



**Université Libre de Bruxelles**

*Institut de Recherches Interdisciplinaires  
et de Développements en Intelligence Artificielle*

**Evolution of Solitary and Group Transport  
Behaviors for Autonomous Robots  
Capable of Self-Assembling**

Roderich Groß and Marco Dorigo

**IRIDIA – Technical Report Series**

Technical Report No.  
TR/IRIDIA/2006-022

August 2006

**IRIDIA – Technical Report Series**  
ISSN 1781-3794

Published by:

IRIDIA, *Institut de Recherches Interdisciplinaires  
et de Développements en Intelligence Artificielle*  
UNIVERSITÉ LIBRE DE BRUXELLES  
Av F. D. Roosevelt 50, CP 194/6  
1050 Bruxelles, Belgium

Technical report number TR/IRIDIA/2006-022

Revision history:

TR/IRIDIA/2006-022.001	August 2006
TR/IRIDIA/2006-022.002	May 2008

The information provided is the sole responsibility of the authors and does not necessarily reflect the opinion of the members of IRIDIA. The authors take full responsibility for any copyright breaches that may result from publication of this paper in the IRIDIA – Technical Report Series. IRIDIA is not responsible for any use that might be made of data appearing in this publication.

# Evolution of Solitary and Group Transport Behaviors for Autonomous Robots Capable of Self-Assembling

Roderich Groß<sup>1,2</sup>

roderich.gross@ieee.org

Marco Dorigo<sup>2</sup>

mdorigo@ulb.ac.be

<sup>1</sup> Ant Lab, School of Biological Sciences, University of Bristol, UK

<sup>2</sup> IRIDIA, Université Libre de Bruxelles, Brussels, Belgium

May 2008

## Abstract

Group transport is being performed in many natural systems and has become a canonical task for studying cooperation in robotics. We simulate a system of simple, insect-like robots that can move autonomously and grasp objects as well as each other. We use artificial evolution to produce solitary transport and group transport behaviors. We show that robots, even though not aware of each other, can be effective in group transport. Group transport can even be performed by robots that behave as in solitary transport. Still, robots engaged in group transport can benefit from behaving differently from robots engaged in solitary transport. The best group transport behaviors yielded by half of the evolutions let the robots organize into self-assembled structures. This provides evidence that self-assembly can provide adaptive value to individuals that compete in an artificial evolution based on task performance. We conclude the paper by discussing potential implications for evolutionary biology and robotics.

**Keywords** group transport · solitary · social behavior · evolution of cooperation · self-assembly · autonomous robots · evolutionary robotics · swarm robotics · swarm intelligence · evolutionary biology

## 1 Introduction

Group transport can be defined as the “conveyance of a burden by two or more individuals” [55, page 227]. Group transport is being performed in many natural systems [22, 55, 81, 41]. It has also become a canonical task for studying cooperation in artificial systems [45, 43, 61, 29, 73]. When compared to solitary transport, it offers the advantage of being more reliable and in addition more powerful, as a group may exert higher forces onto an object than each of the participating individuals alone.

We use evolutionary algorithms to produce solitary transport and group transport behaviors for a system of simulated robots [58, 38]. Similar to many animals, our robots can move based on their own propulsion and have a mechanism by which they can grasp an object to be transported or another robot. Their cognitive capabilities are very limited and similar to those of solitary animals: they can neither perceive teammates nor communicate with them directly. The task is to move the object in an arbitrary direction.

The aim of the study presented in this paper is two-fold. Firstly, we want to understand to what extent individuals that as a group have to accomplish a cooperative task can benefit from behaving differently from individuals that have to accomplish tasks on their own. In particular, we examine whether individuals engaged in group transport can benefit from behaving differently from individuals engaged in solitary transport. Answers to this research question are relevant for the design of robotic systems. A deeper understanding of the relation

between solitary and social behavior can also shed light on the evolution of cooperation in animal groups and societies, which is one of the most important unanswered questions in evolutionary biology.

Secondly, we examine whether evolutionary algorithms can yield behaviors that let individuals organize into physically connected pushing or pulling structures to accomplish a task. This *self-assembly* ability, without being explicitly favored by the fitness function of the evolutionary algorithm, can evolve if providing an adaptive value for the individuals. This can lead to new insight for the design of robotic systems. Moreover, gaining some basic understanding of the factors that favor self-assembly formation is one focus of current research in biology [77, 4].

Group transport and self-assembly are both widely observed phenomena in ant colonies [22, 55, 4]. In both cases, ants temporarily organize into functional units at a level intermediate to that of the individual and that of the colony [3]. In designing group transport systems, we follow some basic principles that are used to explain the behavior of such functional units, the behavior of whole colonies, and collective animal behavior in general [30, 76, 10, 75, 28]. We emphasize system properties such as: (i) *decentralization*, that is, robots follow rules in a fully autonomous and distributed manner; (ii) *locality of perception / indirect communication*, that is, robots perceive objects in a limited range only, and communicate through physical interactions with their environment; and (iii) *redundancy*, that is, the robotic system can continue to function even when faced with a moderate reduction in its workforce.

Evolutionary algorithms (and other population-based meta-heuristic optimization algorithms) have already proven successful in reproducing, with (simulated or real) autonomous robots, several collective capabilities known from social insects and other animals, including aggregation [17, 64], flocking or schooling [66, 83, 65, 5, 71], foraging [60, 62], and inspection/patrolling [50, 91]. Recently, Pérez-Urbe [62] and Floreano *et al.* [20] studied evolutionary conditions for the “emergence” of cooperative behavior in groups of robots that perform a foraging task. They found that cooperation (in the form of information transfer using pre-existing communication devices) evolved best when groups consisted of genetically identical individuals and when selection acted at the level of groups.

Different from other work in the literature, we study the evolution of cooperation by investigating the relation between solitary behavior and social behavior. We do so by focusing on relatively primitive forms of cooperation that are likely to be available when solitary individuals encounter each other. We use a physics-based 3-D simulator, in which robots can influence each other by means of physical interaction, both directly and indirectly, that is, through the object being manipulated. Coordination can also be implicit, for instance, when the robots’ behaviors exploit invariants or cues present in the environment.

In the evolutionary algorithm used, a “genotype” encodes a simple recurrent neural network which is cloned and copied to each robot within a group. Thus, all members of a group are genetically identical. However, this does not preclude *variability* in their behaviors, as each of the robot in a group can be different, in terms of its phenotype<sup>1</sup> and of the experience gained during its life-time. We consider a population of genotypes, that is, a population of groups. Selection acts at the level of genes [37, 13].

We consider two types of environments. The first one is used to evolve solitary transport behaviors. Only one robot and one light object are present, and consequently behavior is not selected for being social. The second type of environment is used to evolve group transport behaviors. In this setup, two robots and a heavy object are present, and the object can not be moved without cooperation. Any behavior must be social, either mutually beneficial or spiteful [84]. Note that, in this setup, selection at the level of genes is equivalent to between-group selection [8, 42, 21].

The paper is organized as follows. Section 2 details the methods. Section 3 describes the results of our study on the evolution of solitary and group transport behavior. Section 4 overviews the related work. Section 5 discusses the results and concludes the paper.

---

<sup>1</sup>Randomness affects properties of the robot (e.g., its perceptual range). This may account for imprecision associated with building robotic hardware, and is not genetically determined.

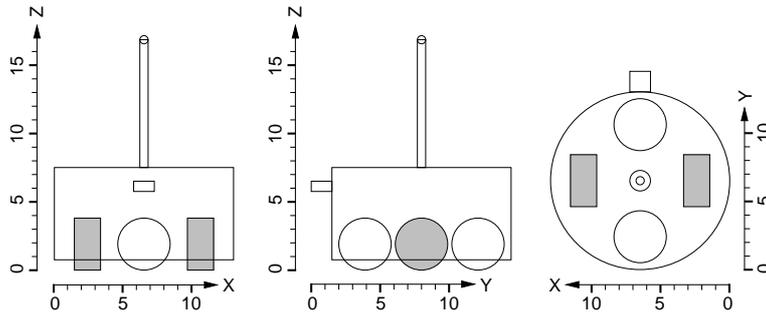


Figure 1: The simulation model of the robot: front, side, and top view (units in cm).

## 2 Methods

In the following, we detail the task, the simulation model, the robot’s controller, and the evolutionary algorithm.

### 2.1 Task

We study solitary and group transport of an object, hereafter also called the *prey*. The robots’ environment consists of a flat ground, the prey, and a light source. The prey is modeled as a cylinder, 10 cm in height, and 12 cm in radius. For solitary transport (i.e., 1 robot), we use a prey of mass 250 g, for group transport (i.e., 2 robots) we use a prey of mass 500 g. The 500 g prey can not be moved by a single robot. The light source represents an environmental cue and as such can be exploited by the robots to coordinate their actions. Initially, the robots are put at random positions near the prey. The task is to move the prey in an arbitrary direction (the farther the better).

### 2.2 Simulation Model

The simulator models the kinematics and dynamics of rigid, partially constrained, bodies in 3-D using the Vortex<sup>TM</sup> simulation toolkit (CMLabs Simulations, Inc., Canada). Frictional forces are calculated based on the Coulomb friction law [12]. The model of the robot is illustrated in Figure 1. It is an approximation of a physical robot, called *s-bot*, that was designed and implemented in the context of the Swarm-bots project [16, 56, 18]. The model is composed of five bodies: two spherical wheels, two cylindrical wheels, and a cylindrical torso. The torso is composed of several parts that are rigidly linked: a cylindrical body, a protruding cuboid (in what we define to be the robot’s front), and a pillar fixed on top. The spherical wheels are linked to the chassis via ball-and-socket joints. The cylindrical wheels are linked to the chassis via hinge joints. Each wheel weighs 20 g. The robot has a total mass of 660 g.

The robot’s actuators and sensors are summarized in Table 1. The cylindrical wheels are motorized, and can be moved with angular speeds  $w_l$  and  $w_r$  (in rad/s) in the range  $[-M, M]$  ( $M \approx 15$ ). The cuboid heading forward represents a connection device (e.g., a gripper). If it is in contact with either the cylindrical body of another robot or the (cylindrical) prey, a physical connection can be established ( $c = 1$ , and  $c = 0$  otherwise). Connections can be released at any time. In particular, this will happen if the intensity of the force transmitted by the connection mechanism exceeds a certain threshold. As a consequence, it is not possible for the robots to form very long pulling chains. The robot is equipped with an omni-directional camera mounted on a pillar support that is fixed at the center of the torso’s top. The camera is able to detect the angular position  $\alpha$  of the light source. Moreover, it provides the angular position  $\beta$  and distance  $d$  of the prey, if the latter resides within the sensing range ( $R \approx 50$  cm). In simulation, angles and distances can be calculated from the positions of other objects in the scene.<sup>2</sup> A connection sensor enables a robot to perceive whether it is connected to another

<sup>2</sup>On the physical s-bot robot, instead, we obtain information on angles and distances by processing the image taken from the camera by feature extraction algorithms [31, 59].

