Measuring Coordination as Entropy Decrease in Groups of Linked Simulated Robots

Gianluca Baldassarre, Domenico Parisi, Stefano Nolfi
Institute of Cognitive Sciences and Technologies
National Research Council of Italy (ISTC-CNR)
baldassarre@ip.rm.cnr.it, s.nolfi@istc.cnr.it, d.parisi@istc.cnr.it

1. Introduction

Within the field of collective robotics [Bonabeau et al. 1999, Martinoli 1999, Grabowski et al. 2003, Baldassarre et al. 2003 a] the implementation of *distributed coordination* of groups of cooperating robots that have to accomplish common tasks is receiving increasing attention [Ijspeert et al. 2001, Quinn et al. 2002, Dorigo et al. 2004]. When a group of robots has to accomplish a task, the *coordination* of the robots' behaviors has a critical importance since each robot has to take into account the time and space dependencies of its behavior with those of the rest of the group to contribute to the success of the group itself. These dependences require that the robots follow a common strategy to achieve the overall goal, accomplish some actions at the same time, synchronize cyclical behaviors, play different complementary roles, and so on. *Distributed* coordination of groups of robots refers to situations where the group is not centrally coordinated by one or more "leaders". In some tasks, distributed coordination might have advantages with respect to centralized coordination in terms of communication, robustness to failure of members, and scalability with respect to group size. For these reasons it is of interest for both engineering and scientific purposes.

Distributed coordination requires that decisions on coordination are made collectively. One way to accomplish this is through self-organization principles based on local interactions between the members of the group. For example, collective systems might coordinate cyclical activities through adjustable rhythmic mechanisms [Strogatz 2003], might avoid "overcrowding" through negative feedback, or might select a common particular strategy to achieve a goal through *positive feedback*, as we will discuss in this paper [Camazine et al. 2001]. Although most interesting work in distributed coordination of groups of robot has been based on self-organizing principles [Kube & Bonabeau 1998, Holland & Melhuish 1999, Ijspeert et al. 2001, Quinn et al. 2002], the self-organizing principles used have usually not been clearly isolated and quantitatively described. The goal of this paper is to study in detail the particular example of distributed coordination described in Baldassarre et al. (2004, in press) in order to identify the principles of self-organization at the basis of the observed collective behavior, and describe them quantitatively.

The distributed coordination problem studied here is a coordinated motion task. In this task, groups of physically linked simulated robots, each endowed with a very simple sensory system that allows each robot to sense the motion of linked robots, have to move away from an initial position in any direction as fast as possible. In this scenario, the success of the group depends on the capacity of its members to select a common direction of motion. Since the robots do not have a leader and do not possess dedicated communication channels, the common direction of motion has to emerge from the physical interactions of the robots themselves. Evolutionary techniques were used to evolve the neural-network controller of the robots to check if this automatic search process was capable of producing robots capable of harmonizing their motion on the basis of self-organizing principles. Previous work [Baldassarre et al. 2004, Baldassarre et al. in press] showed these groups of evolved robots are able to negotiate and converge toward a common direction of movement on the basis of a combination of a conformist tendency, that is a tendency to conform to the average direction of movement of the group, and of an autonomous tendency of each robot to move straight. In this paper we show that it is possible to measure the increasing organization of the group on the basis of Boltzmann entropy, and to use this measure to identify and describe the effects of the operation of the positive feedback mechanism that is at the basis of the observed evolved groups' behavior.

Section 2 presents the problem, the simulated robots, and the genetic algorithm used to evolve their neural controllers. Section 3 describes the functioning of the evolved controller at the individual level and the global behavior emerging from the interactions between individual behaviors. Section 4 illustrates how Boltzmann entropy can be used to measure the degree of organization of the group. Section 5 analyzes the collective behavior of groups of evolved robots by using this index. Finally, section 6 presents the conclusions.

2. Experimental setup

The scenario used for the experiments consists of a group of simulated robots (from 4 to 36, see Figure 5) set in an open arena. The robots are physically linked between them and have to harmonize their direction of motion in order to move together as far as possible from the initial position in a given amount of time.

