larger or smaller teams of s-bots (from 2 up to 10 s-bots) assembled so as to form a linear structure (recall that the different s-bots of a team have identical neural controllers, hence for each team made up by a given number of s-bots, we had to create the same number of clone controllers to guide them). The results of these tests show that s-bots maintain their ability to negotiate a single direction and to produce a coordinated movement along it independently of the size of the team (results not shown). In the case of two s-bots, in few tests, depending on the initial orientation of their chassis, s-bots fail to converge toward a single direction and circle around the group's centre. This situation is a dynamic equilibrium because the centripetal force caused by the link between the two s-bots causes a traction toward the group's centre that makes the s-bots turn toward it. This dynamic equilibrium has also been observed in other experiments (see Section 3.3). In the majority of cases, however, s-bots generalize successfully to variations of the team size.

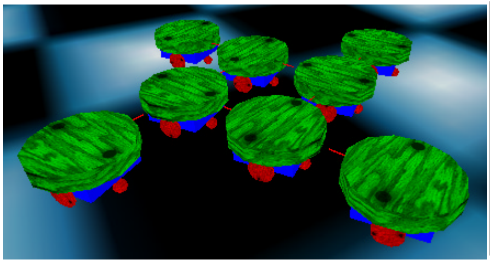
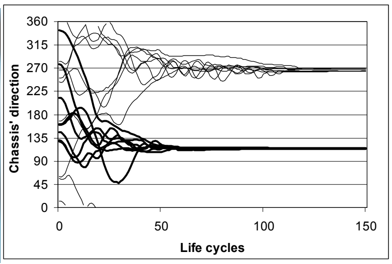

Fig. 5. *Left:* Eight s-bots connected by rigid links so as to form a star formation. *Right:* The orientation angle of the chassis of the eight s-bots of a star formation (bold lines) and snake formation (thin lines) in 150 steps

S-bots also display an ability to produce coordinated movements when assembled so as to form topologies that are different from the linear topology with which they have been evaluated during evolution. By testing a team of eight s-bots connected so as to form the star structure shown in the left part of Fig. 5, we observed that they displayed an ability to negotiate a unique direction and to move toward such emergent direction (see Fig.5, right).

Finally, s-bots also display an ability to produce coordinated movement when assembled by means of flexible instead of rigid links. Flexible links are made up by two segments connected by a hinge joint that allows the connected s-bots to rotate between them on the ground plane. By testing eight s-bots connected by flexible links so as to create a linear structure (snake formation), we observed that s-bots were able to negotiate a unique direction and to produce a coordinated movement along such direction also in this case. At the beginning of each trial the formation deformed as a consequence of the different orientation of the s-bots' chassis, but after some time it settled to a stable configuration and stable direction (Fig. 5, right). Given that in flexible assembled structures, the motor action produced by s-bots affects both the shape of the swarm and the traction perceived by the s-bots, these results seem to indicate that the evolved strategy is extremely robust and allows the s-bots to coordinate even when the traction sensors provide incomplete information about the movements of the team.

3.2 Individual and Collective Obstacle Avoidance

By placing s-bots in an environment with obstacles we observed that they display individual and collective obstacle avoidance behaviour. This can be explained by considering that collision with obstacles, by generating a traction force toward the direction opposite with respect to the direction of movement, might allow s-bots to turn away from obstacles.

In a first test, eight unconnected s-bots were placed in a squared arena surrounded by walls and including 4 large cylindrical obstacles (Fig. 6, left). In this situation, s-bots move straight when far from obstacles and turn avoiding obstacles and other s-bots when collisions occur (Fig. 6, right). In addition, after sometime all s-bots move in the arena following a clockwise or anticlockwise direction. This can be explained by considering that team of s-bots are provided with identical control systems and each evolved s-bots tend to avoid obstacles by turning left or right.

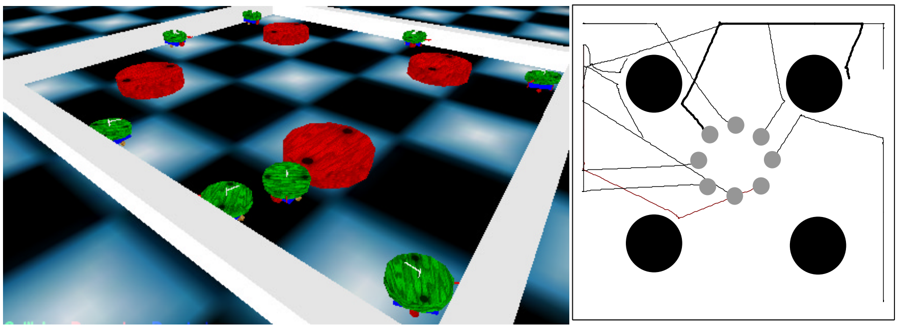

Fig. 6. Individual obstacle avoidance. *Left:* The arena with the eight s-bots, the walls, and the obstacles (larger cylinders). *Right:* Traces left by the s-bots (thin lines) in 300 cycles. The bold line is the trace left by an s-bot that hits an obstacle, the wall, a companion, and the wall again along its path

