
Towards group transport by swarms of robots

Roderich Groß*

Laboratoire de Systèmes de Robotiques,
Ecole Polytechnique Fédérale de Lausanne,
EPFL – STI – IMT – LSRO,
Station 9, 1015 Lausanne, Switzerland
E-mail: roderich.gross@ieee.org
*Corresponding author

Marco Dorigo

IRIDIA, CoDE,
Université Libre de Bruxelles,
Ave. F. Roosevelt 50,
CP 194/6, 1050 Brussels, Belgium
E-mail: mdorigo@ulb.ac.be

Abstract: We examine the ability of a swarm robotic system to transport cooperatively objects of different shapes and sizes. We simulate a group of autonomous mobile robots that can physically connect to each other and to the transported object. Controllers – artificial neural networks – are synthesised by an evolutionary algorithm. They are trained to let the robots self-assemble, that is, organise into collective physical structures and transport the object towards a target location. We quantify the performance and the behaviour of the group. We show that the group can cope fairly well with objects of different geometries as well as with sudden changes in the target location. Moreover, we show that larger groups, which are made of up to 16 robots, make possible the transport of heavier objects. Finally, we discuss the limitations of the system in terms of task complexity, scalability and fault tolerance and point out potential directions for future research.

Keywords: artificial neural networks; bio-inspired computation; collective structures; cooperation; cooperative transport; evolutionary robotics; foraging; group transport; modular robots; object manipulation; reconfigurable robots; self-assembling robots; self-assembly; swarm intelligence; swarm robotics.

Reference to this paper should be made as follows: Groß, R. and Dorigo, M. (2009) 'Towards group transport by swarms of robots', *Int. J. Bio-Inspired Computation*, Vol. 1, Nos. 1/2, pp.1–13.

Biographical notes: Roderich Groß received his Diploma in Computer Science from Universität Dortmund, Dortmund, Germany in 2001, the Diplôme d'Etudes Approfondies in Applied Science from Université Libre de Bruxelles, Brussels, Belgium in 2003 and his Doctorate in Engineering Sciences from Université Libre de Bruxelles, Brussels, Belgium in 2007. He was a JSPS Research Fellow at the Tokyo Institute of Technology, Tokyo, Japan in 2005, a Research Assistant at the University of Bristol, Bristol, UK in 2006 and a Marie Curie Fellow at Unilever R&D Port Sunlight, Bebington, UK in 2007. His research interests include artificial life, computational biology, robotics and swarm intelligence.

Marco Dorigo received his Doctorate in Information and Systems Electronic Engineering in 1992 from Politecnico di Milano, Milan, Italy. Since 1996, he has been a tenured Researcher of the FNRS, the Belgian National Fund for Scientific Research and a Research Director of IRIDIA, the artificial intelligence laboratory of the Université Libre de Bruxelles. He is the inventor of the ant colony optimisation metaheuristic. He is the Editor-in-Chief of *Swarm Intelligence*. In 2003, he was awarded the 'Marie Curie Excellence Award', in 2005, the 'Dr A. De Leeuw-Damry-Boullart award in applied sciences' and in 2007, the 'Cajastur International Prize for Soft Computing'.

1 Introduction

Over the last two decades, group transport has become a canonical task for studying cooperation in multi-robot systems. Instead of relying on a single robot that is powerful

enough to transport a given object, group transport relies on multiple, typically less powerful robots. If these robots act in concert, they can, as a group, exert larger forces than any of the individual robots alone.

The most common approach to group transport is to employ a decentralised control system (e.g., see Donald et al., 1997; Pereira et al., 2004). In many decentralised systems, the robots take their actions on the basis of local information – a common exception for this is when some or all robots possess information about the target location. Some of these systems have a leader-follower architecture (Stilwell and Bay, 1993; Kosuge and Oosumi, 1996; Kosuge et al., 1998; Wang et al., 2004). In most other systems, the robots are homogeneous in control (Kube and Zhang, 1993, 1997; Aiyama et al., 1999; Yamada and Saito, 2001). Studies of such homogeneous systems are of particular interest to researchers in the field of swarm robotics (Dorigo and Şahin, 2004), especially when the coordination of large numbers of robots is addressed.

In Groß and Dorigo (2004a, 2008a), we investigated self-assembly (Whitesides and Grzybowski, 2002; Groß and Dorigo, 2008c) as a potential coordination mechanism in group transport. We considered a system of simulated robots that can physically connect to each other and to the transported object. We used an evolutionary algorithm to design artificial neural network controllers. The controllers were trained to let two simulated robots transport a cylindrical object. Two distinct strategies were evolved:

- 1 Strategy 1 lets the robots push directly the object (without self-assembling)
- 2 Strategy 2 lets the robots self-assemble and push the object, both directly and indirectly.

We then increased the number of robots from two to five and at the same time we increased the mass of the object by the same factor (but kept constant the object’s dimensions). In this new set-up, Strategy 1 (i.e., pushing directly the object) was incapable of moving the object. Strategy 2 (i.e., self-assembling and pushing the object) could move the object, but the performance was poor. When we also increased the diameter of the object, Strategy 1 once again reached a satisfactory level of performance. Strategy 2 also performed better in this set-up, but still the performance was poor.

Kube and Bonabeau (2000) made similar observations. They evaluated how sensitive a system of physical robots was in regard to the object’s geometry. In an experiment where the robots had to move a small cuboid (with a side length that was approximately twice the robot’s diameter), the task took longer to complete when more robots were involved (groups with two to six robots were considered). This loss of efficiency was attributed to the increased competition for the limited surface of the object when robots were part of larger groups. Computer simulations confirmed this effect. Note that in the most successful set-ups that were investigated by Kube and Bonabeau (2000) and by Groß and Dorigo (2008a), the object was cylindrical and there was similarly low competition for the surface of the object.

A second problem that is typically related to the object’s geometry is that the object, while being manipulated, might rotate in an undesired manner (Mason, 1986). As noted by

Matarić et al. (1995), for instance, the ‘performance of pushing an elongated box toward a goal region is difficult for a single robot and improves significantly when performed cooperatively, but requires careful coordination between the robots’.

At present, multi-robot systems are fairly sensitive to the shape and size of the object to be moved. Moreover, the performance of the systems does not scale well with the number of robots. In our view, this lack of flexibility and scalability limits the practical value of current group transport systems.

In this paper, we investigate whether robots that are capable of self-assembling may overcome these limitations. We aim at controlling a swarm of robots to transport objects of different shapes and sizes. We make use of an evolutionary algorithm to design artificial neural network controllers. The objects to be transported are of cuboid as well as of cylindrical shape, with footprints that differ in size by factors up to 6.25. In addition, the objects can be of two different heights (and potentially cause visual occlusions). The objects’ masses can vary up to a factor of two and are independent of their geometries.

The paper is organised as follows. Section 2 details the used methods. Section 3 presents a series of results obtained in computer simulations. In particular, we analyse the performance and the behaviour of the system and quantify its sensitivity to

- 1 the object’s geometry
- 2 changes in the target location
- 3 the problem size, that is, the mass of the object and the number of robots.