The simulation of the robots was carried out with a C++ program based on Vortex™ SDK, a set of commercial libraries that allow programming realistic simulations of dynamics and collisions of rigid bodies in three dimensions. The simulation of each robot is based on the prototype of an hardware robot that is being built within the project SWARM-BOTS funded by the European Union [Mondada et al. 2003] (Figure 1). Each robot was composed of a cylindrical turret with a diameter of 5.8 cm and a chassis with two motorized wheels at the two sides and two caster wheels at the front and at the rear for stability. The simulated robot is half the size of the hardware robot: this decreases the weights of the simulated bodies and so allows increasing the simulation step of Vortex and decreasing the computational burden of the simulations (see below). The chassis was capable of freely rotating with respect to the turret through a further motor. This motor was activated on the basis of the difference of the activation of the motors of the two side wheels to ease the robots' turning while being physically linked to other robots (see Baldassarre et al. 2004 for details). The turret was provided with a gripper through which the robot could grasp other robots: this gripper was simulated through a rigid joint between the robots since our work focused on the behavior of groups of robots that were physically linked between them during the whole duration of the experiments. The gravitational acceleration coefficient was set at 9.8 cm/s^2 and the maximum torque of the wheels' motors was set at 70 dynes/cm. These low parameter settings, together with the small size of the robots, allowed using a relatively long integration time step in Vortex lasting 100 ms. This was desirable since simulations based on Vortex are computationally very heavy. The speed of the wheels was updated by the robots' controllers every 100 ms and could vary within ±5 rad/s.

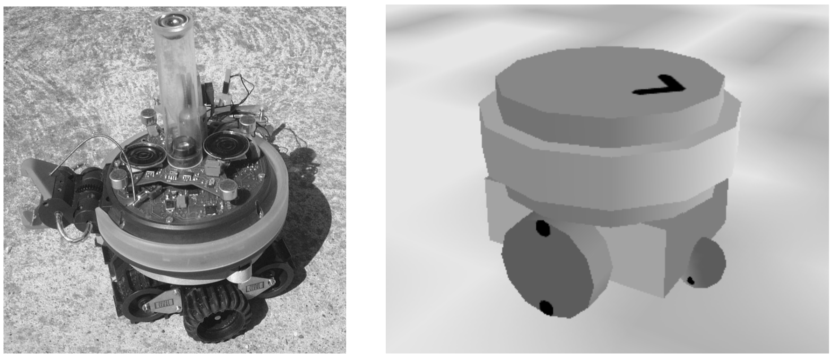

Figure 1. Left: The hardware prototype of an individual robot. Right: Each simulated robot is made up by a chassis to whom two motorized cylindrical wheels and two smaller caster wheels are attached (the visible dark-gray caster wheel marks the front of the chassis). The chassis supports a cylindrical turret (the arrow on the turret indicates its orientation).

Each robot had only a sensor, a special sensor called *traction sensor* (introduced for the first time in Baldassarre et al. 2003). This sensor was placed between the turret and the chassis. The sensor indicated to the robot the angle (with respect to the chassis orientation) and the intensity of the force that the turret exerted on the chassis. During the tests this force is caused by the physical interactions between the robots, in particular by the mismatch of the direction of movement of the chassis of the robot with respect to the movement of the robots attached to its turret. Notice that if one assumes a perfect rigidity of the physical links, the turrets and the links of the robots of the group form a whole solid body, so the traction measures the mismatch of movement between the robot's chassis and the rest of the group. Traction, seen as a vector, was affected by a 2D noise of ±5% of its maximum length. The state of the robots' traction sensor was updated every 100 ms.

The controller of each robot was a two-layer feed-forward neural network. The input layer was composed of four sensory units that encoded the traction force from four different preferential orientations with respect to the chassis's orientation (rear, left, front and right). When the angle was within ±90 degrees, each of these units had an activation proportional to the cosine of the angle between the unit's preferential orientation and the traction direction. With angles different from ±90 degrees, the units had a zero activation. The activation was also scaled by the intensity of traction. The last unit of the input layer was a bias unit that was constantly activated with 1. The two sigmoid output units were used to activate the wheels' motors by mapping their activation onto the range of the desired speed motor commands that varied in ±5 rad/s.

The connection weights of the neural controllers were evolved through an evolutionary algorithm [Nolfi & Floreano 2001]. Initially the algorithm created a population of 100 random genotypes. Each genotype contained a binary encoding of the ten connection weights of one neural controllers (the weights ranged over ±10). The neural controller encoded by a genotype was duplicated for a number of times equal to the number of robots forming a group, and these identical controllers were used to control the robots themselves (so the robots were "clones").

Groups of four robots connected to form a line were used to evolve the controllers. Each group was tested in five epochs each lasting 150 cycles (15 s). At the beginning of each epoch the robots were assigned random chassis' orientations. The 20 genotypes corresponding to the groups with the best performance of each generation were used to generate five copies each. Each bit of these copies was mutated (flipped) with a probability of 0.015. The whole cycle composed of these testing, selecting and reproducing phases was repeated 100 times (generations). The whole evolutionary process was replicated 30 times by starting with different populations of randomly generated genotypes. Notice that in this evolutionary algorithm one genotype corresponds to one robots' group, and the robots' groups compete and are selected as wholes (the group is the unit of selection of the genetic algorithm). This allows obtaining groups composed of highly cooperating individuals so avoiding the risk of the emergence of "free rider" individuals within them.