Table 1: Summary of the robot’s actuators and sensors. Units are in cm, rad, and rad/s. For details see text.

ACTUATORS	
left wheel (angular speed)	$w_l \in [-M, M]$
right wheel (angular speed)	$w_r \in [-M, M]$
connection mechanism	$c \in \{0, 1\}$
SENSORS (EXTEROCEPTIVE)	
light source (angular position)	$\alpha \in [0, 2\pi]$
prey (angular position)	$\beta \in [0, 2\pi]$
prey (distance)	$d \in [0, R]$
SENSORS (PROPRIOCEPTIVE)	
connection mechanism	$c' \in \{0, 1\}$

object ( $c' = 1$ ) or not ( $c' = 0$ ). The robot is not equipped with any sensor capable of detecting the presence of a teammate.

Random noise affects the characteristics of the robot’s actuators and sensors (i.e., variables  $w_l$ ,  $w_r$ ,  $M$ ,  $\alpha$ ,  $\beta$ ,  $d$ , and  $R$ ). We model two different types of random noise: (i) random variables  $\xi_1, \xi_2, \xi_3, \dots$  generated for each robot only once at the beginning of its life-time; they model differences among the hardware of the robots; (ii) random variables  $\xi_1^t, \xi_2^t, \xi_3^t, \dots$  generated for each robot at each time step  $t$  during its life-time; they model temporary fluctuations in the behavior of the robot’s actuators and sensors. Let  $N(\mu, \sigma^2)$  denote the normal distribution with mean  $\mu$  and variance  $\sigma^2$ . The value  $w_l$  (in rad/s) is initially within the range  $[-15, 15]$  and modified by multiplication with  $\xi_1 \xi_1^t$ , where  $\xi_1$  follows  $N(1, 0.02^2)$  and  $\xi_1^t$  follows  $N(1, 0.05^2)$ . If  $w_l$  is less than  $\xi_2$ , which follows the uniform distribution  $U(0.1, 0.5)$ , then the speed is set to zero. The value  $w_r$  is modified in a similar way using independent random variables  $\xi_3, \xi_3^t$ , and  $\xi_4$ . The camera is assumed to be calibrated with no bias in the error. The only exception to this is the sensing range  $R$  (in cm), which is set to  $\xi_5$ , which follows  $N(50, 1)$ . Angle  $\alpha$  (in rad) is modified by adding  $\xi_6^t$ , which follows  $N(0, 0.1^2)$ , and angle  $\beta$  is modified by adding  $\xi_7^t$ , which follows  $N(0, 0.1^2)$ . Distance  $d$  (in cm) is modified by adding  $\xi_8^t$ , which follows  $N(0, 1)$ .

### 2.3 Controller

All the robots of a group are initially assigned an identical controller. Every 100 ms a control loop spreads activation in a neural network taking input from the robot’s sensors, and uses the outputs as motor commands. The neural network is illustrated in Figure 2. It is a simple recurrent neural network [19] and has an input layer of five neurons ( $i_1, i_2, i_3, i_4$ , and  $i_5$ ), a hidden layer of five (fully inter-connected) neurons, and an output layer of three neurons ( $o_1, o_2$ , and  $o_3$ ). The weights of the synaptic connections of the network are genetically encoded parameters. The activations of the hidden and output neurons are mapped into the range  $(0, 1)$  using the sigmoid function  $f(x) = \frac{1}{1 + \exp(-x)}$ .

The activations of the five input neurons are computed based on the robot’s sensor readings (see Table 1):

$$i_1 = \begin{cases} (1 - \frac{d}{R}) \sin \beta & \text{if } d < R; \\ 0 & \text{otherwise,} \end{cases} \quad (1)$$

$$i_2 = \begin{cases} (1 - \frac{d}{R}) \cos \beta & \text{if } d < R; \\ 0 & \text{otherwise,} \end{cases} \quad (2)$$

$$i_3 = \sin \alpha, \quad (3)$$

$$i_4 = \cos \alpha, \quad (4)$$

$$i_5 = c'. \quad (5)$$

The activations of the three output neurons are used to set the motor commands (see

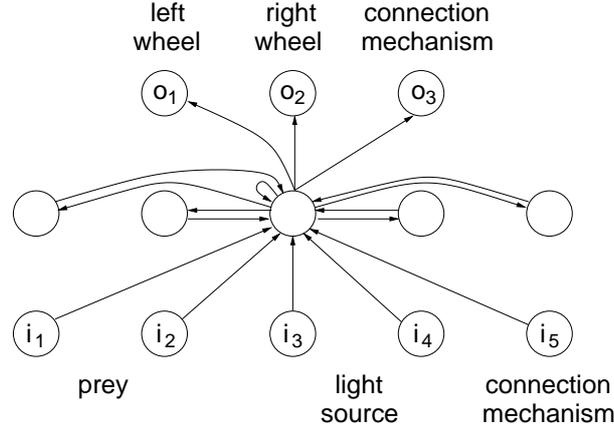


Figure 2: The neural network controller comprising five input neurons (bottom), five hidden neurons (center), and three output neurons (top). Only the synaptic connections to and from the neuron in the center of the hidden layer are illustrated. An additional bias neuron (not shown), providing a constant input of 1, is connected to each neuron of the hidden layer and the output layer.

Table 1):

$$w_l = M(2o_1 - 1), \quad (6)$$

$$w_r = M(2o_2 - 1), \quad (7)$$

$$c = \begin{cases} 0 & \text{if } o_3 < 0.5; \\ 1 & \text{otherwise.} \end{cases} \quad (8)$$

## 2.4 Evolutionary Algorithm

The used evolutionary algorithm is a self-adaptive version of a  $(\mu + \lambda)$  evolution strategy [69, 7]. Each individual<sup>3</sup> is composed of  $n = 73$  real-valued object parameters  $x_1, x_2, \dots, x_n$  specifying the connection weights of the neural network controller, and the same number of real-valued strategy parameters  $s_1, s_2, \dots, s_n$  specifying the mutation strength used for each of the  $n$  object parameters.

The initial population of  $\mu + \lambda$  individuals is constructed randomly. In each generation all individuals are assigned a fitness value. The best-rated  $\mu$  individuals are selected to create  $\lambda$  offspring. Subsequently, the  $\mu$  parent individuals and the  $\lambda$  offspring are copied into the population of the next generation. Note that the  $\mu$  parent individuals that are replicated from the previous generation get re-evaluated. We have chosen  $\mu = 20$  and  $\lambda = 80$ .

Each offspring is created by recombination with probability 0.2, and by mutation, otherwise. In either case, the parent individual(s) is selected randomly. As recombination operators we use intermediate and dominant recombination [7], both with the same probability. The offspring is subjected to mutation. The mutation operator changes the object parameter  $x_i$  by adding a random variable  $\xi_a$  which follows the normal distribution  $N(0, s_i^2)$ :

$$x_i = x_i + \xi_a. \quad (9)$$

Prior to the mutation of object parameter  $x_i$ , the ‘mutation strength’ parameter  $s_i$  is multiplied by a random variable that follows a log-normal distribution [69, 7]:

$$s_i = s_i \exp(\xi_g + \xi_s), \quad (10)$$

where  $\xi_g$ , which is generated once for all strategy parameters, follows the normal distribution  $N(0, \tau_g^2)$ , and  $\xi_s$ , which is generated for each of the strategy parameters  $s_1, s_2, s_3, \dots$ , follows

<sup>3</sup>For simplicity, by *individual* we refer to the genotype. Note that the genotype encodes a neural network controller which is cloned and copied to each robot of a group.

the normal distribution  $N(0, \tau_s^2)$ .  $\tau_g$  and  $\tau_s$  are commonly set to  $\frac{1}{\sqrt{2n}}$  and  $\frac{1}{\sqrt{2\sqrt{n}}}$ , respectively (see [88] and references therein). To prevent premature convergence, a lower bound of 0.01 for the strategy parameters is applied [47].

### 2.4.1 Fitness Computation

The fitness of individuals is assessed using simulations. Each trial lasts  $T = 20$  simulated seconds. Initially, the prey is placed in what we refer to as the center of the environment. The light source is placed at a random position 300 cm away from the prey. This is less than the distance the prey can be moved within the simulation time  $T$ .  $N \in \{1, 2\}$  robots are placed at random positions and orientations, but not more than  $\frac{R}{2} = 25$  cm away from the perimeter of the prey. This ensures that the prey can initially be detected by each robot.

The measure of quality  $\mathcal{Q}$  accounts for the ability of the individual to let the robots remain in the vicinity of the prey, and transport it, the farther the better, in an arbitrary direction. It is defined as:

$$\mathcal{Q} = \begin{cases} \mathcal{C} & \text{if } \mathcal{T} = 0; \\ 1 + (1 + \mathcal{T}^{\rho_1})\mathcal{C}^{\rho_2} & \text{otherwise,} \end{cases} \quad (11)$$

where  $\mathcal{C} \in [0,1]$  reflects the clustering performance,  $\mathcal{T} \in [0, \infty)$  reflects the transport performance, and  $\rho_1 = 0.5$  as well as  $\rho_2 = 5$  are parameters that were determined by trial and error.

The clustering performance  $\mathcal{C}$  is defined as

$$\mathcal{C} = \frac{1}{N} \sum_{i=1}^N \mathcal{C}_i, \quad \text{with} \quad (12)$$

$$\mathcal{C}_i = \begin{cases} 0 & \text{if } d_i^T > R; \\ 1 & \text{if } d_i^T < \frac{R}{2}; \\ \frac{R-d_i^T}{R/2} & \text{otherwise,} \end{cases} \quad (13)$$

where  $d_i^T$  denotes the distance between robot  $i$  and the perimeter of the prey at time  $T$  (see Table 1). If the prey at time  $T$  is not within the sensing range  $R$  of a robot, the latter receives the lowest possible reward (i.e., 0). Robots that at the end of the trial are still within the initial range ( $\frac{R}{2} = 25$  cm) around the prey receive the maximum reward (i.e., 1). Note that  $\mathcal{C}$  does not impose any bias on the transport strategy: any pulling or pushing arrangement of two robots is assigned the maximum clustering performance.

The transport performance  $\mathcal{T}$  is defined as

$$\mathcal{T} = \Delta(\mathcal{X}^0, \mathcal{X}^T), \quad (14)$$

where  $\mathcal{X}^t$  denotes the position of the prey at time  $t$ , and  $\Delta(\cdot, \cdot)$  is the Euclidean distance.

The performance of an individual is evaluated in  $S = 5$  independent trials. For each trial, the start configuration (e.g., specifying the initial locations of the robots and of the light source) is randomly generated. Every individual within the same generation is evaluated on the same sample of start configurations. The sample is changed once at the beginning of each generation. Let  $\mathcal{Q}_i$  be the quality observed in trial  $i$ , and  $\phi$  be a permutation of  $\{1, 2, \dots, S\}$  so that  $\mathcal{Q}_{\phi(1)} \leq \mathcal{Q}_{\phi(2)} \leq \dots \leq \mathcal{Q}_{\phi(S)}$ . Then the fitness  $\mathcal{F}$ , which is to be maximized, is defined as

$$\mathcal{F} = \frac{2}{S(S+1)} \sum_{i=1}^S (S-i+1) \mathcal{Q}_{\phi(i)}. \quad (15)$$

Note that in this way the trial resulting in the lowest transport quality value (if any) has the highest impact on  $\mathcal{F}$ . Thereby, individuals are penalized for exhibiting high performance fluctuations.