Section 4 discusses the results and concludes the paper.

2 Methods

In this section, we detail the simulation model, the controller, and the evolutionary algorithm.

2.1 Simulation model


The simulation environment that we built models the kinematics and dynamics of rigid, partially constrained, bodies in 3-D using the Vortex™ simulation toolkit (CMLabs Simulations, Inc., Canada). Frictional forces are calculated based on the Coulomb friction law (Coulomb, 2001). Moments of inertia of composite objects are calculated using the parallel axis theorem (Weisstein, 2008).

The object to be transported, hereafter called ‘prey’, is modelled either as a cylinder or as a cuboid. If the prey is modelled as a cylinder, its radius (in cm) is chosen in the interval [6, 15]. If the prey is modelled as a cuboid, the length l (in cm) of one horizontal side is chosen in the interval [12, 40] and the length of the other horizontal side is chosen in [12, $\min(l, 20)$]. The height of the prey is either 12 cm or 20 cm (the robot’s height is 16.85 cm). Prey of height 12 cm have a uniform distribution of mass. Prey of

height 20 cm are modelled as a stack of two geometrically identical objects of different masses: the mass of the object on top of the stack is 25% of the total mass. This prevents tall prey from toppling over when pushed or pulled by the robots. Tall prey are of special interest as they can prevent robots from perceiving the target location. The mass of the prey (in grams) is chosen in $\{500, 625, 750, 875, 1,000\}$. In every simulation trial, the geometry and mass of the prey are chosen independently and at random (using uniform distributions). Depending on its mass, the prey requires the cooperative behaviour of either two or three robots to be moved.

The simulation model of the robot is illustrated in Figure 1. It is an approximate model of the ‘s-bot’, the robot of the swarm-bot system (Dorigo et al., 2004, 2005; Mondada et al., 2005b). The robot is composed of an upper part (the turret) and a lower part (the chassis). The turret can rotate with respect to the chassis by means of a motorised axis. The turret 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 chassis is composed of a cylindrical body and of four wheels: two wheels are cylindrical and linked to the chassis via hinge joints; two wheels are spherical and linked to the chassis via ball-and-socket joints. The turret weighs 280 g, the cylindrical body of the chassis weighs 300 g and each of the four wheels weighs 20 g. The complete simulation model has a mass of 660 g and a height of 16.85 cm.

Figure 1 The simulation model of the robot: front, side and top view



Notes: Units are in cm.

The robot’s abilities are summarised in Table 1. The cylindrical wheels are motorised and can be moved with angular speeds w_l and w_r (in rad/s) in the range $[-M, M]$, where $M = 8$. The cuboid heading forward represents a connection device (i.e., a gripper). If it is in contact with either the turret of another robot or with the prey, a physical connection can be established. Connections can be released at any time. A rotating base actuator enables the robot to align its turret (with respect to its chassis) to angular offsets $\bar{\gamma}$ (in rad) in the range $[0, 2\pi]$; at the start of a simulation trial, the angular offset is π (this situation is depicted in Figure 1). Note that the physical response is delayed; the angular speed (in rad/s) of the rotation is two. The pillar on top of the robot represents a camera system, which, in principle, provides an omni-directional view of the scene.

The robot can detect the angular position (α) of the target location (a light source), unless the latter is not visible (it can be occluded by either a tall prey or another robot). The robot scans for the prey and for other robots on a virtual ray heading in a controllable, horizontal direction ($\bar{\beta}$). The scanning range is limited to $R = 50$ cm. Two scans are performed. They stop at the first (i.e., the closest) intersection point (if any) with the prey and with the turret of another robot, respectively. The corresponding distances (d and e) can be calculated. In our system, the prey cannot be occluded by other robots. However, a robot can be occluded by all types of prey; note that a robot can only be perceived if its turret is visible. A connection sensor enables a robot to perceive whether it is connected to another object ($c = 1$) or not ($c = 0$). Moreover, the robot can perceive the current angular position of the rotating base (γ) and of the camera heading (β).

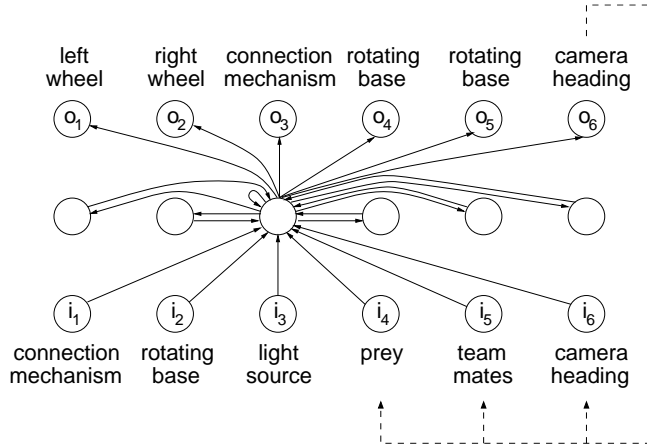
Table 1 Summary of the robot’s abilities

<i>Actuators</i>	
Left wheel (angular speed; rad/s)	$w_l \in [-M, M]$
Right wheel (angular speed; rad/s)	$w_r \in [-M, M]$
Connection mechanism	$\bar{c} \in \{0, 1\}$
Rotating base (angular position; rad)	$\bar{\gamma} \in [0, 2\pi]$
Camera heading (angular position; rad)	$\bar{\beta} \in [0, 2\pi]$
<i>Sensors (exteroceptive)</i>	
Light source (angular position; rad)	$\alpha \in [0, 2\pi]$
Prey (distance; cm)	$d \in [0, R]$
Team mates (distance; cm)	$e \in [0, R]$
<i>Sensors (proprioceptive)</i>	
Connection mechanism	$c \in \{0, 1\}$
Rotating base (angular position; rad)	$\gamma \in [0, 2\pi]$
Camera heading (angular position; rad)	$\beta \in [0, 2\pi]$

Random noise affects the robot’s actuators (i.e., w_b , w_r , M , $\bar{\gamma}$ and $\bar{\beta}$) and sensors (i.e., α , d , e , R , γ and β); for details, see (Groß and Dorigo, 2008a).

2.2 Controller

All the robots of a group are initialised with an identical controller. Every 100 ms, a control loop executes a simple recurrent neural network (Elman, 1990) taking input from the robot’s sensors and uses the outputs as motor commands. The neural network is illustrated in Figure 2. It has an input layer of six neurons (i_1 , i_2 , ..., and i_6), a hidden layer of six (fully inter-connected) neurons and an output layer of six neurons (o_1 , o_2 , ..., and o_6).

Figure 2 The neural network controller

Notes: Only connections to and from the third neuron of the hidden layer (from the left) are illustrated. Dashed arcs illustrate the feedback loop corresponding to the camera heading (i.e., the controllable, horizontal direction in which the surrounding is scanned).