The genetic algorithm selected the best 20 genotypes (groups) of the population of each generation on the basis of a fitness criterion capturing the ability of the groups to move as straight and as fast as possible. In particular, the Euclidean distance covered by each group from the starting point was measured and averaged over the five epochs. To normalize the value of the fitness between [0.0, 1.0] the distance averaged over the five epochs was divided by the maximum distance covered by a single robot moving straight at maximum speed in 15 s (one epoch).

3. The evolved behavior

Figure 2 shows how the fitness of the best group and the average fitness of the population of 100 groups increase throughout the generations. Testing for 100 epochs the best groups of the last generation of each of the 30 evolution replications shows that the best and worst group have a performance of respectively 0.91 and 0.81. This means that all the evolutionary runs produce groups that are very good in coordinating and moving together.

Now the functioning of the evolved behavior will be briefly described at the individual level and then at the collective level. Overall, the behavior of single robots can be described as a "conformist behavior": the robots tend to follow the movement of the group as signaled by their traction sensors. Figure 3 shows more in detail the commands that the controller issues to the motors of the wheels in correspondence to different combinations of intensities and angles of traction. If a robot is moving towards the same direction of motion of the group, the robot perceives a zero or low traction from the front (around 180°): in this case the robot keeps moving straight. If the robots is moving in one direction and the groups moves towards its left hand side, it tends to perceive a traction from the left (around 90°) and as a consequence turns left. Similarly, if the robots is moving in one direction and the groups moves towards its right hand side, it tends to perceive a traction from the right (around 270°) and as a consequence turns right. Finally, if the robot moves in the opposite direction with respect to the group's movement, it perceives a traction from the rear (around 0°): in this case the robot tends to move straight, but since this is an unstable equilibrium state situated between the behaviors of turning left and right, the robot soon escapes it due to noise. It is important to notice that in some tests where the robots' chassis have particular initial orientations, the group starts to rotate around its geometrical center. This collective behavior is a stable equilibrium for the group since the robots perceive a slight traction towards the center of the group itself, that makes them to keep moving in circle around it. The experiments show that the stronger the symmetry of the group with respect to its center, the more likely that it falls into this stable state.

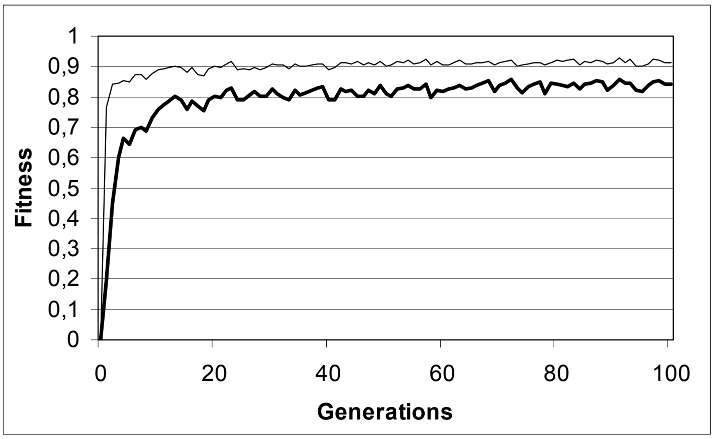

Figure 2. Fitness of the evolution runs throughout 100 generations. Thin line: fitness of the best group of each generation. Bold line: average fitness of the population. The two curves indicate the performance averaged over the 30 replications of the evolution process.

The observation of robots' individual behaviors suggests that the distributed coordination of robots performed by the evolved controller relies upon the self-organizing mechanism of positive feedback. In general, positive feedback allows a set of entities that form a system and have to select one option among many different available alternatives, to converge on the same choice without the need of centralized coordinators. In the case considered here, the robots forming a group have to select a common direction of motion without "leader robots" taking the decision. The behavior that underlies positive feedback is, at the individual level, "Conform to the behavior of the group". As we have seen, this behavior indeed emerged during the evolution illustrated above.