In a second set of tests, we placed assembled s-bots in the same environment surrounded by walls and including obstacles, and observed that swarm-bots displayed a collective obstacle avoidance behaviour independently from the topology with which they were assembled and the rigidity or flexibility of the connections used. Fig. 7 shows the behaviour of: a star formation connected with rigid links (top); a circular formation connected with flexible links (middle); a snake formation connected with flexible links (bottom). In all cases, the swarm-bots display an ability to coordinate and to collectively avoid walls. In the case of the two formations provided with flexible links, swarm-bot tend to change their shape after colliding with obstacles. However, given that they also tend to persevere in their direction of movement, they are also able to pass narrow passages eventually deforming their shape according to the configuration of the obstacles. In many cases (i.e. the control systems evolved in some of the replications of the experiments described in section 2) they never got stuck during the long time of observation (up to 4 hours: with the computer we used, Vortex performed about 30 cycles per second). This can be explained by considering that swarm-bots assembled through flexible links are dynamical systems that keep changing shape until they disentangle from the obstacles (small variations are also observed in the case of swarm-bot connected through rigid links due the fact that these might also bend of few degrees when subjected to significant forces). By moving s-bots change their relative positions with respect to other s-bots so that the whole swarm-bot always generates new configurations and has extremely reach dynamics.

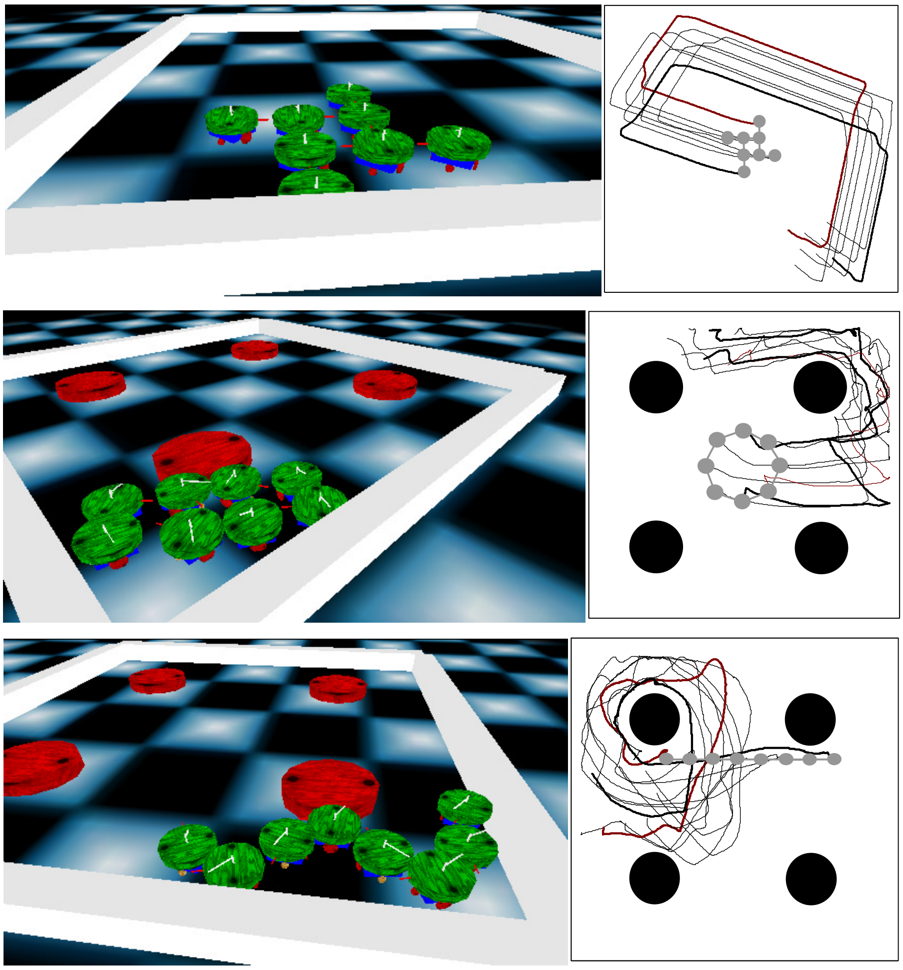

Fig. 7. Swarm-bots displaying collective obstacle avoidance. *Top*: A star formation assembled through rigid links. *Middle*: A circle formation with flexible links. *Bottom*: A snake formation with flexible links. *Left*: The light parallelepipeds are walls while the large cylinders are obstacles inside the arena. *Right:* The small gray circles and lines represent the initial shape of the swarm-bots. The large full circles are the obstacles inside the arena. The lines display the traces left by the s-bots showed in the left part of the figure during 600 cycles (bold lines highlight the traces left by two particular s-bots)