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

$$i_1 = c, \quad (1)$$

$$i_2 = \frac{\gamma}{2\pi}, \quad (2)$$

$$i_3 = \begin{cases} 0 & \text{if light source not visible;} \\ 0.5 + 0.5 \frac{\alpha}{2\pi} & \text{otherwise,} \end{cases} \quad (3)$$

$$i_4 = \begin{cases} 0 & \text{if prey not visible;} \\ 1.0 - 0.9 \frac{d}{R} & \text{otherwise,} \end{cases} \quad (4)$$

$$i_5 = \begin{cases} 0 & \text{if team mates not visible;} \\ 1.0 - 0.9 \frac{e}{R} & \text{otherwise,} \end{cases} \quad (5)$$

$$i_6 = \frac{\beta}{2\pi}. \quad (6)$$

The activations of the six output neurons are used to set the motor commands (see Table 1).

$$w_l = M(2o_1 - 1), \quad (7)$$

$$w_r = M(2o_2 - 1), \quad (8)$$

$$\bar{c} = \begin{cases} 0 & \text{if } o_3 < 0.5; \\ 1 & \text{otherwise,} \end{cases} \quad (9)$$

$$\bar{\gamma} = \pi + (o_4 - o_5)\pi, \quad (10)$$

$$\bar{\beta} = \left(\beta + \frac{\pi}{5}(o_6 - 0.5) \right) \pmod{2\pi}. \quad (11)$$

The robot determines the initial camera heading ($\bar{\beta}$) by scanning once its entire surrounding and then choosing the direction for which the minimum distance to the prey was observed.

2.3 Evolutionary algorithm

The used evolutionary algorithm is a self-adaptive version of a $(\mu + \lambda)$ evolution strategy (Schwefel, 1975; Beyer, 2001). In the following, we refer to the genotype by 'individual'. Each individual is composed of $n = 120$ 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, a fitness value is assigned to each individual. The best-rated μ individuals are selected to create λ offspring. Subsequently, the μ parent individuals and the λ offspring are copied into the population of the next generation. Note that the μ parent individuals that are replicated from the previous generation get re-evaluated. We have chosen $\mu = 20$ and $\lambda = 60$.

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 (Beyer, 2001), both with the same probability. The offspring is subject to mutation. The object parameter x_i is mutated by adding a random variable from the normal distribution $N(0, s_i^2)$. 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 [for details, see (Groß and Dorigo, 2008a)].

2.3.1 Fitness computation

The simulated environment consists of a flat ground, a prey and a light source. The light source represents the target location. A simulation trial lasts $T = 20$ simulated seconds. Initially, $N = 4$ robots are put at random positions and with random orientations in the neighbourhood of the prey. The placement strategy is illustrated in Figure 3. Recall that depending on its mass, the prey cannot be moved by less than two or three robots, respectively.

The performance of an individual is evaluated in $S = 5$ independent trials. Every individual within the same generation is evaluated on the same sample of start configurations. At each generation a new sample is generated. Let \mathcal{P}_i be the performance observed in trial i . Then, the fitness is given by

$$\mathcal{F} = \frac{1}{S} \sum_{i=1}^S \mathcal{P}_i. \quad (12)$$

The fitness values of the individuals have to be maximised.

The performance measure \mathcal{P} is defined as:

$$\mathcal{P} = \begin{cases} \mathcal{A} & \text{if } \mathcal{T} \leq 0; \\ 1 + (1 + \mathcal{T}^{\rho_1}) \mathcal{A}^{\rho_2} & \text{otherwise,} \end{cases} \quad (13)$$

where $\mathcal{A} \in [0, 1]$ reflects the assembly performance, $\mathcal{T} \in [0, \infty)$ reflects the transport performance and $\rho_1 = 0.5$ as well as $\rho_2 = 2$ are parameters that were determined by trial and error.

Let \mathcal{Z} be the set of control steps after $\frac{T}{2} = 10$ s have elapsed – note that as the robots start from separate positions and up to $R = 50$ cm away from the prey, some time is required to approach the prey and to establish a connection. Then, the assembly performance \mathcal{A} can be defined as:

$$\mathcal{A} = \frac{1}{|\mathcal{Z}|N} \sum_{t \in \mathcal{Z}} \sum_{i=1}^N \mathcal{A}_i^t, \quad (14)$$

where $\mathcal{A}_i^t \in [0, 1]$ is defined by

$$\mathcal{A}_i^t = \begin{cases} 1 & \text{if } i \in \mathcal{M}^t; \\ 0 & \text{if } (i \notin \mathcal{M}^t) \wedge (\tilde{d}_i^t > R); \\ 0.75 & \text{if } (i \notin \mathcal{M}^t) \wedge (\tilde{d}_i^t < \frac{R}{2}); \\ 0.65 \frac{R - \tilde{d}_i^t}{R/2} + \frac{1}{10} & \text{otherwise.} \end{cases} \quad (15)$$


\mathcal{M}^t is the set of robots that are physically linked with the prey at time t . It comprises robots both directly and indirectly connected to the prey. \tilde{d}_i^t is an estimate of the minimum distance between robot i and the prey at time t . If no prey is detected within the sensing range, we set $\tilde{d}_i^t = R + 1$.

The transport performance measure \mathcal{T} reflects the distance (in cm) the prey has been moved towards the light source. It is defined as:

$$\mathcal{T} = \max(0, \mathcal{D}^0 - \mathcal{D}^T), \quad (16)$$

where $\mathcal{D}^t, t \in \{0, T\}$ denotes the distance (in cm) between the prey and the light source at time t .

Figure 3 Example of initial placement

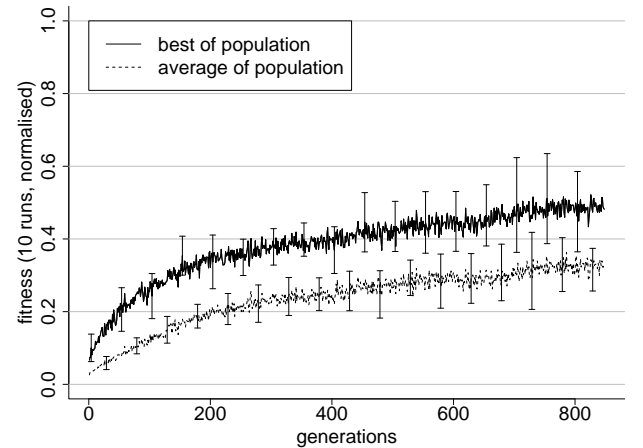


Notes: The prey (black disk) has to be transported towards the light source, which represents a target location (e.g., a nest). The robots approach the prey from the half space on the side of the light source: they are placed randomly within a semi-circle 40 cm to 50 cm away from the prey.

3 Results

We have performed ten independent evolutionary runs of 850 generations each. This corresponds to 68,000 fitness evaluations per run. Each run lasted about three weeks on a machine equipped with an AMD Athlon XP™ 2400+ processor. In Figure 4, the average and the maximum fitness over time are presented.