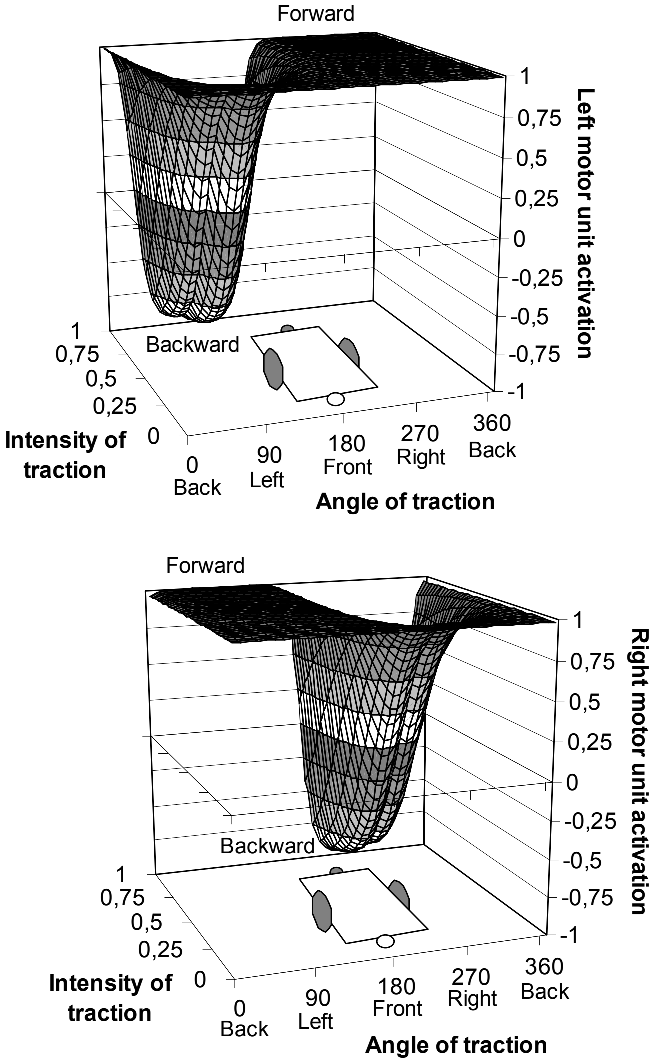

Figure 3. The graphs show how a robot's left motor (top) and right motor (bottom) react to a traction force with different angles and intensities. The vertical axis indicates the activation of the two motor units that, when considered together, indicate whether the robot goes straight, turns left or turns right, and with which speed. The schematic little picture represents a chassis and should aid the reader to understand how the robot perceives and reacts to traction: the white little wheel of the schematic chassis represents the rear of the chassis and corresponds to an angle of traction of 0°, measured clockwise.

When the evolved robots are tested together, one can observe that the robots start to pull and push in different directions selected at random. In fact initially there is a symmetry in the distribution of the motion directions over 360°. Chance causes some robots to move toward similar directions. If one of these random fluctuations eventually gains enough intensity, it break the initial symmetry: other robots start to follow those robots, and in so doing they further increase the traction felt by the non-aligned robots toward that direction. The whole group will hence rapidly converge toward the same direction of motion. What happens is that the positive feedback mechanism amplifies one of the initial random fluctuations so causing an avalanche effect that will rapidly lead the whole group to coordinate.

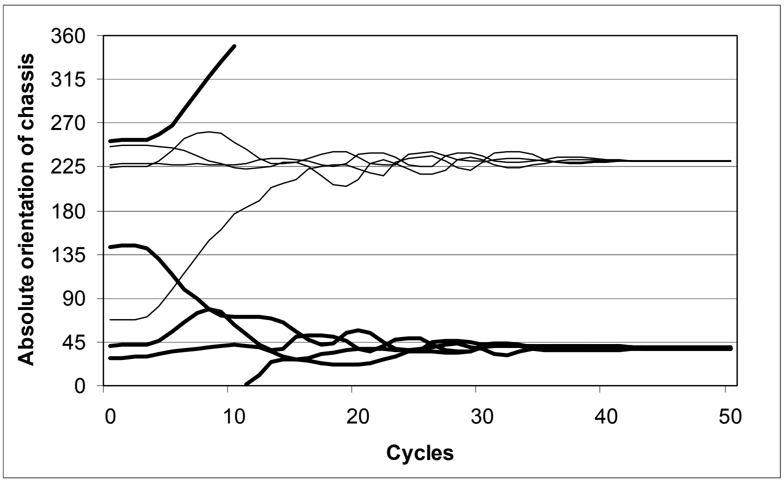

Figure 4. Absolute angles (with respect the environment) of the chassis' orientations of the four robots forming a group, in two tests where the initial orientations are randomly selected.

Before closing this section, it is important to notice that the common direction of motion that emerges in one coordinated motion test is the result of a collective decision that depends on the robots' initial random orientations and their fluctuations. Indeed, as shown in Figure 4, if the test is repeated more times the group's direction of motion that emerges is always different.

4. Self-organization as entropy decrease

This section introduces an index that measures the system's order based on Boltzmann entropy. This index is used in section 5 to measure the increase of order in groups of evolved robots engaged in coordinated motion tasks. The index has been chosen for its general applicability, as illustrated more in detail in the last section.

Boltzmann entropy refers to a system made up by N distinct elements that can assume a state in a given state space S. In our example the robots can assume a given motion direction among the possible orientations ranging over 360°. Since in our case the dimensions of S are continuous, the application of Boltzmann entropy requires to divide S in a given arbitrary number C of cells, each having a constant "volume". Furthermore, the index requires that we assume that the elements composing the system move randomly in S due to noise. As a consequence, at any time step an element can occupy any cell with a constant probability equal to 1/C.