## 3 Results

We conducted 30 independent evolutionary runs for 150 generations each. This corresponds to 15,000 fitness evaluations per run. This limit was defined in order to keep the execution time

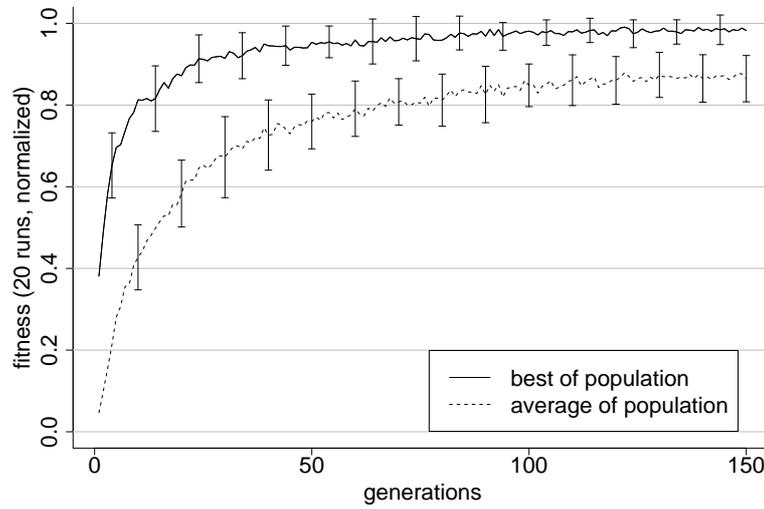


Figure 3: Evolution of transport behaviors with one robot and a 250 g prey. Development of the population best and population average fitness. Each curve corresponds to the average of 20 evolutionary runs with different random seeds. Bars indicate standard deviations.

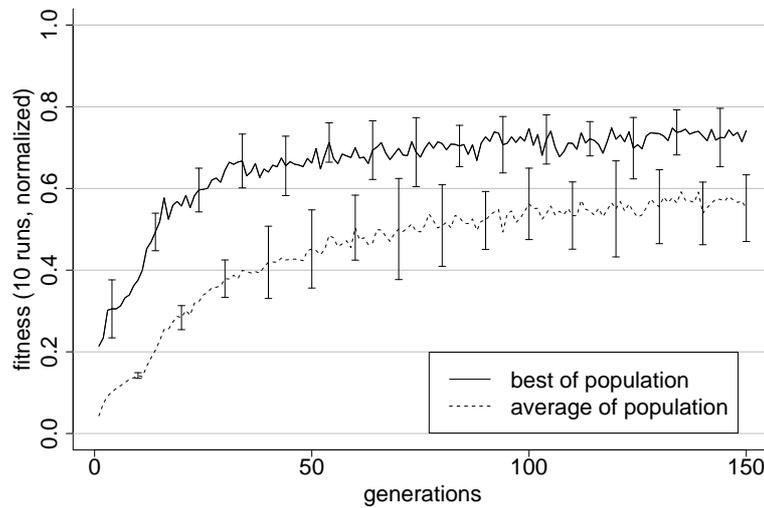


Figure 4: Evolution of transport behaviors with two robots and a 500 g prey. Development of the population best and population average fitness. Each curve corresponds to the average of 10 evolutionary runs with different random seeds. Bars indicate standard deviations.

per run within a time frame of 1–4 days. In 20 runs, the fitness of individuals reflected the performance in solitary transport (i.e., simulations with a single robot and a prey of weight 250 g), whereas in the other 10 runs, the fitness reflected the performance in group transport (i.e., simulations with two robots and a prey of weight 500 g).<sup>4</sup> Recall that the 500 g prey can not be moved by a single robot. Figures 3 and 4 present the corresponding average and maximum fitness time histories. The curves correspond respectively to the average of 20 and 10 runs with different random seeds. The values are normalized in the range  $[0,1]$ . The lower bound is tight, and represents trials in which the prey was not moved and the robots lost visual contact to it (fitness zero). The upper bound corresponds to the maximum distance a robot that is pre-assembled with the lighter (250 g) prey can push the latter within  $T = 20$  s. To compute the upper bound, we disabled any random noise affecting the actuators. The upper bound so computed was 152 cm. We assume this to be also an upper bound for the maximum distance two robots can transport the heavier (500 g) prey during the same time period. By comparing the figures, we can see that the fitness values obtained in the one-robot evolutions (see Figure 3) are higher than the fitness values obtained in the two-robot evolutions (see Figure 4).

### 3.1 Quantitative Analysis

The fitness assigned to a group depends not only on the genotype, but also on other factors, including the robots' initial positions and orientations, the position of the light source in the environment, and the noise affecting the robots' sensors and actuators. Thus, there is a very large number of possible configurations to test. However, the genotype is evaluated only in five trials (per generation) during the evolutionary design phase.

To select the *best* individual of each evolutionary run, we post-evaluate the  $\mu = 20$  best-rated (parent) individuals of the final generation on a random sample of 500 start configurations. For every evolutionary run, the individual exhibiting the highest average performance during the post-evaluation is considered to be the best one.

To allow for an unbiased assessment of the performance of the selected individuals, we post-evaluate each of them for a second time on a new random sample of 500 start configurations.

Let us first consider the performance of the best individuals from the evolutionary runs in which a single robot was simulated.

#### 3.1.1 Individuals Evolved for Solitary Task Performance

Figure 5 illustrates the transport performance of the individuals evolved for solitary task performance using a box-and-whisker plot. The gray boxes correspond to the distances (in cm) the 250 g prey was moved by a single robot in the 500 trials of the post-evaluation. The average distances (in cm) range from 95.0 to 137.9. This is 62.5% to 90.7% of the upper bound. The standard deviations are in the range  $[9.7,35.3]$ . Note that the performance in some trials exceeds the upper bound (indicated by the bold horizontal line). This is caused by the random differences among the actuators of the robots (e.g., differences in the maximum speed  $M$  of a wheel). Recall that to compute the upper bound, any form of random noise was disabled.

We now examine the ability of a group of robots—each acting as in solitary transport—to transport a prey that requires cooperation to be moved. Note that the robots can not perceive each other, nor have they been trained in situations that involve multiple robots. For each individual, we assessed the performance of a group of two robots on 500 start configurations with the 500 g prey. All robots of the group were initially assigned a copy of the same neural network controller. The results are shown in Figure 5 (white boxes). The average distances (in cm) range from 30.4 to 70.1. This is 20.0% to 46.1% of the upper bound. The standard deviations are in the range  $[38.3,53.9]$ . The performance obtained with two robots and the heavy prey is significantly worse than the performance obtained with one robot and the light prey (two-sided Mann-Whitney test, 5% significance level).

Let us now consider the performance of the best individuals from the evolutionary runs in which two robots were simulated.

---

<sup>4</sup>Note that the computational costs may increase super-linearly with the number of robots being simulated. This is particularly the case if the robots physically interact with each other.

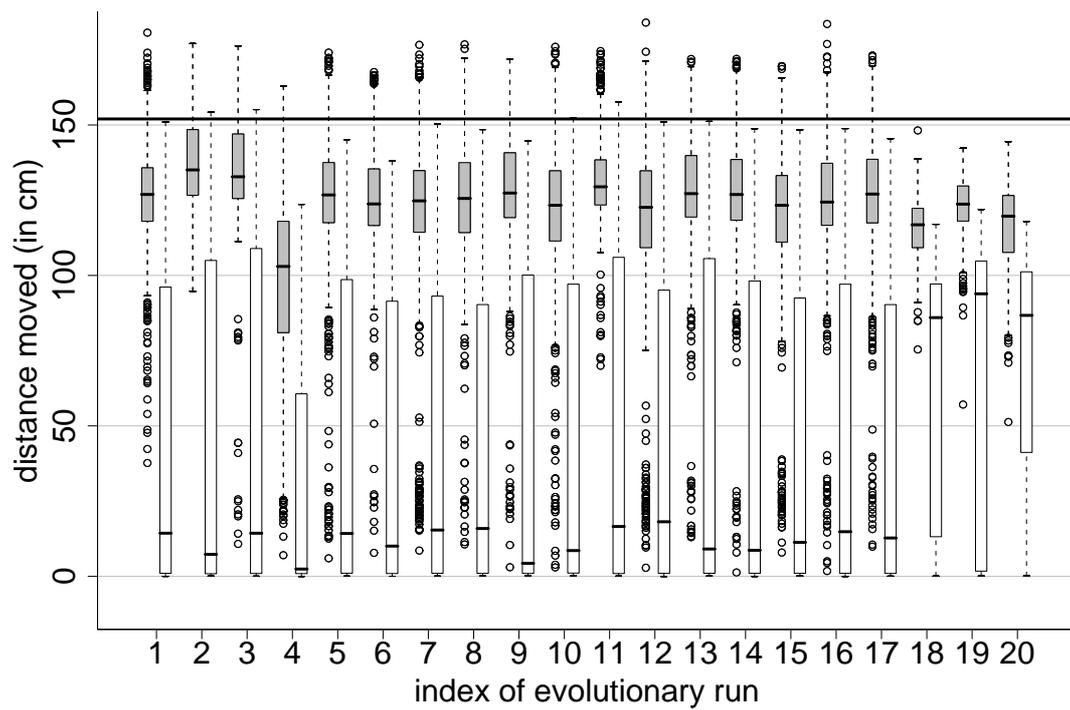


Figure 5: Solitary and group transport performance of the best individuals evolved for solitary transport. Box-and-whisker plot [6] of the distance (in cm) the prey was moved by each individual (500 observations per box). Each box comprises observations ranging from the first to the third quartile. The median is indicated by a bar, dividing the box into the upper and lower part. The whiskers extend to the farthest data points that are within 1.5 times the interquartile range. Outliers are indicated as circles. Characteristics about two types of observations for each individual are displayed: (i) gray boxes refer to solitary transport simulations (one robot, 250 g prey); (ii) white boxes refer to group transport simulations (two robots, 500 g prey).

