

Evolving a Cooperative Transport Behavior for Two Simple Robots

Roderich Groß and Marco Dorigo

IRIDIA, Université Libre de Bruxelles
Avenue Franklin D. Roosevelt 50
CP 194/6, 1050 Bruxelles, Belgium
{rgross,mdorigo}@ulb.ac.be

Abstract. This paper addresses the problem of cooperative transport of an object by a group of two simple autonomous mobile robots called *s-bots*. S-bots are able to establish physical connections between each other and with an object called the prey. The environment consists of a flat ground, the prey, and a light-emitting beacon. The task is to move the prey as far as possible in an arbitrary direction by pulling and/or pushing it. The object cannot be moved without coordination. There is no explicit communication among s-bots; moreover, the s-bots cannot sense each other. All experiments are carried out using a 3D physics simulator.

The s-bots are controlled by recurrent neural networks that are created by an evolutionary algorithm. Evolved solutions attained a satisfactory level of performance and some of them exhibit remarkably low fluctuations under different conditions. Many solutions found can be applied to larger group sizes, making it possible to move bigger objects.

1 Introduction

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 social insects such as ants.

Almost half a century ago, Sudd (1960) studied the transport of prey by single ants and by groups of ants of the species *Pheidole crassinada*. Although he observed that single ants would mostly behave similar to those engaged in group transport, he reported that group transport “showed co-operative features” (Sudd, 1960).

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 by itself, but a group of robots would have a complex and efficient behavior. Cooperation could be achieved without any direct communication among robots, as in some biological systems that rely on interactions via the environment or that exhibit particular individual

behaviors (Grassé, 1959; Deneubourg and Goss, 1989). In a transport task, for instance, coordination could occur by inter-individual competition and/or via the object to be moved.

In a remarkable series of works, Kube and Zhang (1993a, 1993b) and Kube et al. (1993) studied a decentralized approach to let a group of simple robots push an object that was too heavy to be moved by a single robot. The approach did not make use of explicit communication among the robots. Taking inspiration from ant colonies, they extended the model adding a stagnation recovery mechanism (Kube and Zhang, 1995). Later, Kube and Bonabeau (2000) provided a first formalized model of cooperative transport in ants. They used a 2D simulator and a real robotic system for validation. Kube and Zhang, as well as Sudd, reported a relative lack of efficiency of the behaviors they observed.

The aim of the experiment described in the following of this paper is to study to what extent a group of two simple s-bots equipped with very limited cognitive and computational abilities is able to exhibit coordination of individual activities and to exploit environmental signals in order to transport an object as far as possible within a fixed time period. To do so, we synthesize individual control policies via an evolutionary algorithm with the objective of obtaining highly efficient group behaviors.

2 Experimental Framework

2.1 The Environment and the s-Bot

The experiment is carried out using a 3D physics simulator.¹ The simulated environment consists of a flat ground of infinite size, a light-emitting beacon, and a prey. The prey is modeled as a cylinder of mass 500 g, 10 cm in height, and 12 cm in radius. The friction coefficient between the prey and the ground is set to 0.25. To calculate frictional forces, an approximation based on the Coulomb friction law is used. Gravity is set to 981 cm s^{-2} .

The s-bot model is shown in Figure 1. It is an approximation of a real s-bot, currently under construction within the “Swarm-bots” project (Mondada et al., 2002; Dorigo et al., 2003; Trianni et al., 2003; see also www.swarm-bots.org). The body is composed of a cylindrical torso, a *gripper* element that is fixed on the torso’s edge heading forward, a long cylindrical pillar placed on the top of the torso to support a visual camera system, and four wheels: two active wheels on the left and right, and two passive wheels, one in the front, and one in the back. The s-bot has a height of 16.85 cm and a mass of 660 g.

The control system reads the sensors status and sets actuator activations every 0.1 seconds. In the following, the actuators and sensors of the s-bots are detailed.

¹ The core simulator system has been developed in cooperation with IDSIA (Dalle Molle Institute of Artificial Intelligence Studies), Switzerland. It is based on libraries of the commercial physics engine VortexTM from CMLabs.