Figure 4 Evolution of group transport behaviours for groups of four robots: development of fitness in ten evolutionary runs



Note: Error bars are shown.

Each curve corresponds to the average of ten runs with different random seeds. The fitness values are normalised in the range $[0, 1]$; bounds for the performance were computed as described in Groß and Dorigo (2008a). Altogether, the best and average fitness values continuously increase for the entire range of 850 generations. Note, however, that the attained fitness level varies among the different runs.

3.1 Performance analysis

The fitness assigned to an individual (i.e., to a genotype) is the outcome of a stochastic process that is influenced by external factors, including the robots' initial positions and orientations, the geometry of the prey and the noise affecting the robots' sensors and actuators. 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 this 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 2,500 start configurations.

In the following, we analyse the assembly performance and the transport performance of the best individuals. In this analysis, we do not use the performance metrics that were employed during the fitness computation [see equations (14) and (16)]. These metrics are relatively complex because they were partially designed to help bootstrapping the evolutionary process. Instead, we look at two related performance metrics:

- $|\mathcal{M}^T|$, that is, the number of robots that are physically linked with the prey at the end of the trial (these robots are either directly or indirectly connected to the prey)
- $D^0 - D^T$, that is, the distance (in cm) by which the prey approached the target location during the trial.

Table 2 Number of robots connected to the prey at the end of a trial for a group of four robots controlled by the best individual of each evolutionary run (average over 2,500 observations; runs presented, from the top to the bottom, with an increasing average performance)

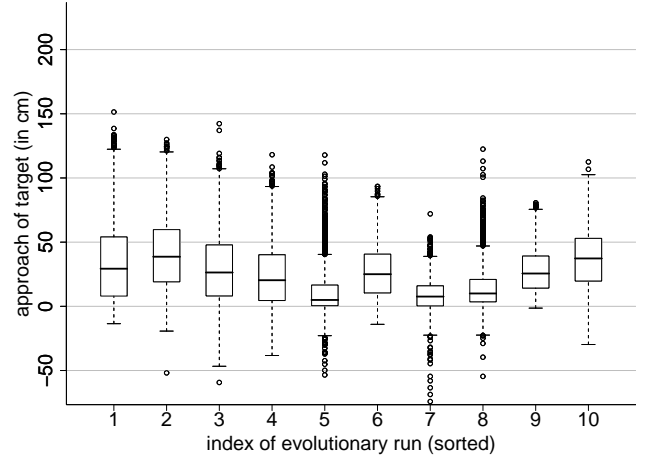
<i>Index of run</i>	<i>Average of \mathcal{M}^T</i>
1	3.05
2	3.13
3	3.15
4	3.19
5	3.20
6	3.21
7	3.28
8	3.44
9	3.45
10	3.57

Notes: The number comprises robots both directly and indirectly connected to the prey.

The assembly performance of the best individuals is shown in Table 2. The evolutionary runs are presented, from the top to the bottom, with an increasing average performance of their best individuals. In each run, on average, more than 3.0 (out of the four) robots connected to the prey. The highest assembly performance is 3.57 (run 10).

The transport performance of the best individuals is shown in Figure 5. The evolutionary runs are presented in the same order as in Table 2. The two best individuals belong to runs two and ten: the average distances (in cm) are 40.5 and 37.0, respectively. The corresponding standard deviations are 27.6 and 22.0.

Figure 5 Box-and-whisker plot visualising the distance (in cm) prey (of different shapes and sizes) were moved by a group of four robots controlled by the best individuals of each run (2,500 observations per box)




Notes: Each box comprises observations ranging from the first to the third quartile. The median is indicated by a bar. The whiskers extend to the farthest data points that are within 1.5 times the interquartile range. Outliers are indicated as circles.

In the following, we focus on the best individual of run ten. This individual exhibited a relatively high transport performance and at the same time relatively few fluctuations in performance. Moreover, it achieved the best assembly performance. For simplicity, we refer to this individual as to the best individual.

Figure 6 plots the distance (in cm) by which prey of different mass approached the target location, as observed in the 2,500 trials for the best individual. Upper bounds for the performance are indicated by the bold line. The average distances (in cm) are 53.8, 44.5, 37.0, 29.0 and 20.9. This is respectively 24.1%, 21.9%, 20.2%, 17.8% and 14.6% of the upper bounds for prey of mass (in grams) 500, 625, 750, 875 and 1,000. Note that the upper bounds are not tight; they correspond to situations in which all four robots start from an optimal configuration in which they are already pre-assembled with each other and the prey; for details, see (Groß and Dorigo, 2008a). The standard deviations are astonishingly low; they range from 14.2 to 21.6. Recall that the fluctuations in performance are partially caused by the nature of the task (e.g., differences in the initial placement of the robots).

Figure 6 Box-and-whisker plot visualising the distance (in cm) prey of different mass were moved by a group of four robots controlled by the best individual (500 observations per box)