The index is based on the definition of microstates and macrostates of the system. A *microstate* of the system corresponds to a situation where each of the N elements is in a particular state among the possible ones, that is it occupies a particular cell in S. Notice that since each of the N elements can be in one of C possible different states, the total number of different microstates is C^N. A *macrostate* of the system corresponds to a particular distribution of the elements over the cells, for example zero elements in the first cell, two elements in the second cell, one element in the third cell, etc. Each macrostate is usually composed of several possible microstates. For example, given a system with two elements and two cells, the macrostate with one element in each cell is composed of two microstates, one where the first element occupies the first cell and the second element occupies the second cell, and a second one where the first element occupies the second cell and the second element occupies the first cell. The other two macrostates, respectively with both elements either in the first or the second cell, are composed of only one microstate each. Boltzmann entropy E_m of a given macrostate m is defined as:

$$E_m = k \ln[w_m]$$

where w_m is the number of microstates that compose m, ln[.] is the natural logarithm and k is a scaling constant.

Since at any instant the probability of having any microstate is constant, the probability of a given macrostate is proportional to the number of microstates that compose it. In particular, the probability that the system falls into the macrostate m as a consequence of noise is w_m/C^N. The important point for Boltzmann entropy is that *noise* tends to spontaneously drive the system to macrostates composed of many microstates, that is, macrostates with high probability. For this reason, since Boltzmann entropy has a positive relationship with the number of microstates that compose the macrostate which the system is in, it can be used as a *measure of the disorder of the system* caused by the random forces operating on its composing elements. Notice that highly disordered macrostates correspond to situations where the elements of the system are equally distributed over the cells, hence to situations where the systems is highly *symmetric*.

Now consider the possibility that an *ordering mechanism* starts to operate on the elements of the system that previously were only subject to noise. An ordering mechanism is a mechanism that operates against noise by driving the system's elements towards *asymmetric conditions*, for example towards macrostates where the elements gather in few cells. The important property of ordering mechanisms is that they drive the system toward macrostates that are composed of few microstates, that is, macrostates with low levels of Boltzmann entropy. This implies that Boltzmann entropy can be used as an index to detect the presence of an ordering mechanism operating in the system, and to measure its efficacy in ordering the system.

In order to calculate the value of Boltzmann entropy for a particular macrostate m of the system, one has to compute the number w_m of microstates that compose it by using the following formula:

$$w_m = \frac{N!}{N_1!\ N_2!\ \ldots\ N_C!} \qquad \sum_{i=1}^{C} N_i = N$$

where N_i is the number of elements in the cell i, and "!" is the factorial operator. The formula relies upon the fact that there are $((N)(N-1)\ldots(N-N_1+1))/N_1!$ different possible sets of elements that can occupy the first cell, there are $((N-N_1)(N-N_1-1)\ldots(N-N_1-N_2+1))/N_2!$ different sets of elements that can occupy the second cell for each set of elements occupying the first cell, and so on. The formula of w_m is given by the multiplication of the number of these possibilities referring to the C cells. The index can then be computed as follows:

$$E_m = k \ln[w_m] = k \ln\left[\frac{N!}{N_1!\ N_2!\ \ldots\ N_C!}\right] = k\left(\ln[N!] - \sum_{i=1}^{C} \ln[N_i!]\right)$$

Since once N and C are given, the maximum entropy is equal to the entropy of the macrostate where the N elements are equally distributed over the cells, it is possible to set k to one divided the maximum entropy so as to normalize the index in [0, 1]:

$$E_m = k \ln[w_m] = \frac{1}{\ln\left[\frac{N!}{((N/C)!)^C}\right]} \ln[w_m] = \frac{1}{\ln[N!] - C\ln[(N/C)!]}\left(\ln[N!] - \sum_{i=1}^{C}\ln[N_i!]\right)$$

The calculation of the index can avoid computing the factorials, computation that soon becomes very heavy for increasing integers, by using Stirling's approximation:

$$\ln[n!] \approx \left(n + \frac{1}{2}\right)\ln[n] - n + \ln\left[\sqrt{2\pi}\right]$$

Stirling's approximation gives increasingly good approximations for integers n of increasing size (the error of approximation is already below 0.5% for $n > 20$).

