

Group Transport of an Object to a Target That Only Some Group Members May Sense

Roderich Groß and Marco Dorigo

IRIDIA - Université Libre de Bruxelles - Brussels, Belgium
{rgross,mdorigo}@ulb.ac.be

Abstract. This paper addresses the cooperative transport of a heavy object, called *prey*, towards a sporadically changing target location by a group of robots. The study is focused on the situation in which some robots are given the opportunity to localize the target, while the others (called the *blind* ones) are not. We propose the use of relatively simple robots capable of self-assembling into structures which pull or push the prey. To enable a blind robot to contribute to the group's performance, it can locally perceive traction forces, and whether it is moving or not. The robot group is controlled in a distributed manner, using a modular control architecture. A collection of simple hand-coded and artificially evolved control modules is presented and discussed. For group sizes ranging from 2 to 16 and different proportions of blind robots within the group, it is shown that controlled by an evolved solution, blind robots make an essential contribution to the group's performance. The study is carried out using a physics-based simulation of a real robotic system that is currently under construction.

1 Introduction

The transport of heavy objects by groups of robots can be motivated by low cost of manufacture, high robustness, high failure tolerance, or high flexibility – all desirable properties for a robotic system to have. Instead of using one robot powerful enough to transport the object (called *prey* hereafter) on its own, the task is accomplished in cooperation by a group of less powerful robots.

This paper addresses the situation in which the target location may be localized by some robots, while the others (the *blind* ones) cannot perceive it. In reality, such a constraint may exist due to various reasons. For a robot the target could be not visible because it is shadowed by the prey, or because of environmental conditions such as fog, steam or smoke. Or, it may happen that the robot's sensory system responsible for the perception of the target fails.

We propose the use of a group of relatively simple, autonomous robots (hereafter called *s-bots*) capable of self-assembling into structures which pull or push the prey. Thus, this paper addresses the self-organization of s-bots into assembled structures and the transport of heavy prey by groups of assembled s-bots. The robot group is controlled in a distributed manner. To enable a blind s-bot to contribute to the transport performance, it can locally perceive traction forces, and whether it is moving or not.

The paper is organized as follows. Section 2 provides a brief overview of some works related to this study. In Section 3, the simulation model is introduced. Section 4 details the different control policies that are used. Finally, in Sections 5 and 6 the results are discussed and conclusions are drawn.

2 Related Work

The field of distributed robotics has received growing attention by researchers within the last 15 years. Multi-robot systems have been studied in various topic areas and in different application domains [Parker, 2000]. Several works considered the cooperative transport of objects by a group of mobile robots. Some of these have been inspired by studies of social insect behavior.

Deneubourg et al. [1990] proposed the use of self-organized approaches for the collection and transport of objects by robots in unpredictable environments. Each robot unit could be simple and inefficient in itself, but a group of robots could exhibit complex and efficient behaviors. Cooperation could be achieved without any direct communication among robots [Grassé, 1959, Deneubourg and Goss, 1989]. Indirect communication is prevalent in the robotic system realized by Aiyama et al. [1999] in which two autonomous legged robots are carrying a common load in a pre-defined direction. In their system, the robots were communicating via the object to be moved. In order to synchronize its own actions with its teammate, each robot is provided with a sensor to measure the force exerted by the common load on the robot itself.

The use of force sensors for the coordinated motion of a group of pre-assembled robots has been studied by Dorigo et al. [2004]. In their system, the group was also able to move objects, if not too heavy, in an arbitrary direction.

Kube and Zhang [1993] studied a distributed approach to let a group of simple robots find and push a box towards a light. The box was too heavy to be moved by a single robot. Inspired by the observation that the behavior of ants during transport (e.g., changes in their spatial arrangement) can be associated with the detection of the stagnation of movement, Kube and Zhang [1995] extended their system with a stagnation recovery mechanism.

3 Simulator

3.1 Environment and Task

The simulated environment that we use in our experiments consists of a flat plane, a prey, and eight beacons which are uniformly arranged in a circle at a distance of 500 cm from a position called the *center*. The prey emits green light. It is modeled as a cylinder with a radius (in cm) in the range [5.8, 10], and a mass (in grams) in the range [200*N*, 300*N*], where *N* denotes the number of s-bots used. For each simulation, radius and mass are chosen randomly according to uniform distributions.




Fig. 1. The s-bot: (a) front, side and top view of the s-bot model and (b) the control scheme. For details see Section 3.2 and Section 4.