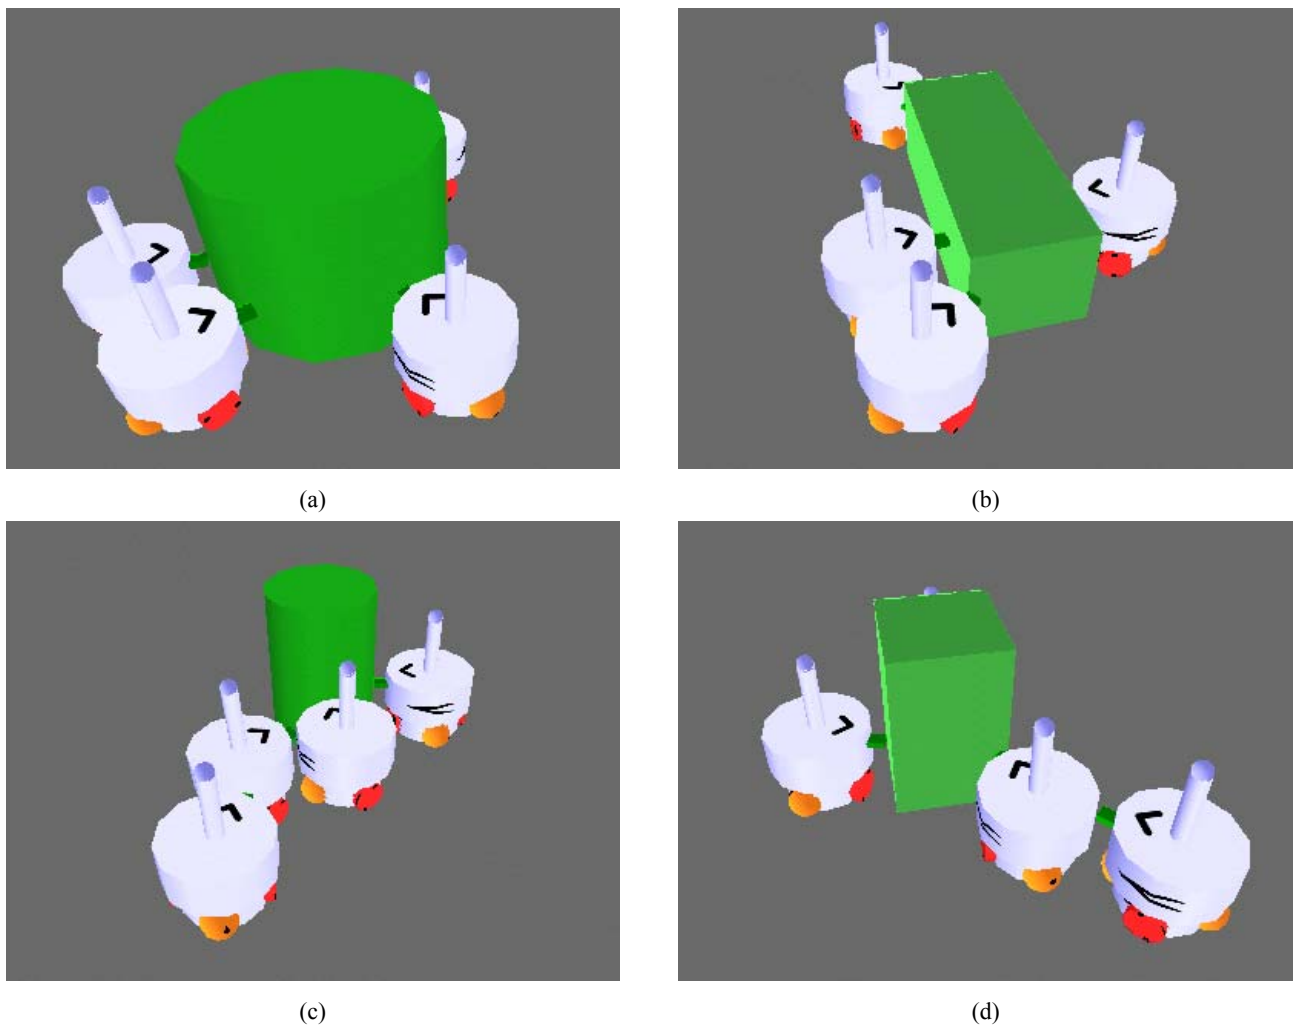
Note: The bold line indicates upper bounds for the performance

3.2 Behavioural analysis


Figure 7 shows several example structures that were formed during the transport of prey of different shapes and sizes. In Figures 7(a) and 7(d), the target location is occluded by the prey for one of the robots. Figure 7(c) shows a chain of three assembled robots. In three cases, see Figures 7(d)–7(f), chains of two assembled robots were formed.

Figure 8 illustrates the 17 topologies in which up to four robots and a prey can be organised. All topologies occurred at least once during the post-evaluation of the best individual (the frequencies are indicated in the figure). Over the 2,500 trials, $4 \times 2,500 = 10,000$ times a robot was controlled to transport the prey. In 89.3% of the cases, the robot was connected either directly or indirectly to the prey at the end of the trial. In 44.9% of the cases, the robot was part of an assembled structure, in other words, a structure of two or more robots that were directly connected to each other. In 70.4% of the 2,500 trials, the robots formed at least one such assembled structure (see cases 4, 6–8 and 10–17 in Figure 8).


Figure 7 Group transport of prey of different sizes and shapes towards the target location (towards the bottom, outside the range of these images) (see online version for colours)



Note: All four robots are controlled by identical recurrent neural networks.

Figure 7 Group transport of prey of different sizes and shapes towards the target location (towards the bottom, outside the range of these images) (continued) (see online version for colours)

(e)



(f)

Note: All four robots are controlled by identical recurrent neural networks.

Figure 8 Topology into which up to four robots (transparent disks) and a prey (gray disk) can be assembled and frequencies by which these topologies were observed at the end of a trial (2,500 observations in total)

topology	frequency		
			9 14.12% (353)
	1 1.12% (28)		10 37.96% (949)
	2 3.64% (91)		11 6.56% (164)
	3 10.56% (264)		12 6.52% (163)
	4 2.60% (65)		13 0.08% (2)
	5 0.20% (5)		14 12.52% (313)
	6 0.72% (18)		15 1.32% (33)
	7 0.08% (2)		16 0.20% (5)
	8 0.04% (1)		17 1.76% (44)

So far, we have considered only structures that are connected to the prey (see Figure 8). However, robots could also form structures that are not connected to the prey. In only 26 out of 2,500 trials, one out of the four robots remained separate from the prey (see cases 5–8 in Figure 8). It happened much more frequently that none of the four robots remained separate (2,026 trials), or that two out of the four robots remained separate from the prey (329 trials). This suggests that the connection attempts are not independent of each other. We speculate that in most trials where multiple robots remained separate from the

prey, these robots were connected to each other, and thereby did not succeed in connecting also to the prey.

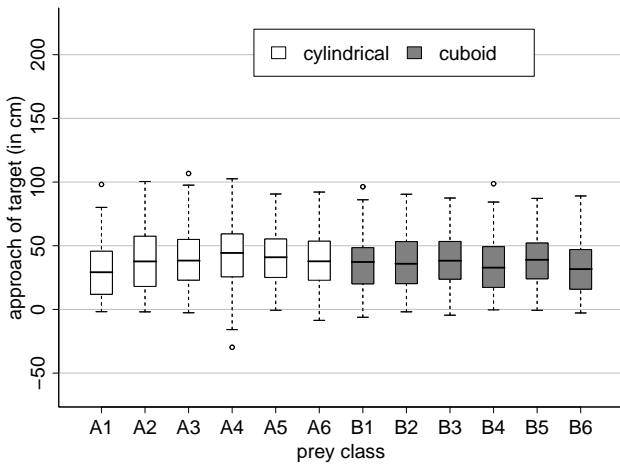
3.3 Different object geometries

To assess the flexibility of the best individual with respect to the shape and the size of the prey, data of 2,000 trials is partitioned according to different classes of prey (all units in cm):

- Prey classes A1, A2, ..., A6 comprise the data concerning cylindrical prey. A_i corresponds to trials in which the cylinders' radii reside in the range $[6 + 1.5(i - 1), 6 + 1.5i)$.
- Prey classes B1, B2, ..., B6 comprise the data concerning prey of cuboid shape. B1, B2, ..., B6 refer to trials in which the length of the longer side is in the ranges [12, 16), [16, 20), [20, 25), [25, 30), [30, 35) and [35, 40]. The length of the shorter side is always in the range [12, 20].

Figure 9 plots the distance (in cm) by which prey of different classes approached the target location (i.e., $D^0 - D^T$), as observed in the 2,000 trials for the best individual. Overall, the individual seems to perform fairly robustly with respect to all combinations of size and shape; the average performance for each prey class is within $\pm 17.7\%$ of the average performance across all classes of prey. However, we detected seven significant differences when testing each of the 66 pairs of classes (2-tailed Mann-Whitney tests, $p = 0.05$ before Bonferroni correction). Differences were observed for pairs A1 and A4, A1 and A5, A1 and A6, A1 and B3, A1 and B5, A4 and B4 and A4 and B6.