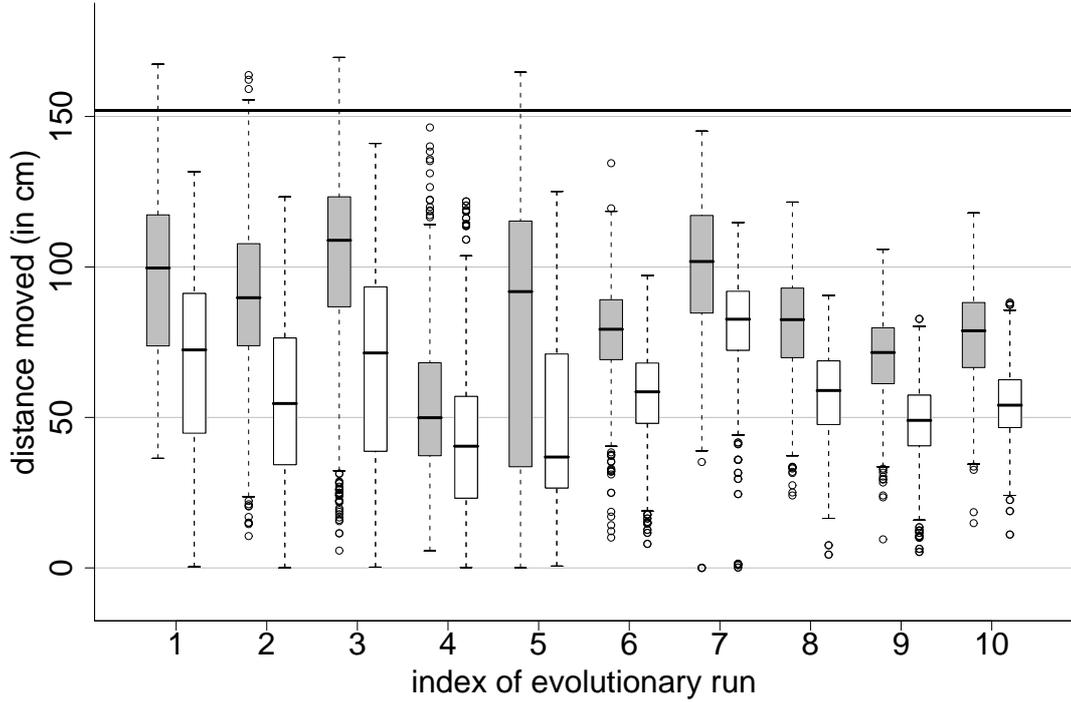


Figure 6: Solitary and group transport performance of the best individuals evolved for group transport. The box-and-whisker plot is explained in the caption of Figure 5.

### 3.1.2 Individuals Evolved for Cooperative Task Performance

Figure 6 illustrates the transport performance of the individuals evolved for group transport using a box-and-whisker plot. Once again, we evaluated both the performance in solitary transport and the performance in group transport in 500 trials each. The gray boxes correspond to the distances (in cm) the 250 g prey was moved by a single robot. The average distances (in cm) range from 53.9 to 101.4. This is 35.4% to 66.7% of the upper bound. The standard deviations are in the range [15.1,40.9]. For the trials with two robots and a 500 g prey (see white boxes), the average distances (in cm) range from 41.6 to 80.9. This is 27.4% to 53.2% of the upper bound. The standard deviations are in the range [12.2,35.6].

Although during evolution two robots were present, the individuals perform consistently better when tested alone (two-sided Mann-Whitney tests, 5% significance level). This latter result supports our intuition that group transport is more complex than solitary transport. The presence of multiple robots is likely to lead to interferences that cause a decrease in performance. Moreover, group transport requires a coordinated action as the members of the group have to push or pull the object in similar directions.

In regard to solitary transport performance, the evolutionary algorithm in which solitary transport behavior is selected for generates better performing individuals than the evolutionary algorithm in which group transport behavior is selected for (one-sided Mann-Whitney test, 5% significance level). In regard to group transport performance, the evolutionary algorithm in which group transport behavior is selected for generates better performing individuals than the evolutionary algorithm in which solitary transport behavior is selected for (one-sided Mann-Whitney test, 5% significance level).

## 3.2 Behavioral Analysis

In the following, we analyze the behaviors of robots when controlled by the neural networks whose parameters are specified by the individuals evolved for solitary and group transport, respectively. We identify proximate mechanisms that cause the coordination of robots in the

group. In particular, we examine the formation of assemblages.

### 3.2.1 Individuals Evolved for Solitary Task Performance

Concerning the 20 runs for the evolution of solitary transport, 17 out of 20 of the best neural networks let the robot grasp and push the prey by moving forward. Over the 500 trials with two robots in the environment with the 500 g prey, in 96.4% to 100.0% of the cases, depending on the neural network used, a robot was connected either directly or indirectly to the prey at the end of the trial. Rarely, self-assemblages—that is, structures of robots being directly connected to each other—were formed (in 0.0 to 8.8% of the trials, respectively). In the majority of cases, the robots failed to push effectively the prey in a common direction. The reason for this poor performance is that the robots’ behavior and the collective structures they form (via connections with either the prey or with each other) are not suited for the accomplishment of the transport task.

The remaining three neural networks (indexed 18–20 in Figure 5) let the robots push the prey with their bodies by moving backward. These networks display a high median performance, even in group transport. To achieve coordination, the robots do not take advantage of the light source.<sup>5</sup> Instead, they interact with each other through the prey. If we assume that each robot pushes towards the center of the prey with the same intensity, the combined force of the two exceeds (in intensity) the force of any of the two, as long as their pushing directions differ by less than  $120^\circ$ . As the two robots are initially randomly distributed around the prey, such degree of coordination is present in about 2/3 of the trials. In most of these cases, the resulting force is sufficient to start moving the prey at low speed. As the robots’ pushing directions intersect with each other, once the prey is in motion the robots approach each other sliding along the perimeter of the prey.<sup>6</sup> As the robots continuously adjust their pushing directions according to the position of the prey (and thus to each other), they self-organize into an effective pushing arrangement.

The latter result shows that behaviors that evolve for solitary task performance can provide mutual benefit once robots start acting in groups. In our case, 15% of the individuals that evolved for solitary transport, once put together with a clone of themselves, exhibit social behavior (mutual benefit); they physically interact with each other (either directly, or through the prey) and thereby enhance their degree of coordination.

### 3.2.2 Individuals Evolved for Cooperative Task Performance

Concerning the 10 runs for the evolution of group transport, 5 out of 10 of the best neural networks let the robots make use of the connection mechanism (corresponding to the first five pairs of boxes in Figure 6). Four out of five neural networks employ the strategy depicted in Figure 7(a): each robot cycles (with the connection mechanism heading forward) around the prey to reach a side correlated with the direction of the light source (e.g., the opposite side). Some neural networks let the robot cycle either counter-clockwise or clockwise depending on which path is shorter. During this phase, the robot remains distant from the prey, and thereby, also from a potential teammate that is already connected to the prey. Once the side that is correlated with the light source is approximately reached, the robot approaches the prey and potentially the connected teammate, and establishes a connection. In other words, by exploiting the relative position of both the prey and the light source, the two robots organize into a dense formation, potentially a linear chain. Each robot keeps on moving forward, pushing the prey (e.g., towards the light source).

The other five neural networks (corresponding to the five latter pairs of boxes in Figure 6) make no use of the connection mechanism. Their strategy is depicted in Figure 7(b). They control the robot to move backward. The robot cycles around the prey to reach a side

---

<sup>5</sup>Only the neural network from run 20 lets the robots (slightly) correlate their direction of pushing with the direction of the light source. The networks from runs 18 and 19, however, do not let the robots correlate their direction of pushing with the direction of the light source. In fact, they let the robots transport the prey in a direction that is uniformly random (as experimentally verified). Recall that the task is to move the prey, the farther the better, in an arbitrary direction.

<sup>6</sup>Recall that the particular behavior that is discussed here lets the robots make no use of their connection mechanisms. Instead, the robots’ bodies are in physical contact with the prey and push the latter by moving backward.

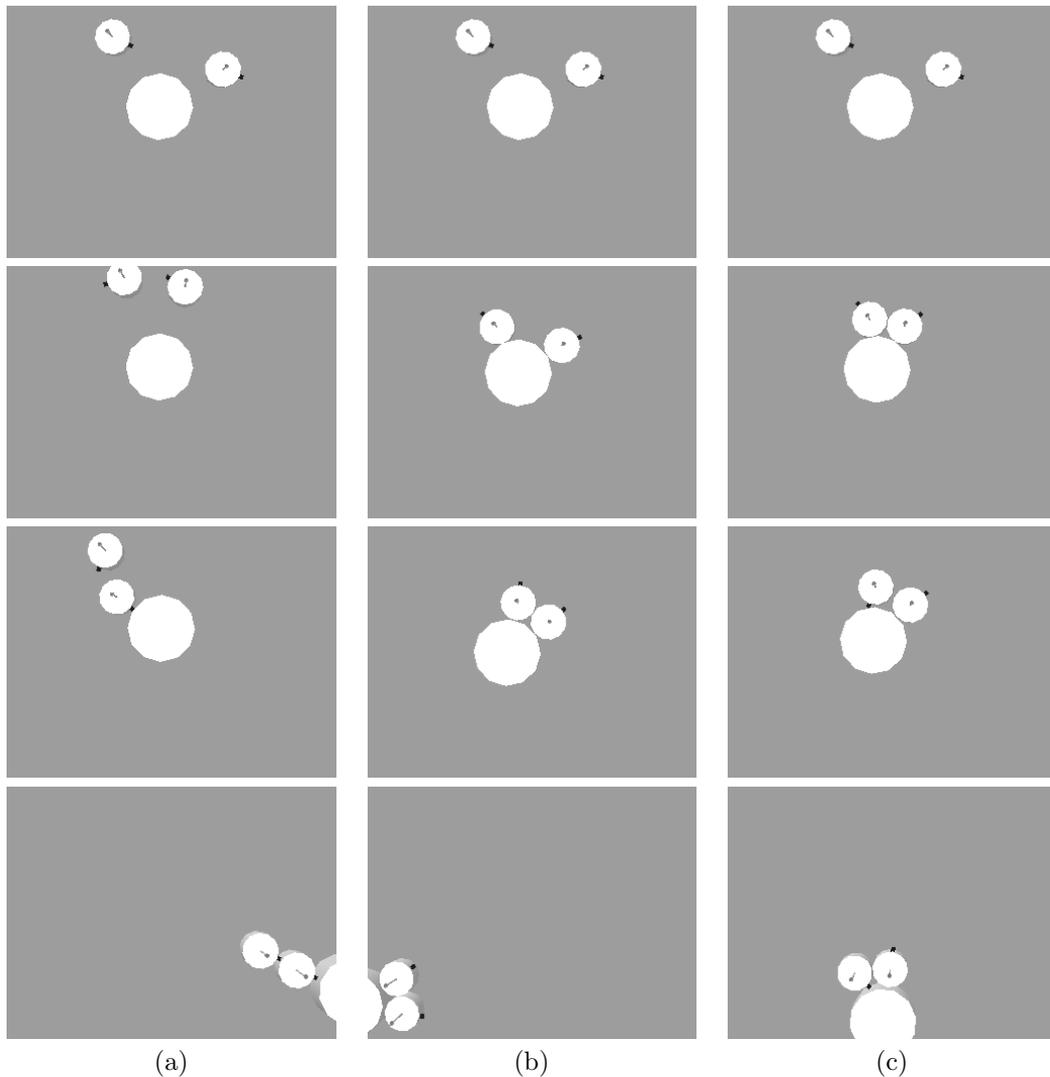


Figure 7: Group transport of a heavy prey in arbitrary direction. The light source is located outside the range of the image. Both robots are controlled by identical recurrent neural networks. Sequences of actions during a trial (from the top to the bottom, at time 0s, 2s, 4s, and 14s, respectively), corresponding to three different neural networks that respectively (a) let the robots assemble with the prey and/or with the teammate and transport the prey by moving forward, (b) let the robots push the prey with their body by moving backward, (c) let each robot either push the prey by moving backward or assemble with the prey or teammate and push by moving forward.

correlated with the light source. Once again, some neural networks let the robot cycle either counter-clockwise or clockwise depending on which path is shorter. Differently from the previous behavior, however, the robot gets into physical contact with the prey while cycling around it. In fact, the robot tries to push the prey with its body, while at the same time sliding along the prey’s perimeter. If multiple robots are present, their behaviors let them organize into a dense, and thus very effective, pushing arrangement (see Figure 7(b)).