A simulation lasts 35 simulated seconds. At any point in time, there is only one beacon emitting white light indicating the target location. The beacon is selected randomly according to a uniform distribution once prior and twice during simulation (after 15 and 25 seconds). Initially, the prey is put in the center, and N s-bots are placed at random positions and with random orientations no more than 50 cm from the prey. In principle, all s-bots are able to sense the light-emitting beacon from any position in the environment. However, the controllers of the blind s-bots are not provided with the corresponding sensory information.

The s-bot controllers are supposed to let the s-bots localize and approach the prey, self-assemble into structures, each of them physically linked to the prey, and pull or push the prey towards the light-emitting beacon.

3.2 S-Bot Model

The s-bot model is illustrated in Figure 1 (a). It is an approximation of a real s-bot, currently under construction within the SWARM-BOTS project [Mondada et al., 2004, Dorigo et al., 2004, see also www.swarm-bots.org]. In the following, the actuators and sensors are detailed. Both the actuators and the sensors are affected by noise.

Actuators. The s-bot is equipped with two motorized wheels, a rotational base, a gripper element and an LED ring. Each motorized wheel can be controlled by setting the desired angular speed (in rad/s) in the range $[-8, 8]$. The s-bot's upper part (the turret) can rotate with respect to the lower part (the chassis) by means of a motorized axis. Therefore, the s-bot can align its turret with respect to its chassis to any angular offset in $[-\pi, \pi]$. The gripper element is modeled as a box heading forward with a small contact plate in the front. If the element is in *gripping* mode, a rigid connection will be established as soon as the contact plate touches the turret of another s-bot or the prey. The LED ring can emit green light. It can either be switched on or off. In this study, each s-bot uses its LED ring to signal whether an object is gripped or not.

Sensors. The simulated s-bot is equipped with a gripper status sensor, two proximity sensors, a rotational base sensor, a traction sensor, and an omnidirectional camera. These sensors have their counterparts on the real s-bot.

The gripper status sensor is able to detect whether the gripper is connected to another object or not. The proximity sensors measure the proximity of objects in the front left and front right directions. The rotational base sensor is able to perceive the angular offset between the turret and the chassis. The traction sensor provides an estimate of the magnitude and orientation of the horizontal component of the traction force that acts on the hinge joint between the turret and the chassis of the s-bot. This force is affected by the s-bot's actions and by the force exerted by all the objects that are physically linked to the s-bot.

The camera sensor is able to estimate the horizontal angle to the light-emitting beacon with respect to the s-bot's forward motion. The camera sensor can detect the presence and the coarse orientation of green objects in the front of the s-bot up to a distance of 60 cm, and 20 cm, respectively. In addition, the s-bot is able to detect whether it is moving or not. For the real s-bot, this ability can be implemented using the camera system, for example by comparing the position of light landmarks in its vicinity over time.

4 Control Policies

In a previous study on cooperative transport by a group of non-blind s-bots, we addressed the entire task utilizing a single neural network controller synthesized by artificial evolution [Groß and Dorigo, 2004]. We observed that the evolutionary process required a large number of generations to provide controllers that performed quite well, and the attained performance level differed essentially between the replications of the experiment.

On the contrary, in this paper we have decomposed the problem of control at the level of an s-bot into sub problems (see Figure 1 (b)): an s-bot can be in one of two main states as indicated by the status of its LED ring: if the s-bot has gripped the (green) prey or another s-bot, the LED ring emits a green light, otherwise it is switched off. Initially all s-bot are not connected and the rotational bases are locked. In this initial status, the *Assembling-Module* procedure is in charge of controlling the s-bot so that it locates a green object, approaches it and eventually establishes a physical connection with it. The s-bot is not allowed to grip objects other than green. Therefore, a green s-bot is, either directly or indirectly, always connected to the prey. Once a connection has been established, the gripper remains closed and the rotational base is unlocked. Unless the s-bot is *blind*, a simple but effective behavior for the robot is to align the chassis towards the target location and to set the speed of the wheels to the maximum value (*Move-Towards-Target* procedure). If the s-bot is *blind*, it locally perceives traction forces, and whether it is moving or not, and it determines the desired orientation of the chassis and a nonnegative speed for the wheels (*Transport-Module* procedure). In the following, several strategies implementing the *Assembling-Module* and *Transport-Module* procedures are detailed.