Figure 9 Box-and-whisker plot visualising the distance (in cm) prey of different shapes and sizes were moved by a group of four robots towards the light source (147 to 184 observations per box; mass of the prey (in grams): 500 to 1,000)



Concerning the transport of prey of different heights (12 cm and 20 cm, respectively), we could not detect a significant difference in the performances observed in 500 independent trials (2-tailed Mann-Whitney test, $p = 0.05$).

Note that the prey's geometry and mass are chosen independently. Moreover, in the set-up used during this experiment (see Figure 3), it might not occur often that the light source is occluded by a (tall) prey, as the robots are initially placed in between the light source and the prey.


3.4 Changes in target location

We examine to what extent a group of robots engaged in the transport of a prey is able to adapt dynamically the direction of transport according to a new target location. The simulation period of each trial is doubled from 20 to 40 seconds. In the first half of the simulation period, the experimental set-up has been retained unchanged. As soon as the first half of the simulation period has elapsed, the light source, which represents the target location, is moved instantaneously to a new location. The (horizontal) angular displacement (in rad) of the target with respect to the prey is chosen randomly in $\left\{ \frac{i}{8} 2\pi \mid i \in \{0, 1, 2, \dots, 7\} \right\}$.

For each angular displacement the distance by which the prey approached the target during the second half of the simulation period is recorded.

Figure 10 shows the performance of the best individual as observed in 4,000 trials. If the target direction does not change, the average distance (in cm) by which the light prey (heavy prey) approaches the target during the second half of the simulation period is 82.3 (36.2). Depending on the angular displacement of the target, the average distance (in cm) may decrease down to 45.7 (13.9).

Figure 10 Ability of a group of four robots to adapt the direction of transport according to a new target location



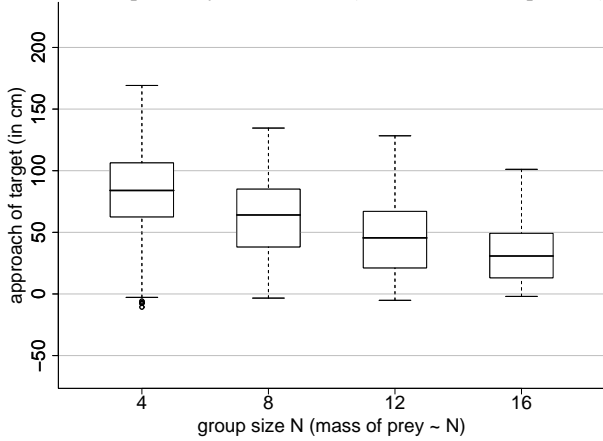
Notes: The box-and-whisker plot shows the distance (in cm) prey (of different shapes and sizes) were moved towards the new target. Eight different angular displacements of the target with respect to the prey are considered (250 observations per box).

3.5 Scalability

We consider groups of 4, 8, 12 and 16 robots. Along with the group size, the mass of the prey (in grams) is increased proportionally from 500 to 1,000, 1,500 and 2,000, respectively. We evaluate the best individual using these set-ups 1,000 times in total. As large groups of robots might require more time to self-assemble, the simulation period (T) is extended to 30 seconds. To ensure a non-overlapping placement of up to 16 robots, the latter are initially put at random positions within a full circle, 25 cm to 50 cm away from the prey (for a comparison, see Figure 3).

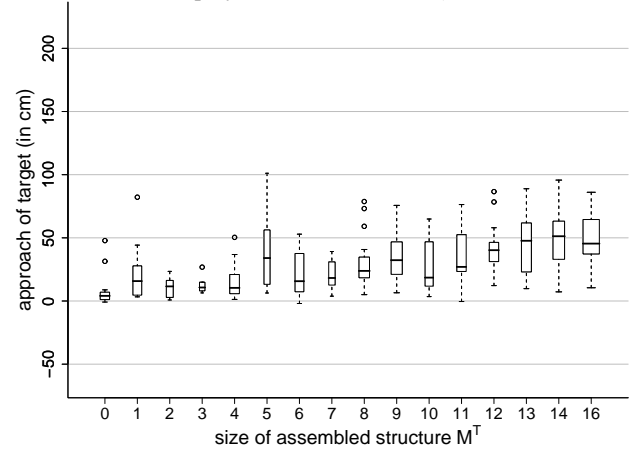
Figure 11 shows the distance (in cm) the prey was moved in each set-up. The average distances (in cm) are 83.2, 61.3, 46.7 and 33.6. This is respectively 25.0%, 18.4%, 14.0% and 10.1% of the upper bound. The standard deviations are in the range [23.6, 34.6]. Overall, the performance decreases with group size and the mass of the prey. We observed that the high density in which the 16 robots are initially put in the semi-circle makes it difficult for the robots to self-assemble into a common structure comprising the prey. However, in those cases in which the majority of robots assembled into a common structure (comprising the prey), the system exhibited its best performance (see Figure 12).

Figure 11 Performance of the best evolved individual for groups of 4, 8, 12 and 16 robots transporting prey of mass (in grams) 500, 1,000, 1,500 and 2,000, respectively, for $T = 30$ s (250 observations per box)



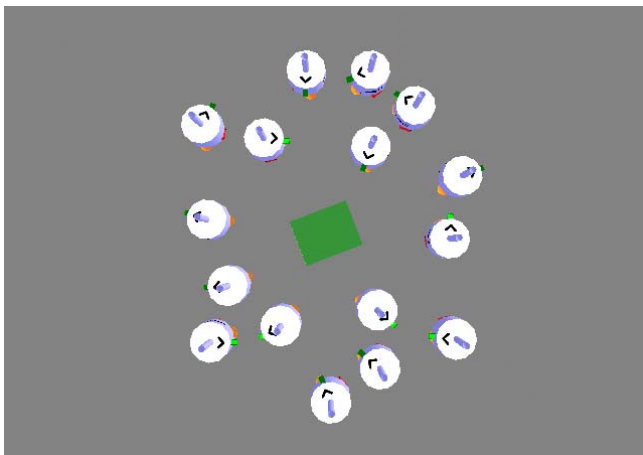
Notes: The robots were initially arranged at random within a full circle, 25 cm to 50 cm away from the prey.

Figure 12 Performance of the best evolved individual for groups of 16 robots transporting prey of different shapes and sizes and of mass 2,000 g for $T = 30$ s; 250 observations, grouped according to $|\mathcal{M}^T|$ (i.e., the number of robots physically linked to each other and the prey at the end of the trial)




Notes: Boxes are drawn with widths proportional to the square-roots of the number of observations in the groups.


Figure 13 Group transport of prey of mass 2,000 g by a swarm of 16 robots towards the target location (towards the bottom, outside the range of these images) (see online version for colours)




(a)



(b)




(c)



(d)

Notes: Robots are controlled by identical recurrent neural networks. (a)–(d): sequence of snapshots taken from a trial where no robot connected to the prey. All robots are connected into a single structure comprising a 10-robot loop, which surrounds the prey. Snapshots were taken at times 0 s (a), 2.5 s (b), 10 s (c) and 25 s (d).