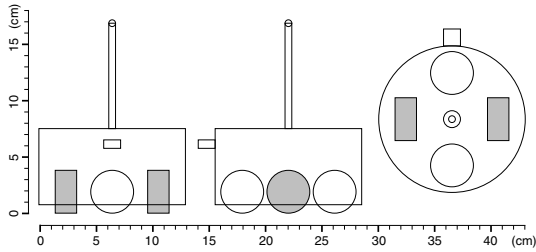


Fig. 1. Front, side and top view of the s-bot: the cylindrical torso is equipped with a small gripper element heading forward, and a camera placed on the pillar. Two passive, spherical wheels (white) are connected to the torso’s back and front part. Two motorized, cylindrical wheels (gray) are connected to the torso on the left and right side.

Actuators. The s-bot is equipped with two active wheels and a gripper element. The active wheels are connected to the torso using a motorized hinge joint. The torque that can act to accelerate an active wheel is limited to $200,000 \text{ dyne cm}^2$. An active wheel can be controlled by setting a desired angular speed $v \in [-15 \text{ rad/s}, 15 \text{ rad/s}]$. To model real actuators’ noise in the simulation, some wheels are biased to turn faster than others, the wheel’s speed is noisy, and activations less than certain random thresholds will not result in any movement of the wheel (for details see Groß, 2003).

The gripper element is a controllable, sticky box heading forward. If the element is in *gripping* mode, a connection will be established as soon as the gripper element is in contact with a grippable object. This is realized by dynamically creating a ball-and-socket joint connecting the s-bot’s torso and the object. The joint is positioned on the gripper element. Once the gripper is set to the *open* mode, this joint will be removed to release the object.

The connection formed by the gripper element will break if a too strong force is transmitted via the corresponding joint. This is a characteristic of the real s-bot and of gripper elements in general. The limit for the force acting on the gripper element is $1,000,000 \text{ dyne}$ (i.e., 10 Newton). This limit still permits s-bots engaged in the transport of a heavy prey to form structures such as small pulling chains.

Sensors. The s-bot is equipped with an omnidirectional camera and a gripper status sensor. These sensors have their counterparts on the real s-bot. A summary of the information provided by each sensor is presented in Table 1.

The camera is mounted on a pillar support that is fixed at the center of the torso’s top. The camera sensor is situated 16.85 cm above the ground. In principle, the camera of the real s-bot is able to sense objects in all horizontal directions having an infinite range. However, the quality of the perceived signal

² $100,000 \text{ dyne}$ correspond to 1 Newton .

decreases when the object’s distance is increased. Therefore, the sensing range is restricted depending on the type of feature that is supposed to be extracted.

The simulated camera provides data based not on recorded images but on information available in the simulator, e.g., the distance to another object. The implicit assumption here is that in case of the real s-bot, such information can be obtained from the camera using feature extraction algorithms.

The camera sensor can provide information about a cylindrical prey in the surrounding, if the horizontal distance between the cylinder’s border and the camera is at most 50 cm. In particular, the camera sensor is able to detect the horizontal distance between the s-bot and the closest point of the prey’s border, and the horizontal angle to the corresponding point with respect to the s-bot’s direction of forward motion.

Moreover, the camera sensor is able to detect the horizontal direction of the highest intensity of light perceived. Light is emitted by a single beacon that is placed in the environment.

Noise affecting the sensors is modeled in simulation in various ways. In case of the light or the prey perceived, an angular deviation r_α is added to the horizontal angle α_{rad} (radian measure) that indicates the target direction to the light or the prey. r_α is generated using the normal distribution $N(0, 0.01)$. Also the measured distance (in cm) to a perceived prey is affected by noise. This is modeled by adding a random variable following the normal distribution $N(0, 1)$.

Table 1. List of information provided by the sensors of an s-bot.

Sensor device	Information provided
gripper status sensor	status of being connected to another object through the gripper element
camera sensor	horizontal angle and distance to a cylindrical prey (horizontal sensing range: 50 cm)
	horizontal angle of highest intensity of light perceived (a constant light signal is emitted by a beacon)

2.2 The Neural Network Controller

The group of s-bots is controlled by a simple recurrent neural network that is synthesized by an evolutionary algorithm. All the s-bots of a group transporting a prey are initially equipped with an identical neural network. Due to recurrent connections, these networks have memory and they are not restricted to purely reactive behaviors. The ability to compare current and previous states of the environment perceived might be beneficial to recognize the presence and the actions of teammates.

We used an Elman network (Elman, 1990) with five input neurons, five (fully inter-connected) hidden neurons and three output neurons. The activations of

the neurons of the input layer correspond to the current s-bot's sensor reading values. The activations of the output neurons are used to control the motorized wheels and the gripper element.