4.1 Hand-Coded Solutions, Lower and Upper Bounds

For the *Assembling-Module* procedure a parameterized, hand-coded controller has been developed: as long as no green object is perceived in its forward direc-

tion, the s-bot turns applying the speed vector $(s_1, -s_1)$ to the wheels. Once an object is perceived, and if it is more than 20 cm away, the s-bot approaches it at maximum speed. As soon as the object is close, the s-bot estimates whether it is more to the left or to the right, in order to approach it while turning in the appropriate direction (applying the speed vector (s_2, s_3) , or (s_3, s_2) , respectively to the wheels). Based on a training set of 200 trials with 2, 4, 6 or 8 s-bots, 1000 combinations of s_1 , s_2 and s_3 have been evaluated, and the one exhibiting the fastest average assembly time has been selected for the analysis.

For the *Transport-Module* procedure we implemented the following two simple policies to control the s-bots that are *blind*: the first one, in which the s-bots do not move their wheels nor their rotational bases; and the second one, in which the s-bots pull the objects gripped with maximum force.

As a lower bound to the performance of *Transport-Module* procedures for which blind s-bots contribute to the group's performance, we consider the case in which the blind s-bots are removed from the simulation. As an upper bound we consider the case in which the blind s-bots are replaced by non-blind ones.

4.2 Evolved Solutions

As an alternative to the hand-coded solutions, implementations of both modules have been obtained by artificial evolution. The evolutionary algorithm used is a self-adaptive $(\mu + \lambda)$ evolution strategy [Schwefel, 1975, Beyer, 2001]. The number of parents is $\mu = 20$ and the number of offspring is $\lambda = 80$.

An individual is composed of real-valued object parameters, and real-valued strategy parameters specifying the mutation strength used for each component of the object parameter vector. The object parameters encode the weights of two neural networks: a simple perceptron for the *Assembling-Module* procedure, and a simple recurrent neural network [Elman, 1990] with 4 hidden nodes for the *Transport-Module* procedure. In the latter case, having memory might be beneficial, since the s-bot can remember past sensory patterns (e.g., the readings from the traction sensor) that are not affected by its current motor actions.

The initial population is generated randomly. With a probability of 0.8 an offspring is created by mutation, otherwise two parent individuals are selected and recombined. In the latter case, the mutation operator is applied to the created offspring. The recombination operator acts on both, object and strategy parameters, and combines two genotypes a) by swapping the parameter sets belonging to the first or second neural network, b) by intermediate recombination, or c) by dominant recombination [Beyer, 2001], each with the same probability.

In the following, the fitness function is described.

Fitness. An individual represents a common controller which is cloned and uploaded to each s-bot in the group. Its fitness is evaluated by performing five tests of 35 simulated seconds each; the group size $N \in [2, 3, 4, 5]$ and the number of blind s-bots in the group $N_B \in [1, 2, \dots, \lceil N/2 \rceil]$ are chosen randomly according to uniform distributions. The fitness score is computed as the average of the quality values observed in the five trials.

The quality Q exhibited in a trial is defined as $Q = cA + (1 - c)B$, where $c = \frac{1}{5}$, A measures the overall assembling performance, and B measures the transport performance of the blind s-bots. Q can be computed based solely on information the s-bots perceive locally during simulation. A is the average of the assembling performance of all s-bots for all control steps, while B is the average of the transport performance of all blind s-bots for a subset of all control steps: only control steps happening in the 6 seconds preceding a target change or the end of the simulation are taken into account, as during the other control steps the s-bots are busy assembling or adapting to the current target location¹.

The assembling performance rewards s-bots for approaching and gripping green objects: $A^t(j) = 1$, if s-bot j grips a green object at time t , and otherwise

$$A^t(j) = \begin{cases} \frac{1}{2} + \frac{1}{4}(1 - h^t(j)) & \text{if s-bot } j \text{ senses a green object } (\leq 20 \text{ cm}); \\ \frac{1}{4} & \text{if s-bot } j \text{ senses a green object } (> 20 \text{ cm}); \\ 0 & \text{otherwise;} \end{cases} \quad (1)$$

where $h^t(j) = \min(1, \alpha^t(j)/\frac{\pi}{6})$, and $\alpha^t(j)$ denotes the angular deviation (in rad) of s-bot j 's heading with respect to the green object perceived at time t .