Now some considerations on the particular evolution that the order of the group should follow if a positive feedback mechanism really underpins the self-organization of the robots engaged in a coordinated motion task are introduced. Note that order can be measured as the complement to one of the entropy normalized in [0, 1]. The order of the group should increase following a typically sigmoid-shaped evolution during the coordination of the group of robots. In fact initially the positive feedback mechanism should produce an increase of the organization of the group with a growing rate. After some time the natural exhaustion of the robots that converge to the common direction of motion should slow down the increase of organization and then stop it at the maximum level when all the robots are aligned. Note that exhaustion of elements is a general mechanism that, as negative feedback, usually works in couple with positive feedback mechanisms within natural systems, and has the power to slow down and then stop the operation of positive feedback mechanisms themselves [cf. Camazine et al. 2001]. The type of evolution just described can be modeled through the following first-order differential equation:

$$\frac{dO}{dt} = \left(k\,(1-O)\right)O$$

where O is the order of the group, dO/dt is the rate of change of order with respect to time (derivative), and k is a positive constant. Starting with level of O close to zero, (1-O) is close to one, so the change rate of the order is equal to k*O: this phase corresponds to the initial grow of order of the group where the size of the set of elements that change their selection is proportional to the size of the elements that have already chosen that selection. When O approaches one, (1-O) converges to zero so that also the rate of change converges to zero: this phase corresponds to the exhaustion of the non-aligned robots. The explicit solution of the differential equation is the following equation (the equation of the entropy E, equal to 1-O, is also reported here):

$$O = \frac{1}{1 + \left(\frac{1}{O_0} - 1\right) e^{-kt}} \qquad E = 1 - \frac{1}{1 + \left(\frac{1}{1 - E_0} - 1\right) e^{-kt}}$$

The next section uses the entropy index to monitor the level of disorder (entropy) of groups of robots engaged in a coordinated motion task, in order to determine if it actually follows the evolution predicted by the theoretical model presented here. This is done to evaluate, at a macro level, if a self-organizing principle based on positive feedback, together with the process of exhaustion of elements, indeed underlies the distributed organization of the groups of robots presented in section 3.

5. Distributed organization as entropy decrease

This section analyzes the dynamics of Boltzmann entropy in groups of robots guided by the evolved controller and engaged in the coordinated movement task.

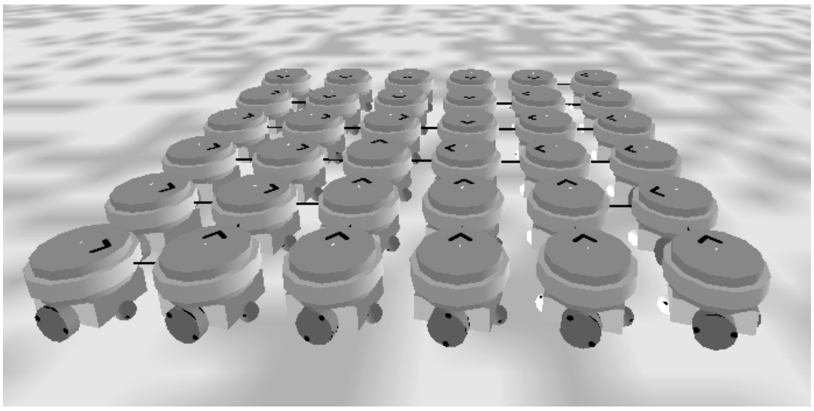

Figure 5. A group of 36 robots engaged in the coordinated motion task. The black segments between the turrets of two robots represent the physical connection between them.

The index introduced in section 4 has been used to measure the decrease of entropy in coordinated motion tests involving groups of 16, 25, and 36 linked robots that form a square lattice (Figure 5). To this purpose, the 360° possible angles of each robot's orientation was considered as the state space of the elements of the system (the robots). This angle was divided in eight cells of 45° each, with the 0° angle corresponding to 22.5° clockwise with respect to the absolute angle of one particular robot chosen as "pivot". While the origin of the angle on the basis of which the cells are computed is arbitrary, the selection done here assures that if the group achieves high coordination, the chassis' orientations of the robots will be located near to the center of the first cell and inside it (minimum entropy). Notice that since the pivot robot is always in the first cell, the number of microstates used to compute the entropy must be computed with respect to N-1 and not N. In order to normalize E_m in [0, 1], the scaling constant k of the index was set to one divided by the maximum value that $\log[w_m]$ can assume for the system considered here, corresponding to a (roughly) uniform distribution of the chassis' orientations over all cells. For example, in the case of groups formed by 36 robots, this value was set at:

$$k = 1 / \ln[35! / (5!\ 5!\ 5!\ 4!\ 4!\ 4!\ 4!\ 4!)] \approx$$
$$1 / \ln[7.509 * 10^{26}] \approx 1 / 61.8843 \approx 0.01615$$

Figure 6 reports the level of entropy measured in 20 coordinated motion tests run with robot's groups of 16, 25, and 36. The results indicate that, independently of the size of the group, the disorganization of the group initially decreases with an increasing rate, then tends to decrease with a decreasing rate, and it finally reaches a null value when all the robots have the same orientation.