2.3 The Evolutionary Algorithm

The evolutionary algorithm used is a self-adaptive version of a $(\mu + \lambda)$ evolutionary strategy (Rechenberg, 1965; Schwefel, 1975). An individual is composed of real-valued object parameters $X = (x_1, x_2, \dots, x_N)$ specifying the weights of the Elman Network used to control the s-bots, and real-valued strategy parameters $S = (s_1, s_2, \dots, s_N)$ specifying the mutation strength used for each component of the object parameter vector.

Before the evolution is started, a *random walk* is performed in order to generate a wide variety of somewhat acceptable solutions to start from. The total number of fitness evaluations during the random walk is equal to the number of evaluations of an evolution lasting ten generations.

In each generation all individuals are assigned a fitness value. The best μ individuals are selected to create λ offspring. 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 λ offspring and the μ parent individuals are copied into the population of the next generation. Note that the parent individuals that are replicated from the previous generation get re-evaluated. In fact, the fitness values are affected by various random components in the fitness estimation procedure. Re-evaluating the parents' fitness values inhibits a potential systematic over-estimation in time due to previous fitness estimations.

The object parameter x_i is mutated by adding a random variable from the normal distribution $N(0, s_i^2)$. Beforehand, the mutation strength parameter s_i is multiplied by a random variable ξ_i that follows a *lognormal* distribution similar to the one proposed by Schwefel (1974). As recombination operator we use intermediate and dominant recombination (Beyer, 2001), both with the same probability.

The number of offspring is $\lambda = 80$ and the number of parents is $\mu = 20$. For a detailed description of the evolutionary strategy used see Groß (2003).

2.4 The Fitness

The task is to control a group of two simple s-bots so that they cooperate to transport as far as possible a heavy prey within a fixed time period. The direction of movement of the prey is not predetermined. Since the prey cannot be moved by a single s-bot, coordination among s-bots is necessary.

Each individual, that is, a common controller for a group of s-bots, is evaluated by performing S tests t_1, t_2, \dots, t_S against a sample composed of S configurations specifying the s-bots' initial placements and the position of the beacon. The sampling size is set in our experiments to $S = 5$. The sample is changed

once per generation, so that all the individuals that compete with each other are evaluated under similar conditions. This should increase the comparability of fitness values among individuals within the same generation. Note that the μ parent individuals that are copied into the next generation by default get re-evaluated based on the new sample.

In each test, the simulation lasts 20 simulated seconds. The prey is placed in the center of the environment. The s-bots are placed at random positions and orientations, but not more than 25 cm away from the prey. This ensures that the prey can initially be detected by each s-bot. The beacon is placed at a random position 300 cm away from the center of the environment. This is less than the distance the prey can be moved within the simulation time of 20 seconds.

For each test $t_i, i \in \{1, \dots, S\}$, a quality measure q_i (i.e., the transport quality) evaluates the exhibited transport performance. The final fitness score f is computed using a weighted average. Let ϕ be a permutation of $\{1, 2, \dots, S\}$ such that $q_{\phi(1)} \leq q_{\phi(2)} \leq \dots \leq q_{\phi(S)}$, then the fitness value is given by

$$f = \frac{2}{S(S+1)} \sum_{i=1}^S (S-i+1) q_{\phi(i)}. \quad (1)$$

In this way, tests resulting in weaker transport quality q_i contribute more than the others to the fitness value so that fluctuations get punished. The intended effect is that individuals that emerge during evolution get more resistant over time to the randomness of the robots' initial placements as well as to the noise in the robots' sensors and actuators.

To favor the evolution of solutions that let all s-bots participate in the transport activities, the transport quality takes into account also the clustering performance of s-bots around the prey. The clustering quality c_i in a test t_i is defined as

$$c_i = \frac{1}{n} \sum_{j=1}^n c_i^{(j)}, \text{ and } c_i^{(j)} = \begin{cases} 0 & \text{if } d_i^{(j)} > 50; \\ 1 & \text{if } d_i^{(j)} < 25; \\ \frac{50-d_i^{(j)}}{25} & \text{otherwise;} \end{cases} \quad (2)$$

where $d_i^{(j)}$ denotes the Euclidean distance (in cm) observed in test t_i between the final positions of the j th s-bot and the center of the prey minus the radius of the prey.