One neural network was capable of letting the robots display a combination of both types of behaviors (see Figure 7(c)). In 33.0% of the cases at the end of the trial, one robot was pushing the prey with the body by moving backward, while the other robot was grasping and pushing the prey, or its teammate, by moving forward (recall that both robots were controlled by an identical neural network). The performance the group achieved in this configuration was significantly higher than the performance the group could achieve in any other configuration when controlled by the same neural network (two-sided Mann-Whitney tests, 5% significance level).

We examine the physical structures that “emerged” in more detail (only for the five neural networks that let the robots make use of the connection mechanism). Over the 500 trials with two robots in the environment with the 500 g prey, in 71.1% to 94.6% of the cases, depending on the neural network used, a robot was connected either directly or indirectly to the prey at the end of the trial. Self-assembled structures were formed in 6.8% to 58.8% of the trials, respectively. Compared to individuals that were evolved for solitary transport (and did also make use of the connection mechanism), the increase in the rate of self-assembly is significant (two-sided Mann-Whitney test, 5% significance level). For the networks whose strategy is depicted in Figure 7(a), self-assembled structures, when formed, were in 77.8% to 93.5% of the cases physically attached to the prey. For the network whose strategy is depicted in Figure 7(c), the respective value is 42.1%.

### 3.3 Scalability

We examine to what extent the observed behaviors are scalable, that is, whether the evolved individuals are able to let robots cooperate in the transport of a heavier prey when the group size becomes larger. We focus on the best individuals evolved for group transport. For each run we take the best individual and evaluate it 200 times using a group of five robots and a time period of 30 s. We keep the geometry of the prey identical, but we increase its weight proportionally to the increase in the number of robots (1250 g). The gray boxes of the plot in Figure 8 show the distance (in cm) the prey was moved during these trials. The individuals from the first five evolutionary runs shown from the left in the figure are those that let the robots make use of the connection mechanism to solve the task. Self-assemblages occurred in respectively 89.0%, 99.0%, 92.0%, 59.5%, and 46.0% of the trials. The other five individuals do not let the robots make use of the connection mechanism. Overall, the individuals making use of self-assembly (the average distances are respectively, 14.4, 18.7, 20.0, 3.7, and 3.4 cm) outperform the other individuals (the average distances are respectively, 1.6, 1.8, 1.7, 1.4, and 2.0 cm); two-sided Mann-Whitney test, 5% significance level. The latter individuals are incapable of achieving the task as the prey does not offer enough contact surface for being pushed effectively by more than two robots (see Figure 9). However, if the perimeter of the prey is scaled by the same factor as the weight and the number of robots has increased, all individuals are able to move the prey, and the ones that let the robots not self-assemble exhibit a better performance (see white boxes in Figure 8).

Video recordings are available at <http://iridia.ulb.ac.be/supp/IridiaSupp2007-007>.

## 4 Related Work

In the following, we briefly review related work on group transport and self-assembly. First, we consider studies that are concerned with groups of social insects and of social spiders. Then, we consider studies that are concerned with groups of robots.

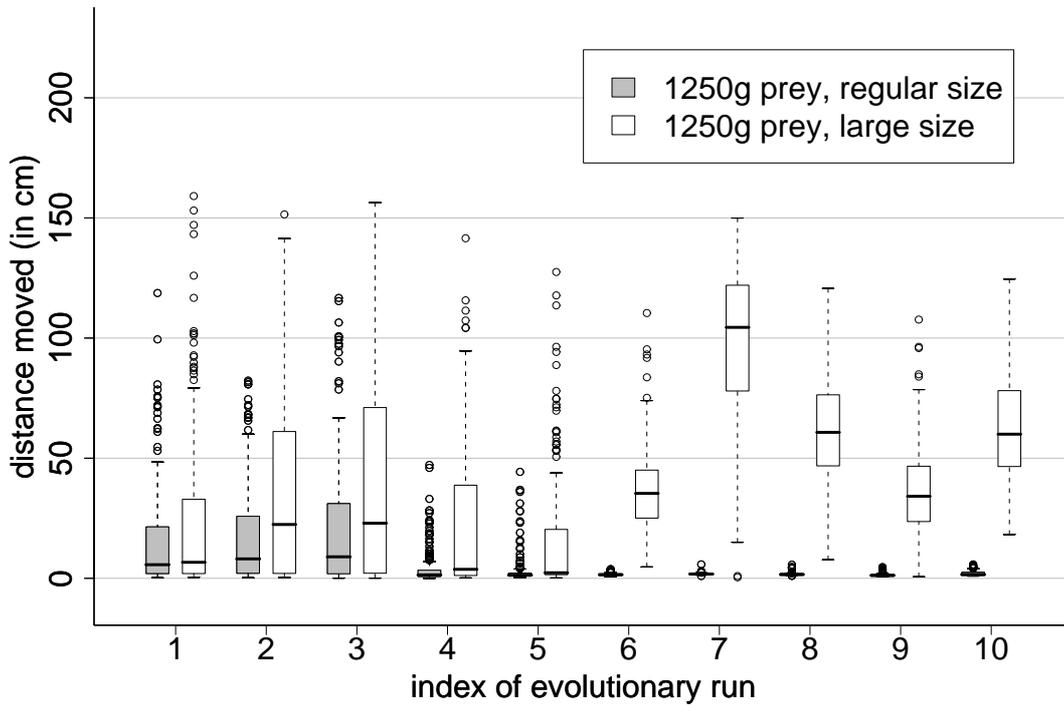


Figure 8: Post-evaluation of the best individuals with groups of five robots transporting a 1250 g prey for 30 seconds (200 observations per box). Individuals labeled 1–5 are those that let robots self-assemble, all other make no use of the connection mechanism. Geometry of the prey: (a) size equal to setup during evolution (radius: 12 cm), (b) size scaled by the factor the prey’s weight as well as the number of robots has increased (radius: 30 cm). The box-and-whisker plot is explained in the caption of Figure 5.

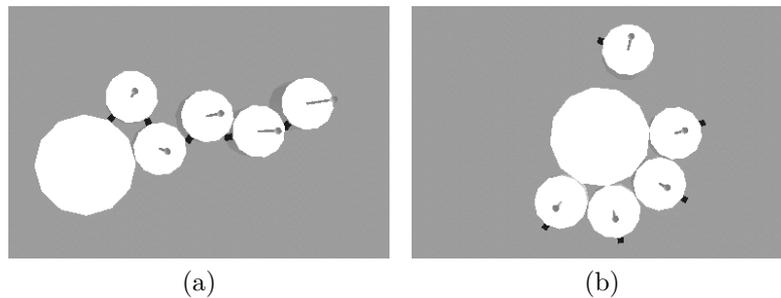


Figure 9: Group transport of a 1250 g prey (radius: 12 cm) by five robots. Snapshots for two different individuals: (a) an individual that let the robots self-assemble. The group is capable of transporting the prey at low speed. (b) An individual that let the robots make no use of the connection mechanism. The group is incapable of moving the prey as the latter offers not enough contact surface for being pushed effectively by more than two robots. For a quantitative analysis see Figure 8.

## 4.1 Groups of Social Insects and of Social Spiders

### 4.1.1 Group Transport

In the literature, group transport is almost exclusively reported in the context of ants. In fact, Moffett [55, page 220] claimed that group transport “is better developed in ants than in any other animal group”. Nevertheless, it “has seldom been recognized as a form of social behavior that is worthy of investigation in its own right” (page 227).

In most ant species, group transport presumably provides adaptive value as reproductive immatures are much bigger than workers, and therefore can not be transported by a single worker alone (e.g., during an emigration). Moffett [55, page 220] states that group transport of “bulky larvae and pupae is probably nearly universal in ants and is likely to have preceded the transport of food by this method”. Moffett [55] lists 39 species of ants for which group transport of food has been reported. He states that “without doubt the group transport of food has arisen independently in numerous phylogenetic lines”. “At least with regard to carrying food, those ants species capable of group transport are unquestionably in the minority” (page 227).

Almost half a century ago, Sudd [72] studied the transport of prey by single ants and by groups of ants of the species *Pheidole crassinoda*. Sudd reported that during transport the ants did not pull steadily but in short successive hauls that were generally associated with changes in the arrangement of ants in the group. “In almost all series involving groups of ants there was an upward trend of the force exerted in successive hauls; where only one ant was pulling however the proportion of hauls with upward and downward trend was about equal” (page 301). Changes in the arrangement of ants in the group were of two types [72]:

In *realignment* the ant altered the orientation of its body without releasing its hold on the prey. Realignment was sometimes the cause and sometimes the effect of rotation of the prey. In *repositioning* however the ant released the prey and returned to it at a different position. Realignment appeared to correspond to the turning movements of a single ant experiencing difficulty in pulling prey, whilst repositioning corresponded to the excursions which were made from the prey before an ant left it to return to the nest. . . . Realignment occurred throughout traction but repositioning involved a sharper change and was more occasional. (pages 301 and 304)

Even though a positive group effect was present, the behavior of individual workers in group transport appeared “to contain no elements of behavior that were not shown by single transporting ants. . . . If cooperative transport existed therefore it resulted from the coordination, within the group, of behavior also shown by individuals working alone” (page 304).

Franks [22] and Franks *et al.* [24, 23] investigated the performance and organization of groups of army ant workers (*Eciton burchellii* and *Dorylus wilverthi*), who cooperate to transport large prey. Army ants carry items by first straddling them so that the item is slung beneath their bodies and, hence, they always face the same direction. In contrast, other ants such as *Pheidole crassinoda* tend to pull the item, and often several ants pull in different directions. Franks [22] and Franks *et al.* [23] showed that in most of the instances involving the army ants, the group was composed of an unusually large front-runner, that presumably steered and determined the direction of transport, and one or more particularly small followers. Anderson and Franks [2] do not consider the front-runner as a leader in any sense. Instead, they hypothesize that “all of the individuals that form a team in army ants are initially using exactly the same rules of thumb” (page 537). Franks [22] reported that the performance in the group was much more than the sum of the performances of its individual members. They could do so probably because by straddling the prey between them the rotational forces (i.e., forces that occur when lifting the prey in a position aside its barycenter) are balanced and disappear.