To quantitatively evaluate if the entropy dynamics follows the evolution predicted by the mathematical model presented in section 4, a non-linear regression based on the model has been applied to the data of each of the 20 experiments run with the group formed by 36 robots. The regression uses the level of entropy as the dependent variable, the number of cycles as the independent variable, and k and E_0 as the parameters to estimate. The models based on the estimated parameters have been plotted in the last graph of Figure 6 (compare this with the third graph of the same figure). The values of the mean square error of the regressions, and their average, are reported in table 1. The table indicates that the models fit the data with small mean square errors (the average for the 20 regressions is 0.000265376), and this supports the hypothesis that a positive feedback mechanism underlies the self-organization of the robots engaged in the coordinated motion tests.

6. Conclusions

In this paper we tried to identify the principles at the basis of self-organization of groups of robots able to coordinate in a distributed fashion on the basis of local sensory information and actions. To this purpose, the paper presented a detailed analysis of the particular example of robots' distributed coordination reported in Baldassarre et al. (2004, in press). This study had the goal of clearly identifying and quantitatively describing the specific self-organizing mechanisms operating within those groups. This type analysis is rarely done in collective robotics research that exploits distributed coordination based upon self-organizing principles. More specifically, (a) the paper proposed a method for measuring the level of organization of groups of robots on the basis of Boltzmann entropy, and (b) it showed that the dynamics of the organization level of the group is in accordance with a dynamics originating from the operation of a positive feedback mechanism.

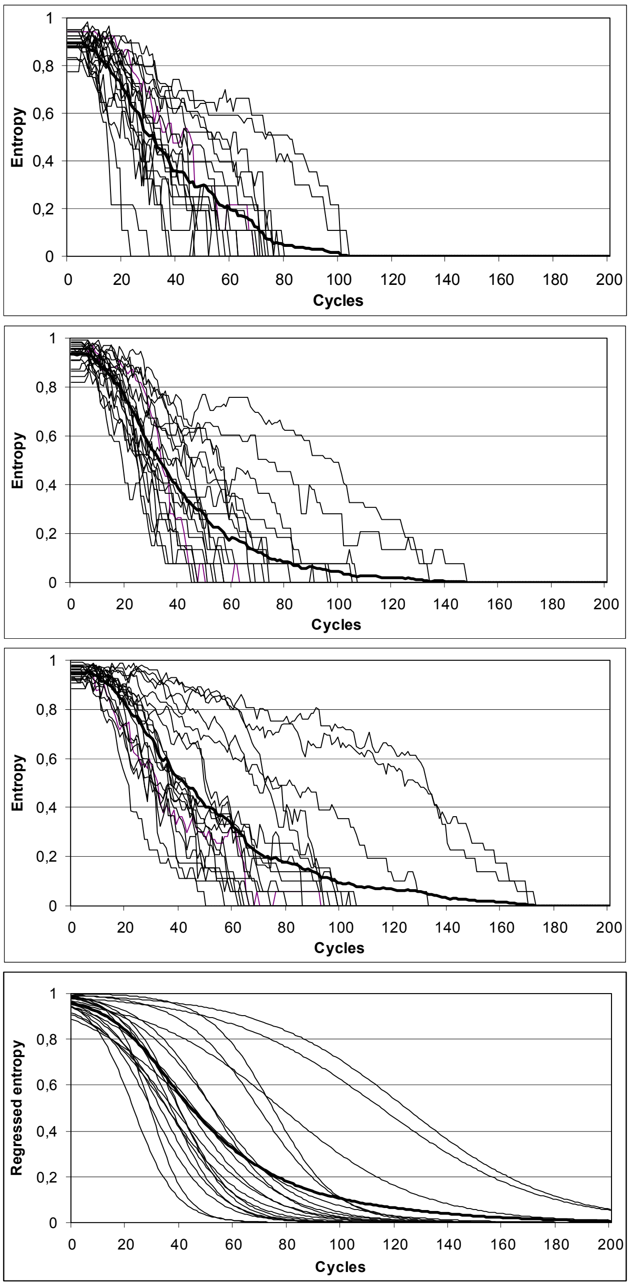

Figure 6. First three graphs: entropy of three groups respectively formed by 16, 25, and 36 robots engaged in a coordinated motion task. Each graph reports the entropy of 20 different tests (thin lines), and their average (bold line), run with different initial random orientations of the robots' chassis. Bottom: regressed models of the data relative to 20 replications of the coordinated motion test run with 36 robots (thin lines) and their average (bold line).

It is important to determine the generality and applicability of the method proposed here. The index based on Boltzmann entropy can be used to evaluate the effect of self-organizing principles that cause a set of elements composing a system to converge to the same choice among many available ones. The index can be used both in the case of continuous and discrete spaces, with any number of dimensions (notice that in the case of spaces with discrete dimensions, it is not necessary to arbitrarily divide the space into equal "cells" as done here). As mentioned, the index requires that in the absence of the ordering mechanisms, that is, the mechanisms that create the order captured by the index, the elements of the system tend to *homogeneously distribute* their choices over the whole space as a consequence of noise. The method has also the advantage of having a clear interpretation in terms of probabilities in that one can always determine the probability that the system falls into a particular macrostate if no ordering mechanisms are operating on it.