Therefore, s-bots that are not more than 25 cm away from the perimeter of the prey get the maximum clustering performance value of 1. In this way, any structure of two collaborating s-bots pushing or pulling the prey cannot be punished. S-bots that are more than 50 cm away cannot sense the prey any more and get the lowest possible clustering performance value of 0.

The transport quality q_i in a test t_i is defined as

$$q_i = \begin{cases} c_i & \text{if } D_i = 0; \\ 1 + (R + \sqrt{D_i})c_i^\rho & \text{otherwise;} \end{cases} \quad (3)$$

where D_i denotes the Euclidean distance (in cm) observed in test t_i between the prey's initial and final positions, $R = 1$ is a constant reward and $\rho = 5$ is a constant exponent in order to punish solutions exhibiting weak clustering performance. The root function is applied as scaling function.

3 Results

The experimental setup described above has been used in ten independent evolutionary runs of 150 generations each. One run lasted a little bit less than a week on a machine equipped with an AMD Athlon XP™ 1800+ processor. In the following, the transport performance of the evolved controllers and their ability to scale when using larger group sizes are discussed.

Quality of transport. The transport quality of a group (as defined by Equation 3) in a test depends on the genotype specifying the neural network controller, the s-bots' initial positions and orientations, the position of the beacon in the environment, the particular offset and threshold values of each sensor and actuator, and the amount of noise that affects sensors and actuators at each time step. Of course, the ultimate goal is to generate genotypes that perform well under every possible condition. However, there is a very large number of possible conditions, and during evolution the solutions are evaluated using a sample of five different conditions per generation.

Figure 2 displays the development of the best and average fitnesses for certain sets of individuals over all ten evolutionary runs:

- The dark gray boxes correspond to the observed fitness values of the best rated individuals for every population. Due to the random component in the fitness evaluation procedure, individuals exhibit performance fluctuations. Thus, the best fitness values are likely to be over-estimated.
- The average fitness value of the μ parent individuals of each population is illustrated by the light gray boxes. To avoid bias in the evaluation of the parent individuals, they get re-evaluated in the subsequent generation. In this way, although the average fitness of the re-evaluated parent individuals may fluctuate due to the noise, there is no systematic over-estimation caused by previous fitness estimations. In the figure, the average fitness value of the parent individuals is computed based on the set of re-evaluated fitness values.
- The white boxes correspond to the average fitness values of the entire population.

Looking at Figure 2 one can see that, in all cases, the fitness values tend to increase. Since the fitness values are noisy and, additionally, they are computed using a weighted average, it is hard to estimate the attained quality level. Therefore, we measured the transport quality of several individuals of each evolutionary run on a sample of 500 different tests. For each evolutionary run, we selected the $\mu = 20$ best (parent) individuals of the final generation. This set

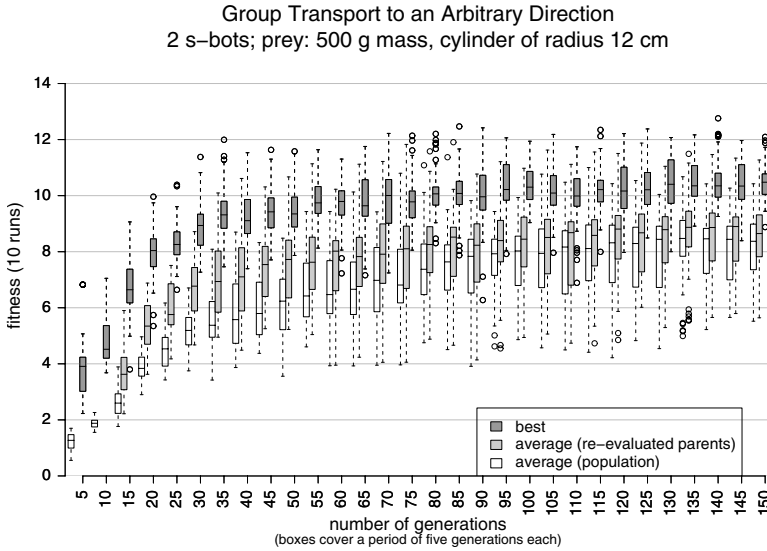


Fig. 2. Box-and-whisker plot visualizing certain characteristics of the evolutionary progress in ten runs. The box comprises observations ranging from the first to the third quartile. The median is indicated by a bar, dividing the box into the upper and lower part. The whiskers extend to the most extreme data point which is no more than 1.5 times the interquartile range. Outliers are indicated as circles. This plot illustrates the development of the best fitness of a population (dark-gray boxes), the average fitness of the set of $\mu = 20$ (re-evaluated) parent individuals (light-gray boxes) and the average fitness of the whole population (white boxes). Each box covers a period of five generations.