Super-efficiency in group transport has also been observed in other ant species [39, 54]. In *Pheidologeton diversus*, for instance, on average an ant engaged in group transport held at least 10 times the weight it did as solitary transporter [54]. Moffett characterized the behaviors in solitary and group transport as follows:

Ants carrying burdens solitarily grasped them between their mandibles, lifting the burdens from the ground and holding the burdens ahead of them as they walked

forward (burdens were rarely slung beneath the body). Group-transporting ants carried burdens differently. One or both forelegs were placed on the burden, and appeared to aid considerably in lifting it. The mandibles were open and usually lay against the burden, but the burden's surface was seldom gripped between them. Burdens small enough to be carried by few ants, as well as small appendages extending from larger burdens, were often carried partly by ants that behaved as described for those engaged in solitary transport. (page 388)

Group transport of prey has also been observed in a few species of social spiders [81]:

During transport, as an aid to the movement of the prey, spiders weave silk that we named 'traction silk', fixed between the prey and the web (in the direction of the shelter) that will permit a slight lifting of the prey. This process will be repeated until the prey has been transported under the shelter. (page 765)

Coordination in group transport by social spiders seems to occur through the item that is transported [81]:

Movement of one spider engaged in group transport is likely to modify the stimuli perceived by the other group members (such as vibration produced, or indirectly, available site on the prey) possibly producing, in turn, recruitment or departure of individuals. . . . Coordination in spider colonies is based on signals that are made inadvertently as side products of their activities. The communal network, as a means of information, seems to be at the origin of cooperation. This supports the hypothesis of a sudden passage from solitary to social life in spinning spiders [80, 9, 86, 63]. (pages 770–771)

#### 4.1.2 Self-Assembly

Following Whitesides and Grzybowski [85], self-assembly can be defined as a process by which pre-existing discrete components organize into patterns or structures without human intervention.

Self-assembly has been widely observed in social insects [70, 4]. Via self-assembly, ants, bees, and wasps can organize into functional units at an intermediate level between the individual and the colony. Anderson *et al.* [4] identified 18 distinct types of self-assembled structures that insects build: “bivouacs, bridges, curtains, droplets, escape droplets, festoons, fills, flanges, ladders, ovens, plugs, pulling chains, queen clusters, rafts, swarms, thermoregulatory clusters, tunnels, and walls” (page 99). In some cases (e.g., an ant raft) the individuals assemble into “a formless random arrangement”, whereas in other cases (e.g., an ant ladder) the individuals assemble into a “particular (required) arrangement” (page 100). The function of self-assemblages “can be grouped under five broad categories which are not mutually exclusive: (1) defense, (2) pulling structures, (3) thermoregulation, (4) colony survival under inclement conditions, and (5) ease of passage when crossing an obstacle” (page 99). Anderson *et al.* [4] claim that in almost all of the observed instances, the function could not be achieved without self-assembly.

Pulling structures have been observed in a few ant species (e.g., *Eciton burchellii*) as well as in honey bees (*Apis mellifera*) [4]. The structures formed generate torque, for instance, to fix a large prey to the floor or to bend a leaf during nest construction. Although a pulling structure may only require a few individuals, often a critical density of individuals may be required to initiate self-assembly and growth [4]. When part of an assembled structure, ants have few degrees of mobility. In some species, worker ants seem even to become motionless as a reaction to being stretched (see [4] and references therein).

At present “virtually nothing is known regarding the rules, signals, and cues used by individuals in formation [of assembled structures] or the physical constraints these structures are under” [4, page 107]. Lioni *et al.* [48, 49] studied mechanisms by which ants of the genus *Ecophylla* form living ladders and bridges by linking with each other. They showed that the ants, if offered two alternative sites to bridge an empty space, typically end up in a single, large aggregate in either one of the two sites. They observed that the process is controlled by the individual probabilities of entering and leaving the aggregates. The probabilities depend on the number of ants in the aggregate.

Theraulaz *et al.* [77] modeled self-assembly processes in *Linepithema humile* using an agent-based approach. The ants aggregated at the end of a rod and formed droplets containing several assembled ants that eventually fell down. The model could be tuned to reproduce some properties of the experimental system, such as the droplet size and the inter-drop interval. The function of this behavior is currently unknown.

## 4.2 Groups of Robots

### 4.2.1 Group Transport

In most studies of transport in robotic groups, the members of the group move an object by pushing it. Pushing strategies have the advantage that they allow the robots to move objects that are hard to grasp. In addition, multiple objects can be pushed at the same time. On the other hand, it is difficult to predict the motion of the object and of the robots, especially if the ground is not uniform.<sup>7</sup> Therefore, the control typically requires sensory feedback. Most studies consider two robots pushing a wide box simultaneously from a single side [52, 74, 15, 61, 29]. To coordinate the robots' actions, robots are specifically arranged [52, 15, 61, 29], control is synchronized [52], relative positions are known [15, 61], explicit communication is used [52, 61], and/or individual tasks are generated by a designated leader agent [29, 74].

Only few studies considered more than two robots, pushing a box simultaneously [45, 87, 46, 44, e.g.]. In these cases, the control is homogeneous and decentralized; the robots make no use of explicit communication. Kube and Zhang [46] and Kube *et al.* [44] reported that if the box is small compared to the size of the pushing robots the performance decreases drastically with group size as the box offers only limited contact surface.

Many studies considered the transport of an object by multiple, mobile robots grasping and/or lifting it. In these studies, typically 2–3 robots are manually attached to the object [14, 43, 1, 73, 53, 82]. To coordinate the robots' actions, often, robots have knowledge of their relative positions. In some systems the desired trajectories are given prior to experimentation to all robots of the group. The object is transported as each robot follows the given trajectory by making use of dead-reckoning [14]. In other systems, the manipulation is planned in real-time by an external workstation, which communicates with the robots [53]. Often, instead of an external computer, a specific robot called the *leader* knows the desired trajectory or the goal location. The leader robot can send explicit high- or low-level commands to the *followers* [73, 82]. However, in many leader-follower systems explicit communication is not required [43, 1].

### 4.2.2 Self-Assembly

Self-reconfigurable robots [89, 68] hold the potential to self-assemble and thus to mimic the complex behavior of social insects. In current implementations [57, 89, 68, 40], however, single modules usually have highly limited autonomous capabilities (when compared to an insect). Typically, they are not equipped with sensors to perceive the environment. Nor, typically, are they capable of autonomous motion. These limitations, common to most self-reconfigurable robotic systems, make it difficult to let separate modules, or groups of modules, connect autonomously. In some systems, self-assembly was demonstrated with the modules being pre-arranged at known positions [90, 92]. Some instances of less constrained self-assembly have been reported (for an overview see [34]). Fukuda *et al.* [25, 27] demonstrated self-assembly among robotic cells using the CEBOT system [26]. In the experiment, a moving cell approached and connected with a static cell. The moving cell was controlled with a finite-state automata. Rubenstein *et al.* [67] demonstrated the ability of two modular robots to self-assemble. Each robot consisted of a chain of two linearly-linked CONRO modules [11]. The robot chains were set up at distances of 15 cm, facing each other with an angular displacement not larger than 45°. The control was heterogeneous, both at the level of individual modules within each robot and at the level of the modular makeup of both robots. Trianni *et al.* [78] evolved a neural network controller for individual and collective phototaxis. The results obtained using embodied simulations demonstrate that groups of up to three mobile robots while performing phototaxis can self-assemble and disassemble in response to the temperature

---

<sup>7</sup>For a theory on the mechanics of pushing see Mason [51].

of the environment. The use of self-assembly was explicitly rewarded by the fitness function. None of these studies demonstrate the utility of self-assembling robots for a concrete task.

## 5 Discussion and Conclusions

This paper studies the evolution of cooperation by investigating the relation between solitary and social behavior: the question that is addressed is to what extent individuals that as a group have to accomplish a cooperative task can benefit from behaving differently from individuals that have to accomplish tasks on their own. Related work from biology suggests that, in some species, individuals show no difference in behavior when engaged in solitary and group transport, respectively. Following this, we hypothesize that group transport in some species evolved from solitary transport.

We designed a study that allows for some initial test on the plausibility of such a transition from solitary to social behavior. In particular, we studied the solitary and group transport of an object called prey using physics-based computer simulations. We simulated simplistic robots that, similarly to ants, can move autonomously and grasp objects in their environment. The robots' cognitive capabilities were very limited and similar to those of many solitary animals: they were not capable of perceiving potential teammates or of communicating with them. We performed two sets of experiments to produce (using evolutionary algorithms) neural networks that control robots when competing based on solitary task performance and cooperative task performance, respectively.

Each of the evolutionary runs we performed—30 in total—yielded an individual of satisfactory performance (the individuals achieved 62.5%–90.7% and 27.4%–53.2% of an upper bound performance for solitary transport and group transport, respectively). Networks evolved for solitary task performance were capable of letting robots engage also in group transport (the individuals achieved 20.0%–46.1% of the upper bound performance). Networks evolved for group transport, were capable of letting robots engage also in solitary transport (the individuals achieved 35.4%–66.7% of the upper bound performance). Our results show also that group transport and solitary transport impose different demands on the robot. In fact, robots engaged in group transport benefit from behaving differently from those engaged in solitary transport. The results support also our intuition that group transport is a more complex task than solitary transport. Furthermore, our study revealed a variety of proximate mechanisms that can cause coordinated behavior. In particular, we observed that robots enhance their degree of coordination by physically interacting either directly or indirectly, that is, via the environment, and by exploiting the presence of visual cues in the environment. This holds not only for all networks evolved for cooperative task performance, but also for 15% of the networks evolved for solitary task performance. That is, some “solitary individuals”, when grouped together with a clone, exhibit social behavior (mutual benefit).

In half of the evolutionary runs with groups of robots, the best-rated neural networks let robots self-assemble. Self-assembled structures among the robots were formed in up to 58.8% of the trials with two robots and up to 99.0% of the trials with five robots. Note that the fitness function of the evolutionary algorithm did neither explicitly reward the robots for self-assembling, nor did it impose any bias concerning the spatial organization of the robots during task performance. The “emergence” of self-assembly is a striking result, confirming that such capability—as in social insects—can provide adaptive value to the group. The analysis revealed the proximate mechanisms that caused the formation of the self-assemblages. In particular, two visual cues present in the environment (the prey and a light source) were sufficient to guide robots in the formation of assembled structures. Moreover, as the assemblages were formed by robots facing in approximately the same direction (and sometimes opposite directions), they were suitable for the accomplishment of the task (which required all robots to pull/push the prey in approximately the same direction). When increasing both the group size and the weight of the object, the performance decreased. Groups that self-assembled were still capable of moving the prey. Groups that did not make use of self-assembly were incapable of moving the prey, unless the prey size was increased proportionally.

Overall, our work supports the hypothesis that in some species group transport has evolved from solitary transport, presumably from situations in which solitary transporters, without being aware of each other, cooperatively transported a common load. Because such behaviors could provide adaptive value, it is likely they got improved through evolution (which

potentially led to specialized behaviors for group transport). To the best of our knowledge, although plausible, this hypothesis has not been further investigated by biologists. Such transition would require (as was the case in our study) that multiple individuals were likely to interact with a same load, and that these individuals were related to each other. Future work will investigate more comprehensive models (e.g., including the spatial distribution of competing colonies, and the spatial distribution of different types of prey), and compare the results against knowledge on biological systems.