Table 1. The mean square errors of the regressions of the entropy of the 20 tests run with 36 robots.

1	**2**	**3**	**4**	**5**	**6**	**7**
0.001928676	0.000419204	0.000809705	0.001894163	0.001895273	0.001364410	0.001016100
8	**9**	**10**	**11**	**12**	**13**	**14**
0.004799802	0.000462556	0.000825678	0.001007695	0.002957842	0.000628843	0.001389428
15	**16**	**17**	**18**	**19**	**20**	**Average**
0.004460255	0.000698715	0.002092300	0.000744724	0.001398520	0.005107594	0.000265376

Notice that the approach proposed here can be used not only to measure *self*-organization of groups (of the type just described), but also organization in general. In fact the index is based on the effects produced by the action of the ordering mechanisms, and not on the nature of them. This implies, in particular, that the index can be used to measure the organization of groups also in the case of centralized organization, for example to measure the effectiveness of the action of a ship dog engaged in keeping a heard of randomly-moving sheep within a limited area.

A final remark concerns the relationship between organization and function. It is important to observe that while organization is usually a necessary condition for function, it is not sufficient for it. In fact function depends on the particular purposes of the system, established by the natural or artificial processes that select the system itself. This means that there might be self-organizing processes that are neutral with respect to the function of the system, or that are even detrimental for it. However, this does not diminish the importance of self-organizing principles for the functioning of collective systems since collective systems need to employ self-organizing principles to implement distributed organization, as done by the evolved groups of robots described in the present paper.

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. We thank Gianluca Massera for the discussions on the relationship between self-organization, organization and function, and Francesco Mannella for his help with data processing.

References

Baldassarre, G., Parisi, D., & Nolfi, S., 2004, Coordination and behavior integration in cooperating simulated robots, *From Animals to Animats 8 - Proceedings of the Eighth International Conference on the Simulation of Adaptive Behavior (SAB-2004)*, 385, MIT Press (Cambridge)

Baldassarre, G., Parisi, D., & Nolfi, S., in press, Distributed coordination of simulated robots based on self-organisation, *Artificial Life*.

Baldassarre, G., Nolfi, S., & Parisi, D., 2003, Evolution of collective behavior in a group of physically linked robots, *Applications of Evolutionary Computing - Proceedings of the Second European Workshop on Evolutionary Robotics*, 581, Springer Verlag (Berlin).

Baldassarre, G., Nolfi, S., & Parisi, D., 2003 a, Evolving mobile robots able to display collective behavior, *Artificial Life*, **255**, 9.

Bonabeau, E., Dorigo, M., & Theraulaz G., 1999, *Swarm intelligence,* Oxford University Press (Oxford).

Camazine, S., Deneubourg, J. L., Franks, N. R., Sneyd J., Theraulaz, G., & Bonabeau, E., 2001, *Self-organization in biological systems*, Princeton University Press (Princeton).

Dorigo, M., Trianni, V., Sahin, E., Gross, R., Labella, T. H., Baldassarre, G., Nolfi, S., Denebourg, J-L, Floreano, D., & Gambardella, L. M., 2004, Evolving self-organizing behavior for a swarm-bot, *Autonomous Robots*, **223**, 17.

Grabowski, R., Navarro-Serment, L. E., & Khosla, P. K., 2003, An army of small robots, *Scientific American*, November, 42.

Holland, O., & Melhuish, C., 1999, Stimergy, self-organization, and sorting in collective robotics, *Artificial Life*, **173**, 5.

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*, **149**, 11.

Kube, R. C., & Bonabeau, E., 1998, Cooperative transport by ants and robots, *Robotics and autonomous systems*, **85**, 30.

Martinoli, A., 1999, *Swarm intelligence in autonomous collective robotics: From tools to the analysis and synthesis of distributed control strategies*, PhD Thesis, Computer Science Department, Ecole Polytechnique Fèdèrale de Lusanne (Lausanne).

Mondada F., Guignard A., Bonani M., Floreano D., Bär D., Lauria M., 2003, SWARM-BOT: From Concept to Implementation, *Proceedings of the International Conference on Intelligent Robot and Systems (IROS 2003)*, 1626, IEEE Press (Piscataway).

Nolfi, S., & Floreano, D., 2001, *Evolutionary Robotics*, MIT Press (Cambridge).

Quinn, M., Smith, L., Mayley, G., & Husband, P., 2002, Evolving teamwork and role allocation with real robots, *Proceedings of the 8th International Conference on Artificial Life*, 302, MIT Press (Cambridge).

Strogatz S., 2003, *Sync: the emerging science of spontaneous order*, Hyperion (New York).