of μ parent individuals comprises all genetic material that would be exploited in subsequent generations in case the evolution would be continued. Performing the post-evaluation, it has been observed that parent individuals of the final generation of the same evolutionary run exhibit a quite similar performance. Therefore, we focus on two types of individuals from each evolutionary run: the one with the highest average performance concerning the 500 tests (type 1), and the one with the lowest observed standard deviation (type 2). For every evolutionary run, both types of individuals are post-evaluated for a second time, on a random, but fixed set of 1,000 tests.

Figure 3 shows a box-and-whisker plot presenting the transport quality values observed in these 1,000 tests for both types of individuals for all the evolutionary runs. The first five solutions displayed make an essential use of the gripper element.

According to Equation 3, the transport quality q_i observed in a test t_i is at least 1 if the prey has been moved, and $c_i \in [0, 1]$ otherwise. Looking at Figure 3, it can be seen that the prey has been moved in almost all cases, since almost all observed transport quality values are in the range 2 to 15. The distance D_i

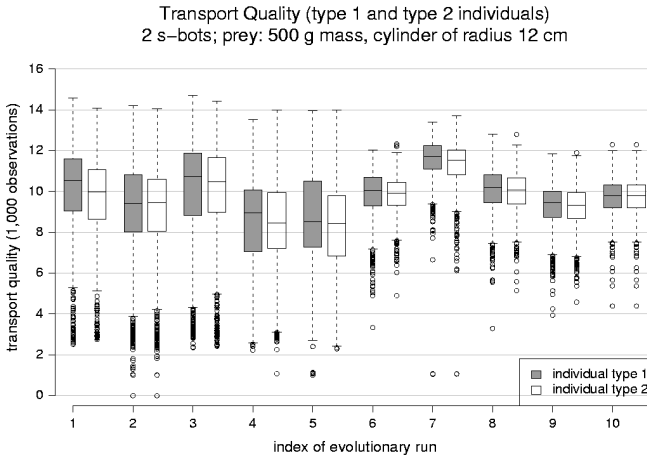


Fig. 3. Box-and-whisker plot visualizing the transport quality observed in 1,000 tests for individuals of type 1 and type 2 in the final generations. The type 1 individual of an evolutionary run is the one with the highest average of the transport quality values observed in the 500 tests during the first post-evaluation and the type 2 individual is the one with the lowest observed standard deviation.

the prey has been moved is at least $(q_i - 2)^2$, if $q_i \geq 2$.³ If $c_i < 1$, D_i is bigger than $(q_i - 2)^2$. Assuming the worst case for the distance moved, that is, a perfect clustering quality $c_i = 1$, some examples for pairs of transport quality q_i and the corresponding moved distance D_i are given in the following table:

q_i	2	3	4	5	6	7	8	9	10	11	12	13	14
D_i (in cm)	0	1	4	9	16	25	36	49	64	81	100	121	144

Let us consider a setup in which a group of two s-bots, structured in a connected, linearly aligned chain, is pulling the prey for a period of 20 seconds. The first s-bot of the chain is connected to the prey right from the start. All s-bots are controlled by a handwritten controller: each s-bot moves backwards with maximum angular speed of the wheels; the gripper element remains closed. If a chain of two s-bots is placed in such a structural configuration, the distance the object can be moved within the simulation time of 20 seconds is around 149.64 cm. This corresponds to a transport quality of approximately 14.23.

However, during the fitness evaluation the s-bots are not placed in such an initial structure: the s-bots are scattered in the environment at random positions and with random orientation. Therefore, the s-bots have to approach the prey and to coordinate themselves before a structure is formed that can be used to

³ Let $q_i \geq 2$: $q_i = 1 + (1 + \sqrt{D_i})c_i^5 \leq 1 + (1 + \sqrt{D_i}) = 2 + \sqrt{D_i} \Rightarrow q_i - 2 \leq \sqrt{D_i} \Rightarrow D_i \geq (q_i - 2)^2$.