Figure 13 Group transport of prey of mass 2,000 g by a swarm of 16 robots towards the target location (towards the bottom, outside the range of these images) (continued) (see online version for colours)



Notes: Robots are controlled by identical recurrent neural networks. (e)-(h): sequence of snapshots taken from a trial where all robots connected to the prey. Fourteen of the robots are connected into a single entity. Snapshots were taken at times 0 s (e), 2.5 s (f), 10 s (g) and 25 s (h).

Figure 13 shows sequences of snapshots that were taken during two trials. In the first trial (see Figures 13(a)–13(d)), no robot is connected to the prey. Nevertheless, all robots form a single structure comprising a 10-robot loop, which surrounds the prey. In the second trial (see Figures 13(e)–13(h)), all robots are connected, directly or indirectly to the prey.

4 Discussions

In this paper, we examined the ability of a swarm robotic system to transport cooperatively objects of different shapes and sizes. The study was conducted using a 3-D physics simulation engine. The robots were autonomous and mobile. Each robot could physically connect to other robots and to the object. Depending on its mass, the object could not be moved by less than two or three robots. Using an evolutionary algorithm, we synthesised artificial neural network controllers that let the robots self-assemble – that is, organise into collective physical structures – and transport the object towards a target location.

Different from our previous studies (Groß and Dorigo, 2004a, 2004b, 2008a), we provided the robot with advanced acting and cognitive abilities and at the same time allowed the neural network to exploit all aspects of these abilities. For example, the neural network could independently control the robot in terms of

- 1 where to move
- 2 where to look
- 3 where to link with other robots or the object.

This strategy is similar to the one employed in Trianni et al. (2004) and Tuci et al. (2008), where it was achieved a smooth transfer of the controllers evolved in simulation to the physical robots. The advantage of this strategy is that it places relatively few constraints on the evolved solutions, thereby making it possible for a variety of complex group behaviours to emerge. For example, in this paper, we reported on a trial where a group of 16 robots self-organised into a cyclic structure, which entirely surrounded the object (without being physically connected to it). This structure

was functional – it pulled the object like an artificial rope [for a collection of video recordings, see Groß and Dorigo (2008b)]. The drawback of this strategy is that the solution search space is relatively complex and large. In our case, the evolutionary process was very demanding in terms of computational resources.

In the following, we summarise our achievements and discuss some of the main challenges that remain in regard to object transportation and manipulation by swarms of self-assembling robots:

- 1 *Task complexity.* We considered objects of cylindrical and of cuboid shapes, with footprints that differed in size by factors up to 6.25. Overall, our system performed fairly robustly with respect to the object's geometry – the average performance for each of the twelve classes of prey was within $\pm 17.7\%$ of the average performance across all classes of prey. It could also cope well with sudden changes in the target location. However, the manipulation space was 2-D and the task required only a linear translation of the object in this space. Moreover, no obstacles were present in the environment. How to design a swarm system that can meet the requirements of 3-D manipulation tasks is an open problem, in particular, if the requirements are not known in advance. Such systems might benefit from a tight coupling between behaviour and morphology (Pfeifer and Bongard, 2006; O'Grady et al., 2007).
- 2 *Scalability.* It was shown that our system is applicable to the transport of heavier objects by larger groups (with up to 16 simulated robots). However, the performance did not scale well with the group size. The results indicate that the mass per capita that can be transported by our system decreases with group size. Results reported by Kube and Bonabeau (2000) indicate that even the total mass that can be transported by their system decreases with group size (groups of 3–6 robots were considered). Presumably, such limitations are present in all physical systems once a critical number of robots are reached. Mondada et al. (2005a) report about super-linear performances in pulling chains of robots of the swarm-bot system. In particular, the force per capita was found first to increase and then to decrease with the length of the chains (depending on the ground condition, the highest per capita performances were achieved by groups of two or three robots). Super-linear performances were also observed in ants (Hölldobler et al., 1978; Franks, 1986). Understanding and exploiting such super-linear performances and pushing the limits of scale to systems with hundreds of robots will be a subject of future research.
- 3 *Fault tolerance.* We exposed our system to various types of noise that affected both the robots' actuators and sensors. However, the robots of our computer simulations were not subject to individual failure. This is a serious limitation as

- 1 robots in swarm robotic systems are error-prone
- 2 a single robot, when part of a physically inter-connected structure, has a semi-global impact on this structure.

Therefore, strategies of fault detection and self-repair need to be investigated (Parker, 1999; Tomita et al., 1999; Bererton and Khosla, 2001; Christensen et al., 2008a, 2008b).

Acknowledgements

This work was supported by the Swarm-bots project, funded by the Future and Emerging Technologies Programme (IST-FET) of the European Commission, under contract IST-2000-31010. The information provided is the sole responsibility of the authors and does not reflect the European Commission's opinion. The European Commission is not responsible for any use that might be made of data appearing in this publication. Roderich Groß acknowledges support by the European Commission in the form of a Marie Curie Intra-European Fellowship, under contract MEIF-CT-2006-040312. Marco Dorigo acknowledges support from the FRS-FNRS, of which he is a Research Director and from the ANTS project, an Action de Recherche Concertée funded by the Scientific Research Directorate of the French Community of Belgium.

References

- Aiyama, Y., Hara, M., Yabuki, T., Ota, J. and Arai, T. (1999) 'Cooperative transportation by two four-legged robots with implicit communication', *Robot. Auton. Syst.*, Vol. 29, No. 1, pp.13–19.
- Bererton, C. and Khosla, P.K. (2001) 'Towards a team of robots with repair capabilities: a visual docking system', in *Proc. of the 7th Int. Symp. on Experimental Robotics*, Vol. 271 of Lect. Notes Contr. Inf., pp.333–342, Springer Verlag, Berlin, Germany.
- Beyer, H-G. (2001) *The Theory of Evolution Strategies*, Springer Verlag, Berlin, Germany.
- Christensen, A.L., O'Grady, R., Birattari, M. and Dorigo, M. (2008a) 'Fault detection in autonomous robots based on fault injection and learning', *Auton. Robot.*, Vol. 24, No. 1, pp.49–67.
- Christensen, A.L., O'Grady, R. and Dorigo, M. (2008b) *From Fireflies to Fault Tolerant Swarms of Robots*, Technical Report TR/IRIDIA/2008-002, IRIDIA, CoDE, Univ. Libre de Bruxelles, Brussels, Belgium.
- Coulomb, C.A. (2001) *Théorie des Machines Simples, en Ayant égard au Frottement de Leurs Parties et à la Roideur des Cordages*, in French, Reprint of an 1821 ed. by Bachelier, Paris, Adamant Media Corporation, Boston.
- Donald, B.R., Jennings, J. and Rus, D. (1997) 'Information invariants for distributed manipulation', *Int. J. Robot. Res.*, Vol. 16, No. 5, pp.673–702.
- Dorigo, M. and Şahin, E. (2004) 'Swarm robotics – special issue editorial', *Auton. Robot.*, Vol. 17, Nos. 2–3, pp.111–113.