The transport performance is defined as $B^t(j) = 0$ if s-bot j is not gripped to any object at time t . Otherwise $B^t(j) = B_1^t(j) + B_2^t(j)$, where $B_1^t(j) = \frac{1}{2} \max(0, o^t(j))$ rewards the s-bot for having its chassis oriented towards the target², while $B_2^t(j) = \frac{1}{4} + \frac{1}{4}s^t(j)o^t(j)$ additionally takes into account the setting $s^t(j) \in [0, 1]$ for the nonnegative speed of the left and right wheels.

5 Results

Ten independent evolutionary runs have been performed for 300 generations each. A single run lasts almost one week on a cluster of 10 computers, each one equipped with an AMD Athlon XPTM 2400+ processor.

Figure 2 shows the development of fitness values for the ten runs performed. The fitness values are very noisy. It is hard for a human observer to assess the attained quality level of the whole transport system by using the given fitness function. Therefore, we introduce another quality measure, the *gained distance*, defined as the difference between the initial and final distances of the prey from the light-emitting beacon. Since the target location changes twice, we take as performance measure the sum of the corresponding three gained distances. Using this new measure, we post-evaluated the $\mu = 20$ best individuals of the final generation of each evolutionary run on a sample of 200 different trials. For each evolution, we consider the individual exhibiting the highest average performance to be the best. We observed that the best individuals of all evolutionary runs exhibit a similar average performance (on average 151 with standard deviation 7.6). In the following, we focus on the performance of the best evolved individual.

Figure 3 (a) shows the performance by means of the gained distance observed for 4, 8, 12 and 16 s-bots, all of which are not blind. Hand-coded and evolved

¹ An s-bot may take up to 2 seconds to realign its rotational base.

² $o^t(j) \in [-1, 1]$; the value 1 (-1) refers to alignment with the target (opposite to it).




Fig. 2. Box-and-whisker plot providing information about the fitness development in ten evolutionary runs. Each box comprises observations ranging from the first to the third quartile. The median is indicated by a bar. The whiskers extend to the farthest data points that are within 1.5 times the interquartile range. Outliers are indicated as circles. The bold line at fitness level 0.5 marks an upper bound to the average performance of robotic systems in which blind s-bots do not contribute to transport: once an s-bot is connected, having its chassis oriented in a random direction and pulling with maximum speed, the expected reward the s-bot receives is 0.5 (see Section 4.2).

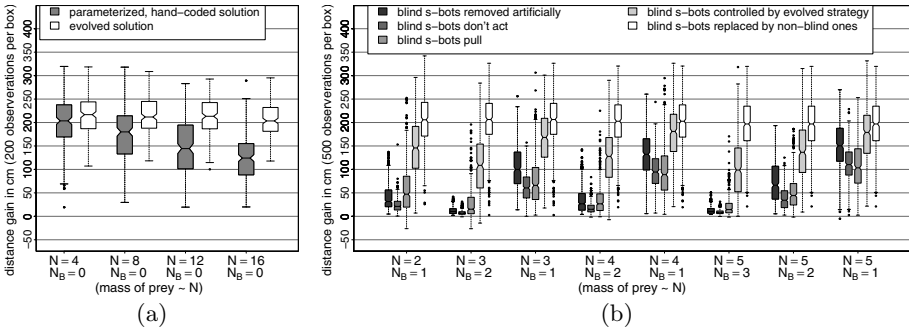


Fig. 3. Performance of solutions for (a) the *Assembling-Module* procedure for groups of 4 to 16 non-blind s-bots, and (b) the *Transport-Module* procedure for N s-bots of which N_B are blind; to assemble, the hand-coded *Assembling-Module* procedure is used.

implementations of the *Assembling-Module* procedure are compared; since no blind s-bots are present, the *Transport-Module* procedure is not required. It can be seen that the evolved solution scales better with group size. The bigger the group, the higher the density of s-bots. Thus, the better performance of the evolved controller might be due to an efficient use of the front left and front right proximity sensors. An s-bot controlled by the hand-coded solution avoids only obstacles shadowing green objects in its forward direction.

Figure 3 (b) shows the performance of several transport strategies for those combinations of group size and number of blind s-bots that occur during evolution. In every case, the parameterized, hand-coded solution with the most successful parameter set is used as *Assembling-Module* procedure.