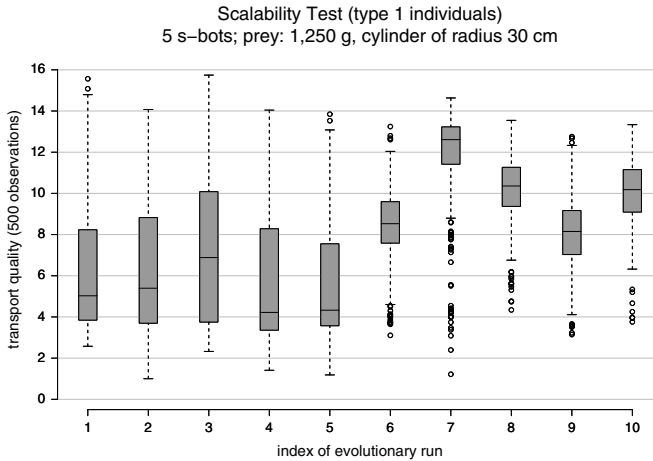


Fig. 4. For each evolutionary run, the individual for which the best average performance was observed in the post-evaluation described previously (the type 1 individual) is evaluated 500 times using a group of five s-bots and a prey of radius 30 cm and mass 1,250 g. The simulation period is extended by ten additional seconds, since it requires more time for each s-bot to move around the larger prey. The first five individuals make essential use of the gripper element.

transport the prey. If we assume that the randomly placed s-bots need ten seconds to form a chain like the one described, this structure could only pull the prey for the remaining ten seconds of simulation time. During this period, the prey can be moved 74.50 cm, corresponding to a transport quality of approximately 10.63.

The median values of transport quality exhibited by the better performing half of the type 1 individuals are in the range [10.03,11.73]. The observed standard deviations concerning the half of the type 2 individuals having lower fluctuations are in the range [0.91,1.30]. In case of the evolutionary runs 6 to 10, we have obtained a number of satisfactory controllers that exhibit low performance fluctuations (see Figure 3). Overall, the controlled groups of s-bots act quite robustly with respect to various kinds of noise concerning the sensors and the actuators and with respect to different initial placements in the environment.

Scalability. The question arises whether these observed behaviors are scalable, that is, whether evolved individuals are able to cooperate in the transport of a heavier prey when the group size becomes larger. To try to give an answer, for each evolutionary run we took the individual with the best average performance and evaluated it using a group of five s-bots, engaged in moving a heavier prey of mass 1,250 g and of the same shape and size as previously. In most cases, the observed performance was low, since, having five s-bots, the strategies seemed not to be able to handle the increased inter-individual competition among s-bots

for the limited space around the prey. However, if the perimeter of the prey is multiplied by the same factor as the number of s-bots is increased, all solutions are able to move the prey, and the ones not relying on the gripper exhibit a satisfactory performance (see Figure 4).

4 Conclusions

In this work, we addressed the problem of controlling a group of two simple, mobile, autonomous robots so that they cooperatively transport a heavy prey as far as possible into an arbitrary direction within a fixed period of time. The robots are able to establish physical connections with each other as well as with the prey. The prey cannot be moved without cooperative behavior. There is no explicit communication among the s-bots. Moreover, the s-bots are not able to directly sense each other. Communication is limited to *interactions via the environment*, that is, stigmergy (Grassé, 1959; Dorigo et al., 2000).

In all of the ten evolutionary runs carried out, controllers were generated which exhibited an acceptable performance under most of the tested 1,000 conditions. In five out of ten runs, the system generated very high quality controllers which reached high performance levels and, at the same time, presented rather low fluctuations in performance. Moreover, many solutions could be applied to larger groups so that prey of bigger size and mass could be moved.

In general, the controllers' performances are very sensitive to the size of the prey. We also observed that the performance of those solutions that make essential use of the gripper element is more fluctuating than that of the other ones. The reason for this behavior is however not clear yet. In the future, we want to encourage the evolution of strategies that let the robots organize into assembled structures in order to transport prey of various shapes and sizes.

Acknowledgments. The authors wish to thank all the members of the “Swarm-bots” project for their support, suggestions and comments. This work was supported by the “Swarm-bots” project, funded by the Future and Emerging Technologies programme (IST-FET) of the European Commission, under grant IST-2000-31010, and by the “Virtual Swarm-bots” project, funded by the Belgian FNRS, contract no. 9.4515.03. Marco Dorigo acknowledges support from the Belgian FNRS, of which he is a Senior Research Associate. 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.