From an engineering perspective, our work showed that cooperative behavior can be achieved by simple means—the robots could neither perceive nor communicate with each other directly. Nevertheless, they could organize into effective transport arrangements.

In general, the study of artificial systems in which self-assembly evolves as providing some functional benefit to the interacting individuals can offer invaluable insight into biological self-assembling systems, and is also a promising direction to be further explored in the evolutionary design of robotic systems. Our work provided some first evidence that self-assembly can provide adaptive value to a group of robots. The value certainly depends on many factors (e.g., mass and size of the prey, number and size of the robots). In ants, self-assembled structures are rarely observed in group transport (yet they are relevant in other context). Individual limitations might prevent individuals, which are organized into an assemblage, from moving collectively by their own propulsion (see related work section). This could explain why self-assemblages are virtually non-existing for group transport in ants. Our scalability analysis indicates that our simple robots might be subject to similar limitations when the group size is increased. For designing group transport systems, one can explore more capable robotic hardware. Following this approach we have designed group transport systems that are effective also for bigger group sizes [33, 32, 79, 35, 36].

## Acknowledgements

This work was supported by the IST FET project “SWARM-BOTS” (in the form of grant number IST-2000-31010), and by the “ANTS” project, an “Action de Recherche Concertée” funded by the Scientific Research Directorate of the French Community of Belgium. Roderich Groß acknowledges support from the BBSRC (in the form of grant number E19832). Marco Dorigo acknowledges support from the Belgian F.R.S.–FNRS, of which he is a Research Director. The authors thank three anonymous referees, Christos Ampatzis, and Nigel R. Franks for helpful comments on a previous version of this paper.

## References

- [1] 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(1):13–19, 1999.
- [2] C. Anderson and N. R. Franks. Teams in animal societies. *Behavioral Ecology*, 12(5):534–540, 2001.
- [3] C. Anderson and D. W. McShea. Intermediate-level parts in insect societies: Adaptive structures that ants build away from the nest. *Insectes Sociaux*, 48(4):291–301, 2001.
- [4] C. Anderson, G. Theraulaz, and J.-L. Deneubourg. Self-assemblages in insect societies. *Insectes Sociaux*, 49(2):99–110, 2002.
- [5] G. Baldassarre, S. Nolfi, and D. Parisi. Evolving mobile robots able to display collective behaviour. *Artificial Life*, 9(3):255–267, 2003.
- [6] R. A. Becker, J. M. Chambers, and A. R. Wilks. *The new S language. A programming environment for data analysis and graphics*. Chapman & Hall, London, 1988.
- [7] H.-G. Beyer. *The theory of evolution strategies*. Springer Verlag, Berlin, Germany, 2001.
- [8] A. F. G. Bourke and N. R. Franks. *Social evolution in ants*. Monographs in behavior and ecology. Princeton Univ. Press, Princeton, NJ, 1995.
- [9] V. Brach. *Anelosimus studiosus* (Araneae: Theridiidae) and the evolution of quasisociality in theridiid spiders. *Evolution*, 31(1):154–161, 1977.

- [10] S. Camazine, J.-L. Deneubourg, N. R. Franks, J. Sneyd, G. Theraulaz, and E. Bonabeau. *Self-organization in biological systems*. Princeton Univ. Press, Princeton, NJ, 2001.
- [11] A. Castano, A. Behar, and P. M. Will. The Conro modules for reconfigurable robots. *IEEE/ASME Transactions on Mechatronics*, 7(4):403–409, 2002.
- [12] C. A. Coulomb. *Théorie des machines simples, en ayant égard au frottement de leurs parties et à la roideur des cordages (in French)*. Adamant Media Corporation, Boston, 2001. Reprint of an 1821 edition by Bachelier, Paris.
- [13] R. Dawkins. *The selfish gene*. Oxford Univ. Press, New York, 1976.
- [14] J. Desai, C.-C. Wang, M. Žefran, and V. Kumar. Motion planning for multiple mobile manipulators. In *Proc. of the 1996 IEEE Int. Conf. on Robotics and Automation*, volume 3, pages 2073–2078. IEEE Computer Society Press, Los Alamitos, CA, 1996.
- [15] B. R. Donald, J. Jennings, and D. Rus. Information invariants for distributed manipulation. *International Journal of Robotics Research*, 16(5):673–702, 1997.
- [16] M. Dorigo. Swarm-bot: An experiment in swarm robotics. In *Proc. of the 2005 IEEE Swarm Intelligence Symp.*, pages 192–200. IEEE Computer Society Press, Los Alamitos, CA, 2005.
- [17] M. Dorigo, V. Trianni, E. Şahin, R. Groß, T. H. Labelle, 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.
- [18] M. Dorigo, E. Tuci, V. Trianni, R. Groß, S. Nouyan, C. Ampatzis, T. H. Labelle, R. O’Grady, M. Bonani, and F. Mondada. SWARM-BOT: Design and implementation of colonies of self-assembling robots. In G. Y. Yen and D. B. Fogel, editors, *Computational Intelligence: Principles and Practice*, pages 103–135. IEEE Computer Society Press, Los Alamitos, CA, 2006.
- [19] J. L. Elman. Finding structure in time. *Cognitive Science*, 14(2):179–211, 1990.
- [20] D. Floreano, S. Mitri, S. Magnenat, and L. Keller. Evolutionary conditions for the emergence of communication in robots. *Current Biology*, 17(6):514–519, 2007.
- [21] K. R. Foster, T. Wenseleers, and F. L. W. Ratnieks. Kin selection is the key to altruism. *TRENDS in Ecology & Evolution*, 21(2):57–60, 2006.
- [22] N. R. Franks. Teams in social insects: Group retrieval of prey by army ants (*Eciton burchelli*, hymenoptera: Formicidae). *Behavioral Ecology and Sociobiology*, 18(6):425–429, 1986.
- [23] N. R. Franks, A. B. Sendova-Franks, and C. Anderson. Division of labour within teams of new world and old world army ants. *Animal Behaviour*, 62(4):635–642, 2001.
- [24] N. R. Franks, A. B. Sendova-Franks, J. Simmons, and M. Mogie. Convergent evolution, superefficient teams and tempo in Old and New World army ants. *Proceedings of the Royal Society B—Biological Sciences*, 266(1429):1697–1701, 1999.
- [25] T. Fukuda, S. Nakagawa, Y. Kawauchi, and M. Buss. Self organizing robots based on cell structures—CEBOT. In *Proc. of the 1988 IEEE Int. Workshop on Intelligent Robots*, pages 145–150. IEEE Computer Society Press, Los Alamitos, CA, 1988.
- [26] T. Fukuda and T. Ueyama. *Cellular robotics and micro robotic systems*. World Scientific Publishing, London, 1994.
- [27] T. Fukuda, T. Ueyama, and K. Sekiyama. Distributed intelligent systems in cellular robotics. In S. G. Tzafestas and H. B. Verbruggen, editors, *Artificial Intelligence in Industrial Decision Making, Control and Automation*, pages 225–246. Kluwer Academic Publishers, Dordrecht, The Netherlands, 1995.
- [28] S. Garnier, J. Gautrais, and G. Theraulaz. The biological principles of swarm intelligence. *Swarm Intelligence*, 1(1):3–31, 2007.
- [29] B. P. Gerkey and M. J. Matarić. Sold!: Auction methods for multirobot coordination. *IEEE Transactions on Robotics and Automation*, 18(5):758–768, 2002.

- [30] P.-P. Grassé. La reconstruction du nid et les coordinations interindividuelles chez *Belligoseria natalensis* et *Cubitermes sp.* La théorie de la stigmergie: Essai d'interprétation du comportement des termites constructeurs (in French). *Insectes Sociaux*, 6(1):41–80, 1959.
- [31] R. Groß, M. Bonani, F. Mondada, and M. Dorigo. Autonomous self-assembly in swarm-bots. *IEEE Transactions on Robotics*, 22(6):1115–1130, 2006.
- [32] R. Groß and M. Dorigo. Cooperative transport of objects of different shapes and sizes. In *Proc. of the 4<sup>th</sup> Int. Workshop on Ant Colony Optimization and Swarm Intelligence (ANTS 2004)*, volume 3172 of *Lecture Notes in Computer Sciences*, pages 107–118. Springer Verlag, Berlin, Germany, 2004.
- [33] R. Groß and M. Dorigo. Group transport of an object to a target that only some group members may sense. In *Proc. of the 8<sup>th</sup> Int. Conf. on Parallel Problem Solving from Nature (PPSN VIII)*, volume 3242 of *Lecture Notes in Computer Sciences*, pages 852–861. Springer Verlag, Berlin, Germany, 2004.
- [34] R. Groß and M. Dorigo. Self-assembly at the macroscopic scale. *Proceedings of the IEEE*, 2008. Accepted for publication.
- [35] R. Groß, F. Mondada, and M. Dorigo. Transport of an object by six pre-attached robots interacting via physical links. In *Proc. of the 2006 IEEE Int. Conf. on Robotics and Automation*, pages 1317–1323. IEEE Computer Society Press, Los Alamitos, CA, 2006.
- [36] R. Groß, E. Tuci, M. Dorigo, M. Bonani, and F. Mondada. Object transport by modular robots that self-assemble. In *Proc. of the 2006 IEEE Int. Conf. on Robotics and Automation*, pages 2558–2564. IEEE Computer Society Press, Los Alamitos, CA, 2006.
- [37] W. D. Hamilton. The genetic evolution of social behaviour. I. *Journal of Theoretical Biology*, 7(1):1–16, 1964.
- [38] I. Harvey, P. Husbands, D. Cliff, A. Thompson, and N. Jakobi. Evolutionary robotics: The Sussex approach. *Robotics and Autonomous Systems*, 20(2–4):205–224, 1997.
- [39] B. Hölldobler, R. C. Stanton, and H. Markl. Recruitment and food-retrieving behavior in *Novomessor* (Formicidae, Hymenoptera). *Behavioral Ecology and Sociobiology*, 4(2):163–181, 1978.
- [40] M. W. Jørgensen, E. H. Østergaard, and H. H. Lund. Modular ATRON: Modules for a self-reconfigurable robot. In *Proc. of the 2004 IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, volume 2, pages 2068–2073. IEEE Computer Society Press, Los Alamitos, CA, 2004.
- [41] S. Klumpp and R. Lipowsky. Cooperative cargo transport by several molecular motors. *Proceedings of the National Academy of Sciences of the United States of America*, 102(48):17284–17289, 2005.
- [42] J. Korb and J. Heinze. Multilevel selection and social evolution of insect societies. *Naturwissenschaften*, 91(6):291–304, 2004.
- [43] K. Kosuge and T. Oosumi. Decentralized control of multiple robots handling an object. In *Proc. of the 1996 IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, volume 1, pages 318–323. IEEE Computer Society Press, Los Alamitos, CA, 1996.
- [44] C. R. Kube and E. Bonabeau. Cooperative transport by ants and robots. *Robotics and Autonomous Systems*, 30(1–2):85–101, 2000.
- [45] C. R. Kube and H. Zhang. Collective robotics: From social insects to robots. *Adaptive Behavior*, 2(2):189–218, 1993.
- [46] C. R. Kube and H. Zhang. Task modelling in collective robotics. *Autonomous Robots*, 4(1):53–72, 1997.
- [47] K.-H. Liang, X. Yao, Y. Liu, C. S. Newton, and D. Hoffman. An experimental investigation of self-adaptation in evolutionary programming. In *Proc. of the 7<sup>th</sup> Int. Conf. on Evolutionary Programming*, volume 1447 of *Lecture Notes in Computer Sciences*, pages 291–300. Springer Verlag, Berlin, Germany, 1998.
- [48] A. Lioni and J.-L. Deneubourg. Collective decision through self-assembling. *Naturwissenschaften*, 91(5):237–241, 2004.