The first, and the last, box in each group represent the lower, and upper, bounds (respectively) for the performance of *Transport-Module* procedures in which blind s-bots contribute to the group’s performance. Blind s-bots are either

artificially removed from the simulation (first box) or replaced by non-blind ones (last box). Thus, the first, and the last, box display to what extent $N - N_B$, or N , non-blind s-bots have transported the prey by pulling with the maximum speed towards the target. By looking at the figure, we observe that in the scenarios in which the percentage of blind s-bots is more than 50% (groups 2 and 6 from left to right) $N - N_B$ s-bots are nearly incapable of moving the prey.

The second and third boxes in each group (from left to right) correspond to the hand-coded solutions for the *Transport-Module* procedure. During transport, the blind s-bots either don't act (second box) or pull with maximum force (third box). In general, if the non-blind s-bots are not powerful enough to overcome the prey's resistance to motion, blind s-bots pulling in random directions can contribute to the group's performance (on average). However, the higher the number of blind s-bots, the less likely it is that all their pulling forces are properly directed. On the other hand, if the non-blind robots are capable of transporting the prey, the blind s-bots, because of their lack of knowledge about the target location, may end up being an additional weight to be carried, determining a decrease in performance.

The fourth box refers to observations in which the *Transport-Module* procedure of the best evolved solution has been used. The performance of this controller clearly outperforms the hand-coded solutions. With this controller blind s-bots make an essential contribution to the group's performance.

So far we have shown that we can control blind s-bots so that they contribute to the transport. However, to what extent there is an advantage by using blind s-bots seems to depend on the values of the parameters N and N_B . Let $P(x, y, z)$ be the performance of a group of x s-bots of which y are blind and whose task is to transport a prey of mass $z \cdot m$ (in grams), where m is chosen uniformly random in [200, 300]. Given the group size N and the number of blind s-bots N_B , we can define the relative system performance as

$$\text{RSP}(N, N_B) = \frac{P(N, N_B, N)}{P(N, 0, N)}. \quad (2)$$

In other words, $\text{RSP}(N, N_B)$ is the ratio between the performance of N s-bots of which N_B are blind and the performance of N non-blind s-bots. In addition, we define the contribution factor of blind s-bots as

$$\text{CF}(N, N_B) = \frac{P(N, N_B, N) - P(N - N_B, 0, N)}{P(N, 0, N) - P(N - N_B, 0, N)}. \quad (3)$$

$\text{CF}(N, N_B)$ is the ratio between the contribution of N_B blind s-bots and the contribution that N_B non-blind s-bots would provide if put in the same situation.

Using the mean distance gain over 200 trials as performance measure, Table 1 lists the relative system performance and the contribution factor of blind s-bots for the case in which both, the *Transport-Module* procedure and the *Assembling-Module* procedure are specified by the genotype of the best evolved individual. It can be seen that blind s-bots contribute to the system's performance unless all s-bots are blind (last column). In all cases in which no more than half of

Table 1. Relative system performance $RSP(N, N_B)$ and contribution factor of blind s-bots $CF(N, N_B)$ (both expressed as percentages) for the best evolved solution, and for different group sizes (N) and different numbers of blind s-bots (N_B).

N_B :		0	1	2	3	$\frac{1}{4}N$	$\frac{2}{4}N$	$\frac{3}{4}N$	N
$N = 4$	RSP	100.00	80.54	64.53	40.57	80.54	64.53	40.57	0.19
	CF	100.00	50.20	57.88	39.14	50.20	57.88	39.14	0.19
$N = 8$	RSP	100.00	93.86	85.01	72.56	85.01	58.91	24.87	-0.43
	CF	100.00	60.41	55.66	53.97	55.66	50.93	22.73	-0.43
$N = 12$	RSP	100.00	94.70	89.01	84.48	84.48	56.95	22.09	-0.70
	CF	100.00	50.54	46.15	50.45	50.45	48.37	19.14	-0.70
$N = 16$	RSP	100.00	96.90	95.34	91.39	82.88	56.96	20.80	0.10
	CF	100.00	51.70	54.29	58.06	40.42	47.41	16.85	0.10

the s-bots are blind, the contribution of the blind s-bots is 40 to 60% of the contribution non-blind s-bots would provide (in average).

6 Conclusions

In this paper a robotic system composed of relatively simple robots capable of self-assembling is proposed for the cooperative transport of a heavy prey. In principle, the advantage of our robotic system with respect to robotic systems used in previous studies in cooperative transport is twofold. First, due to the ability to establish physical links among s-bots, the system is applicable to scenarios in which the number of robots necessary to push/pull the prey is higher than the number of robots which can directly manipulate it. Second, even if some of the robots are blind (i.e., they cannot localize the target), they can contribute to the overall system performance.