3.3 Collective Pushing and Pulling of Objects

We also observed that s-bots connected to an object, or connected so as to form a closed structure around an object, tend to pull/push the object in a coordinated fashion. This can be explained by considering that evolved s-bots tend to follow the direction of the team (for simplicity we will call this the "conformist tendency") but also have a tendency to persevere in their direction of movement if the intensity of the perceived traction is not too high (for simplicity we will call this the "stubborn tendency"). The stubborn tendency is due to the fact that when the direction of the perceived traction is 180 degrees different from the direction of the motion of the s-bots, they tend to go straight (see the flat area around 180 degrees in Fig. 4; the size of this area causing this stubborn tendency varies in different replication of the experiments reported in Section 2). Incidentally, the stubborn tendency might play a role in the ability to produce coordinated movements.

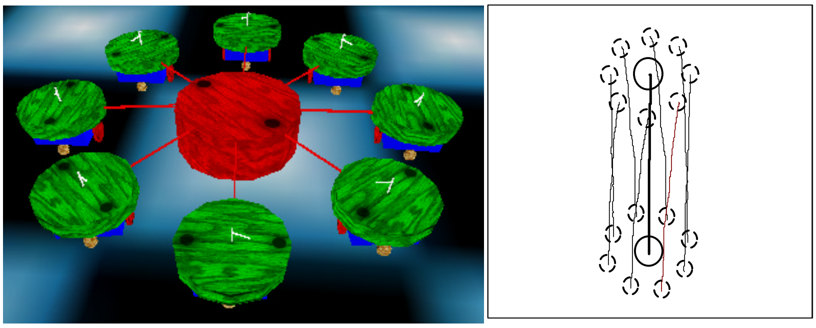

Fig. 8. *Left:* The ants swarm-bot formed by eight s-bots connected to the object with rigid links. *Right:* Traces left by the s-bots (thin lines) and the object (bold line) in 150 cycles. The dotted small circles represent the s-bots' initial (bottom) and final (top) position, the big circles represent the initial and final position of the object

Two sets of tests have been run. In the first set, the s-bots formed a circle and they were individually connected to a cylindrical object through a rigid link but not between them (this will be called the "ants formation", Fig. 8, left). The object had a radius of 4 cm, and a height of 3 cm. The density of the object was varied from 0.01 to 0.3, and the initial orientation of the chassis of the s-bots was set randomly. The test showed that when the weight of the object was small, s-bots were able to negotiate a common direction and to drag the object in that direction in the majority of the cases (cf. Fig. 8, right), while when the weight of the object was higher, s-bots tended to move in circle around it by displaying the behaviour described in Section 3.1. With a density of 0.3 the group succeeded in dragging the object only if the initial chassis' orientation was set to be the same for all s-bots, and only in 6 out of 10 cases (i.e. in 6 out of the 10 teams obtained by using the control strategies evolved in the 10 different replications of the experiments described in Section 2).

In the second set of tests, s-bots were assembled so as to form a circular structure around the object (Fig. 9, left). The results were similar to those obtained with the ants formation (Fig. 9, right). The only difference was that the formation deformed its shape so that some s-bots pushed the object while other s-bots pulled the other s-bots.

These results show that assembled s-bots evolved to move together as fast as possible also display an ability to coordinate through an external object to which they are attached and to pull/push the object in a coordinated fashion.

4 Conclusions and Directions of Future Research

We described how a group of simulated robots that are physically linked so as to form a single physical structure can display coordinated movements. We showed that the problem can be solved in a rather simple and effective way by providing the robots with a sensor that detects the direction and intensity of the traction that the turret exerts on the chassis and by evolving the neural controllers. The evolved strategy exploits the fact that the body of a swarm-bot (i.e. a robot constituted by a group of autonomous but aggregated individual robots) physically integrates the effects of the movements of the individual robots. Once individual robots are provided with traction sensors that allow them to detect the result of this integration, the problem of producing coordinated movements can be easily solved. In fact this sensors allow individual robots to have direct access to global information about what the entire group is doing.