- [49] A. Lioni, C. Sauwens, G. Theraulaz, and J.-L. Deneubourg. Chain formation in *Ecophylla longinoda*. *Journal of Insect Behavior*, 14(5):679–696, 2001.
- [50] A. Martinoli. *Swarm Intelligence in Autonomous Collective Robotics: From Tools to the Analysis and Synthesis of Distributed Control Strategies*. PhD thesis, EPFL, Lausanne, Switzerland, 1999.
- [51] M. T. Mason. Mechanics and planning of manipulator pushing operations. *International Journal of Robotics Research*, 5(3):53–71, 1986.
- [52] M. J. Matarić, M. Nilsson, and K. T. Simsarian. Cooperative multi-robot box-pushing. In *Proc. of the 1995 IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, volume 3, pages 556–561. IEEE Computer Society Press, Los Alamitos, CA, 1995.
- [53] N. Miyata, J. Ota, T. Arai, and H. Asama. Cooperative transport by multiple mobile robots in unknown static environments associated with real-time task assignment. *IEEE Transactions on Robotics and Automation*, 18(5):769–780, 2002.
- [54] M. W. Moffett. Cooperative food transport by an Asiatic ant. *National Geographic Research*, 4(3):386–394, 1988.
- [55] M. W. Moffett. Ant foraging. *National Geographic Research & Exploration*, 8(2):220–231, 1992.
- [56] F. Mondada, L. M. Gambardella, D. Floreano, S. Nolfi, J.-L. Deneubourg, and M. Dorigo. The cooperation of swarm-bots: Physical interactions in collective robotics. *IEEE Robotics and Automation Magazine*, 12(2):21–28, 2005.
- [57] S. Murata, E. Yoshida, A. Kamimura, H. Kurokawa, K. Tomita, and S. Kokaji. M-TRAN: Self-reconfigurable modular robotic system. *IEEE/ASME Transactions on Mechatronics*, 7(4):431–441, 2002.
- [58] S. Nolfi and D. Floreano. *Evolutionary robotics—The biology, intelligence, and technology of self-organizing machines*. MIT Press, Cambridge, MA, 2000.
- [59] S. Nouyan, R. Groß, M. Bonani, F. Mondada, and M. Dorigo. Group transport along a robot chain in a self-organised robot colony. In *Proc. of the 9<sup>th</sup> Int. Conf. on Intelligent Autonomous Systems*, pages 433–442. IOS Press, Amsterdam, 2006.
- [60] L. Panait and S. Luke. Learning ant foraging behaviors. In *Proc. of the 9<sup>th</sup> Int. Conf. on the Simulation and Synthesis of Living Systems (Artificial Life IX)*, pages 569–574. MIT Press, Cambridge, MA, 2003.
- [61] L. E. Parker. Adaptive heterogeneous multi-robot teams. *Neurocomputing*, 28(1–3):75–92, 1999.
- [62] A. Pérez-Urbe, D. Floreano, and L. Keller. Effects of group composition and level of selection in the evolution of cooperation in artificial ants. In *Proc. of the 7<sup>th</sup> European Conf. on Artificial Life*, volume 2801 of *Lecture Notes in Artificial Intelligence*, pages 128–137. Springer Verlag, Berlin, Germany, 2003.
- [63] C. Plateaux-Quénu, A. Horel, and C. Roland. A reflection on social evolution in two different groups of arthropods: halictine bees (Hymenoptera) and spiders (Arachnida). *Ethology, Ecology & Evolution*, 9(2):183–196, 1997.
- [64] J. Pugh and A. Martinoli. Multi-robot learning with particle swarm optimization. In *Proc. of the 5<sup>th</sup> Int. Joint Conf. on Autonomous Agents and Multiagent Systems (AA-MAS’06)*, pages 441–448. ACM Press, New York, 2006.
- [65] M. Quinn, L. Smith, G. Mayley, and P. Husbands. Evolving controllers for a homogeneous system of physical robots: Structured cooperation with minimal sensors. *Philosophical Transactions of the Royal Society A—Mathematical Physical and Engineering Sciences*, 361(1811):2321–2343, 2003.
- [66] C. W. Reynolds. An evolved, vision-based behavioral model of coordinated group motion. In *Proc. of the 2<sup>nd</sup> Int. Conf. on Simulation of Adaptive Behavior*, pages 384–392. MIT Press, Cambridge, MA, 1993.
- [67] M. Rubenstein, K. Payne, P. Will, and W.-M. Shen. Docking among independent and autonomous CONRO self-reconfigurable robots. In *Proc. of the 2004 IEEE Int. Conf. on Robotics and Automation*, volume 3, pages 2877–2882. IEEE Computer Society Press, Los Alamitos, CA, 2004.

- [68] D. Rus, Z. Butler, K. Kotay, and M. Vona. Self-reconfiguring robots. *Communications of the ACM*, 45(3):39–45, 2002.
- [69] H.-P. Schwefel. *Evolutionstrategie und numerische Optimierung (in German)*. PhD thesis, Fachbereich Verfahrenstechnik, Technische Universität Berlin, Berlin, Germany, 1975.
- [70] A. B. Sendova-Franks and N. R. Franks. Self-assembly, self-organization and division of labour. *Philosophical Transactions of the Royal Society B—Biological Sciences*, 354(1388):1395–1405, 1999.
- [71] L. Spector, J. Klein, C. Perry, and M. Feinstein. Emergence of collective behavior in evolving populations of flying agents. *Genetic Programming and Evolvable Machines*, 6(1):111–125, 2005.
- [72] J. H. Sudd. The transport of prey by an ant, *Pheidole crassinoda* Em. *Behaviour*, 16:295–308, 1960.
- [73] T. G. Sugar and V. Kumar. Control of cooperating mobile manipulators. *IEEE Transactions on Robotics and Automation*, 18(1):94–103, 2002.
- [74] H. Sugie, Y. Inagaki, S. Ono, H. Aisu, and T. Unemi. Placing objects with multiple mobile robots—mutual help using intention inference. In *Proc. of the 1995 IEEE Int. Conf. on Robotics and Automation*, volume 2, pages 2181–2186. IEEE Computer Society Press, Los Alamitos, CA, 1995.
- [75] D. J. T. Sumpter. The principles of collective animal behaviour. *Philosophical Transactions of the Royal Society B—Biological Sciences*, 361(1465):5–22, 2006.
- [76] G. Theraulaz and E. Bonabeau. A brief history of stigmergy. *Artificial Life*, 5(2):97–116, 1999.
- [77] G. Theraulaz, E. Bonabeau, C. Sauwens, J.-L. Deneubourg, A. Lioni, F. Libert, L. Passera, and R. Solé. Model of droplet dynamics in the Argentine ant *Linepithema humile* (Mayr). *Bulletin of Mathematical Biology*, 63(6):1079–1093, 2001.
- [78] V. Trianni, E. Tuci, and M. Dorigo. Evolving functional self-assembling in a swarm of autonomous robots. In *Proc. of the 8<sup>th</sup> Int. Conf. on Simulation of Adaptive Behavior*, pages 405–414. MIT Press, Cambridge, MA, 2004.
- [79] E. Tuci, R. Groß, V. Trianni, M. Bonani, F. Mondada, and M. Dorigo. Cooperation through self-assembly in multi-robot systems. *ACM Transactions on Autonomous and Adaptive Systems*, 1(2):115–150, 2006.
- [80] G. Vakanas and B. Krafft. Coordination of behavioral sequences between individuals during prey capture in a social spider, *Anelosimus eximius*. *Journal of Insect Behavior*, 14(6):777–798, 2001.
- [81] G. Vakanas and B. Krafft. Regulation of the number of spiders participating in collective prey transport in the social spider *Anelosimus eximius* (Araneae, Theridiidae). *Comptes Rendus Biologies*, 327(8):763–772, 2004.
- [82] Z. D. Wang, E. Nakano, and T. Takahashi. Solving function distribution and behavior design problem for cooperative object handling by multiple mobile robots. *IEEE Transactions on Systems, Man, and Cybernetics—Part A: Systems and Humans*, 33(5):537–549, 2003.
- [83] C. R. Ward, F. Gobet, and G. Kendall. Evolving collective behavior in an artificial ecology. *Artificial Life*, 7(2):191–209, 2001.
- [84] S. A. West, A. Gardner, and A. S. Griffin. Altruism. *Current Biology*, 16(13):R482–R483, 2006.
- [85] G. M. Whitesides and B. Grzybowski. Self-assembly at all scales. *Science*, 295(5564):2418–2421, 2002.
- [86] W. Wickler and U. Seibt. Pedogenetic sociogenesis via the ‘sibling-route’ and some consequences for *Stegodyphus* spiders. *Ethology*, 95:1–18, 1993.
- [87] S. Yamada and J. Saito. Adaptive action selection without explicit communication for multirobot box-pushing. *IEEE Transactions on Systems, Man, and Cybernetics—Part C: Applications and Reviews*, 31(3):398–404, 2001.

- [88] X. Yao, Y. Liu, and G. Lin. Evolutionary programming made faster. *IEEE Transactions on Evolutionary Computation*, 3(2):82–102, 1999.
- [89] M. Yim, Y. Zhang, and D. Duff. Modular robots. *IEEE Spectrum*, 39(2):30–34, 2002.
- [90] M. Yim, Y. Zhang, K. Roufas, D. Duff, and C. Eldershaw. Connecting and disconnecting for chain self-reconfiguration with PolyBot. *IEEE/ASME Transactions on Mechatronics*, 7(4):442–451, 2002.
- [91] Y. Zhang, E. K. Antonsson, and A. Martinoli. Evolving neural controllers for collective robotic inspection. In *Proc. of the 9<sup>th</sup> Online World Conf. on Soft Computing in Industrial Applications (WSC9)*, volume 34 of *Advances in Soft Computing*, pages 721–733. Springer Verlag, Berlin, Germany, 2006.
- [92] V. Zykov, E. Mytilinaios, B. Adams, and H. Lipson. Self-reproducing machines. *Nature*, 435(7039):163–164, 2005.