References

- Beyer, H.-G.: The Theory of Evolution Strategies. Springer-Verlag, Berlin, Germany (2001)
- Deneubourg, J.-L., Goss, S.: Collective patterns and decision-making. *Ethology, Ecology and Evolution* **1** (1989) 295–311

- Deneubourg, J.-L., Goss, S., Sandini, G., Ferrari, F., Dario, P.: Self-organizing collection and transport of objects in unpredictable environments. In: Proc. of Japan–USA Symposium on Flexible Automation, Kyoto, Japan, ISCI (1990) 1093–1098
- Dorigo, M., Bonabeau, E., Theraulaz, G.: Ant algorithms and stigmergy. *Future Generation Computer Systems* **16** (2000) 851–871
- Dorigo, M., Trianni, V., Şahin, E., Labella, T.H., Groß, M.R., Baldassarre, G., Nolfi, S., Deneubourg, J.-L., Mondada, F., Floreano, D., Gambardella, L.M.: Evolving self-organizing behaviors for a *Swarm-bot*. Technical Report IRIDIA/2003-11, Université Libre de Bruxelles, Belgium (2003)
- Elman, J.: Finding structure in time. *Cognitive Science* **14** (1990) 179–211
- Grassé, P.: 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. *Insect Sociaux* **6** (1959) 41–81
- Groß, R.: Swarm-intelligent robotics in prey retrieval tasks. Technical Report IRIDIA/2003-27, Université Libre de Bruxelles, Belgium (2003) DEA thesis.
- Kube, C.R., Bonabeau, E.: Cooperative transport by ants and robots. *Robotics and Autonomous Systems* **30** (2000) 85–101
- Kube, C.R., Zhang, H.: Collective robotic intelligence. In: Second International Conference on Simulation of Adaptive Behaviour (SAB-1992), Cambridge, MA, MIT Press (1993a) 460–468
- Kube, C.R., Zhang, H.: Collective robotics: from social insects to robots. *Adaptive Behaviour* **2** (1993b) 189–218
- Kube, C.R., Zhang, H.: Stagnation recovery behaviours for collective robotics. In: Proceedings 1994 IEEE/RSJ/GI International Conference on Intelligent Robotics and Systems, Los Alamitos, CA, IEEE Computer Society Press (1995) 1883–1890
- Kube, C.R., Zhang, H., Wang, X.: Controlling collective tasks with an ALN. In: 1993 IEEE/RSJ International Conference on Intelligent Robots and Systems, Los Alamitos, CA, IEEE Computer Society Press (1993) 289–293
- Mondada, F., Pettinaro, G.C., Kwee, I., Guignard, A., Gambardella, L., Floreano, D., Nolfi, S., Deneubourg, J.-L., Dorigo, M.: SWARM-BOT: A swarm of autonomous mobile robots with self-assembling capabilities. In: Proceedings of the International Workshop on Self-organisation and Evolution of Social Behaviour, Monte Verità, Ascona, Switzerland, University of Zurich (2002) 307–312
- Parker, L.E.: Current state of the art in distributed autonomous mobile robotics. In: Distributed Autonomous Robotic System. Volume 4, Tokyo, Japan, Springer-Verlag (2000) 3–12
- Rechenberg, I.: Cybernetic Solution Path of an Experimental Problem. PhD thesis, Royal Aircraft Establishment, Hants Farnborough, UK (1965)
- Schwefel, H.-P.: Adaptive Mechanismen in der biologischen Evolution und ihr Einfluß auf die Evolutionsgeschwindigkeit. Technical Report Re 215/3, Technische Universität Berlin, Germany (1974)
- Schwefel, H.-P.: Evolutionsstrategie und numerische Optimierung. PhD thesis, Technische Universität Berlin, Germany (1975)
- Sudd, J.: The transport of prey by an ant *Pheidole crassinoda*. *Behaviour* **16** (1960) 295–308
- Trianni, V., Groß, R., Labella, T.H., Şahin, E., Dorigo, M.: Evolving aggregation behaviors in a swarm of robots. In: Advances in Artificial Life – Proceedings of the 7th European Conference on Artificial Life (ECAL). Volume 2801 of Lecture Notes in Artificial Intelligence, Springer-Verlag, Berlin, Germany (2003) 865–874