We assessed the performances of simple hand-coded and evolved solutions and we put them in relation to lower and upper bounds. Unless no robot is able to localize the target, having blind robots controlled by an evolved individual is superior in performance to the alternative of removing them from the experiment. On the contrary, in most of the trials, controlling the blind s-bots by the hand-coded solutions is inferior in performance to the alternative of removing them from the experiment. If no more than half of the robots are blind, the contribution of the blind robots, if controlled by the best evolved controller, is 40 to 60% of the contribution non-blind robots would provide (in average) if put in the same situation.

For the best evolved solution the performance scales well with group size, making it possible to transport heavier prey by larger swarms of robots.

Acknowledgments

This work was supported by the Belgian FNRS, of which Marco Dorigo is a Research Director, via the grant “Virtual Swarm-bots” (contract no. 9.4515.03);

by the Scientific Research Directorate of the French Community of Belgium via an “Action de Recherche Concertée” (the “ANTS” project); and by the Future and Emerging Technologies programme (IST-FET) of the European Commission, via the “SWARM-BOTS” project (grant IST-2000-31010). The information provided is the sole responsibility of the authors and does not reflect the Community’s opinion. The Community is not responsible for any use that might be made of data appearing in this publication. The authors also wish to thank Tom Lenaerts and all the members of the “SWARM-BOTS” project.

References

- Y. Aiyama, M. Hara, T. Yabuki, J. Ota, and T. Arai. Cooperative transportation by two four-legged robots with implicit communication. *Robotics and Autonomous Systems*, 29:13–19, 1999.
- H.-G. Beyer. *The Theory of Evolution Strategies*. Springer, Berlin, Germany, 2001.
- J.-L. Deneubourg and S. Goss. Collective patterns and decision-making. *Ethology, Ecology and Evolution*, 1:295–311, 1989.
- J.-L. Deneubourg, S. Goss, G. Sandini, F. Ferrari, and P. Dario. Self-organizing collection and transport of objects in unpredictable environments. In *Proc. of Japan – U.S.A Symp. on Flexible Automation*, pages 1093–1098. ISCIE, Kyoto, Japan, 1990.
- M. Dorigo, V. Trianni, E. Şahin, R. Groß, T.H. Labella, G. Baldassarre, S. Nolfi, J.-L. Deneubourg, F. Mondada, D. Floreano, and L.M. Gambardella. Evolving self-organizing behaviors for a swarm-bot. *Autonomous Robots*, 17(2–3):223–245, 2004.
- J. L. Elman. Finding structure in time. *Cognitive Science*, 14:179–211, 1990.
- P.-P. Grassé. La reconstruction du nid et les coordinations inter-individuelles chez *Bellicositermes natalensis* et *Cubitermes sp.* La théorie de la stigmergie : essai d’interprétation du comportement des termites constructeurs. *Insectes Sociaux*, 6: 41–81, 1959.
- R. Groß and M. Dorigo. Cooperative transport of objects of different shapes and sizes. In *Ant Colony Optimization and Swarm Intelligence – Proceedings of ANTS 2004 – Fourth International Workshop*, volume 3172 of *Lecture Notes in Computer Science*, pages 107–118. Springer Verlag, Berlin, Germany, 2004.
- C. R. Kube and H. Zhang. Collective robotics: from social insects to robots. *Adaptive Behaviour*, 2(2):189–218, 1993.
- C. R. Kube and H. Zhang. Stagnation recovery behaviours for collective robotics. In *1994 IEEE/RSJ/GI Int. Conf. on Intelligent Robotics and Systems*, pages 1883–1890. IEEE Computer Society Press, Los Alamitos, CA, 1995.
- F. Mondada, G. C. Pettinaro, A. Guignard, I. V. Kwee, D. Floreano, J.-L. Deneubourg, S. Nolfi, L. M. Gambardella, and M. Dorigo. SWARM-BOT: A new distributed robotic concept. *Autonomous Robots*, 17(2–3):193–221, 2004.
- L. E. Parker. Current state of the art in distributed autonomous mobile robotics. In *Distributed Autonomous Robotic System 4*, pages 3–12. Springer, Tokyo, Japan, 2000.
- H.-P. Schwefel. *Evolutionsstrategie und numerische Optimierung*. PhD thesis, Technische Universität Berlin, Germany, 1975.