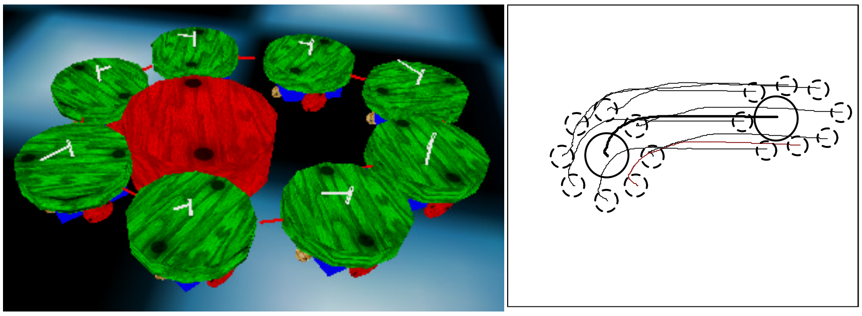

Fig. 9. *Left:* Eight s-bots assembled with flexible links forming a ring around the object. *Right:* Traces left by the s-bots (thin lines) and the object (bold line) in 150 cycles. The dotted small circles represent the swarm-bot's initial (left) and final (right) position, the big circles represent the initial and final position of the object

We also showed how neural controllers evolved for the ability to produce coordinated movement in a swarm-bot of four robots forming a linear structure are able to generalize in rather different circumstances: (a) they produce coordinated movements in swarm-bots with varying size, topology, and type of links; (b) they display individual or collective obstacle avoidance when placed in an environment with obstacles; (c) they spontaneously produce object pushing/pulling behavior when assembled to or around a given object. These results suggest that this strategy might constitute a basic functionality that, complemented with appropriate additional functions, might allow assembled robots to display a large number of interesting behaviors.

In future work we would like to evolve swarm-bots able to move toward a given target and to assemble and disassemble on the basis of their current goal and of the environmental conditions. From this point of view the results reported in this paper on individual and collective obstacle avoidance behavior suggest that the problem of controlling individual robots and teams of assembled robots might be solved with uniform and simple control solutions. Moreover, the results reported in the paper on the ability to generalize to rather different topologies of assembled robots, suggest that these control solutions might scale up to significantly complex setups.

Acknowledgments

This research has been supported by the SWARM-BOT project founded by the Future and Emerging Technologies program (IST-FET) of the European Community under grant IST-2000-31010. The information provided is the sole responsibility of the authors and does not reflect the Community's opinion. The Community is not responsible for any use that may be made of data appearing in this publication. The Swiss Government supported the Swiss project's participants (grant 01.0012).

References

1. Baldassarre, G., Nolfi, S., Parisi, D.: Evolving mobile robots able to display collective behaviour. In: Hemelrijk, C. K. (ed.): International Workshop on Self-Organisation and Evolution of Social Behaviour. Swiss Federal Institute of Technology, Zurich Switzerland (2002)
2. Mondada F., Floreano D., Guignard A., Deneubourg J.-L., Gambardella L.M., Nolfi S., Dorigo M.: Search for rescue: an application for the SWARM-BOT self-assembling robot concept. Technical Report, LSA2 - I2S – STI. Swiss Federal Institute of Technology: Lausanne, Switzerland (2002)
3. Mondada, F., Pettinaro, G.C., Kwee, I., Guignard, A., Gambardella, L.M., Floreano, D., Nolfi, S., Deneubourg, J.-L., Dorigo, M.: SWARM-BOT: A Swarm of Autonomous Mobile Robots with Self-Assembling Capabilities. In: Hemelrijk, C.K. (ed.): International Workshop on Self-Organisation and Evolution of Social Behaviour. Swiss Federal Institute of Technology, Zurich Switzerland (2002)
4. Nolfi, S., Floreano, D.: Evolutionary Robotics. The MIT Press, Cambridge Ma. (2000)
5. Pettinaro, G.C., Kwee, I., Gambardella, L.M., Mondada, F., Floreano, D., Nolfi, S., Deneubourg, J.-L., Dorigo, M.: SWARM Robotics: A Different Approach to Service Robotics. In Proceedings of the 33rd International Symposium on Robotics. International Federation of Robotics, Stockholm Sweden (2002)
6. Quinn, M., Smith, L., Mayley, G., Husband, P.: Evolving teamwork and role allocation with real robots. In Proceedings of the 8th International Conference on The Simulation and Synthesis of Living Systems (Artificial Life VIII) (In press)