- Dorigo, M., Trianni, V., Şahin, E., Groß, R., Labella, T.H., Baldassarre, G., Nolfi, S., Deneubourg, J-L., Mondada, F., Floreano, D. and Gambardella, L.M. (2004) 'Evolving self-organising behaviours for a swarm-bot', *Auton. Robot.*, Vol. 17, Nos. 2–3, pp.223–245.
- Dorigo, M., Tuci, E., Groß, R., Trianni, V., Labella, T., Nouyan, S., Ampatzis, C., Deneubourg, J-L., Baldassarre, G., Nolfi, S., Mondada, F., Floreano, D. and Gambardella, L. (2005) 'The swarm-bots project', in *Proc. of the 1st Int. Workshop on Swarm Robotics at SAB 2004*, Vol. 3342, pp.31–44, Lect. Notes Comput. Sc., Springer Verlag, Berlin, Germany.
- Elman, J.L. (1990) 'Finding structure in time', *Cognit. Sci.*, Vol. 14, No. 2, pp.179–211.
- Franks, N.R. (1986) 'Teams in social insects: group retrieval of prey by army ants (*Eciton burchelli*, Hymenoptera: Formicidae)', *Behav. Ecol. Sociobiol.*, Vol. 18, No. 6, pp.425–429.
- Groß, R. and Dorigo, M. (2004a) 'Evolving a cooperative transport behavior for two simple robots', *Artificial Evolution – 6th Int. Conf., Evolution Artificielle (EA 2003)*, Vol. 2936, pp.305–317, Lect. Notes Comput. Sc., Springer Verlag, Berlin, Germany.
- Groß, R. and Dorigo, M. (2004b) 'Group transport of an object to a target that only some group members may sense', in *Proc. of the 8th Int. Conf. on Parallel Problem Solving from Nature (PPSN VIII)*, Vol. 3242, pp.852–861, Lect. Notes Comput. Sc., Springer Verlag, Berlin, Germany.
- Groß, R. and Dorigo, M. (2008a) 'Evolution of solitary and group transport behaviors for autonomous robots capable of self-assembling', *Adapt. Behav.*, Vol. 16, No. 5, pp.285–305.
- Groß, R. and Dorigo, M. (2008b) 'Group transport of objects of different shapes and sizes', available at <http://iridia.ulb.ac.be/supp/IridiaSupp2008-021>.
- Groß, R. and Dorigo, M. (2008c) 'Self-assembly at the macroscopic scale', *Proceedings of the IEEE*, Vol. 96, No. 9, pp.1490–1508.
- Hölldobler, B., Stanton, R.C. and Markl, H. (1978) 'Recruitment and food-retrieving behavior in *Novomessor* (Formicidae, Hymenoptera)', *Behav. Ecol. Sociobiol.*, Vol. 4, No. 2, pp.163–181.
- Kosuge, K. and Oosumi, T. (1996) 'Decentralized control of multiple robots handling an object', in *Proc. of the 1996 IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, IEEE Computer Society Press, Los Alamitos, CA, Vol. 1, pp.318–323.
- Kosuge, K., Oosumi, T., Satou, M., Chiba, K. and Takeo, K. (1998) 'Transportation of a single object by two decentralized-controlled non-holonomic mobile robots', in *Proc. of the 1998 IEEE Int. Conf. on Robotics and Automation*, Vol. 4, pp.2989–2994, IEEE Computer Society Press, Los Alamitos, CA.
- Kube, C.R. and Bonabeau, E. (2000) 'Cooperative transport by ants and robots', *Robot. Auton. Syst.*, Vol. 30, Nos. 1–2, pp.85–101.
- Kube, C.R. and Zhang, H. (1993) 'Collective robotics: from social insects to robots', *Adapt. Behav.*, Vol. 2, No. 2, pp.189–218.
- Kube, C.R. and Zhang, H. (1997) 'Task modelling in collective robotics', *Auton. Robot.*, Vol. 4, No. 1, pp.53–72.
- Mason, M.T. (1986) 'Mechanics and planning of manipulator pushing operations', *Int. J. Robot. Res.*, Vol. 5, No. 3, pp.53–71.
- Matarić, M.J., Nilsson, M. and Simsarian, K.T. (1995) 'Cooperative multi-robot box-pushing', in *Proc. of the 1995 IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, Vol. 3, pp.556–561, IEEE Computer Society Press, Los Alamitos, CA.
- Mondada, F., Bonani, M., Guignard, A., Magnenat, S., Studer, C. and Floreano, D. (2005a) 'Super-linear physical performances in a swarm-bot', in *Proc. of the 8th European Conf. on Artificial Life*, Vol. 3630, pp.282–291, Lect. Notes Artif. Int., Springer Verlag, Berlin, Germany.
- Mondada, F., Gambardella, L.M., Floreano, D., Nolfi, S., Deneubourg, J-L. and Dorigo, M. (2005b) 'The cooperation of swarm-bots: physical interactions in collective robotics', *IEEE Robot. Autom. Mag.*, Vol. 12, No. 2, pp.21–28.
- O'Grady, R., Groß, R., Christensen, A.L., Mondada, F., Bonani, M. and Dorigo, M. (2007) 'Performance benefits of self-assembly in a swarm-bot', in *Proc. of the 2007 IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, pp.2381–2387, IEEE Computer Society Press, Los Alamitos, CA.
- Parker, L.E. (1999) 'Adaptive heterogeneous multi-robot teams', *Neurocomputing*, Vol. 28, Nos. 1–3, pp.75–92.
- Pereira, G.A.S., Campos, M.F.M. and Kumar, V. (2004) 'Decentralized algorithms for multi-robot manipulation via caging', *Int. J. Robot. Res.*, Vol. 23, Nos. 7–8, pp.783–795.
- Pfeifer, R. and Bongard, J.C. (2006) *How the Body Shapes the Way We Think*, MIT Press, Cambridge, MA.
- Schwefel, H-P. (1975) *Evolutionstrategie und numerische Optimierung (in German)*, PhD thesis, Fachbereich Verfahrenstechnik, Technische Universität Berlin, Berlin, Germany.
- Stilwell, D.J. and Bay, J.S. (1993) 'Toward the development of a material transport system using swarms of ant-like robots', in *Proc. of the 1993 IEEE Int. Conf. on Robotics and Automation*, Vol. 1, pp.766–771, IEEE Computer Society Press, Los Alamitos, CA.
- Tomita, K., Murata, S., Kurokawa, H., Yoshida, E. and Kokaji, S. (1999) 'Self-assembly and self-repair method for a distributed mechanical system', *IEEE Trans. Robot. Autom.*, Vol. 15, No. 6, pp.1035–1045.
- Trianni, V., Tuci, E. and Dorigo, M. (2004) 'Evolving functional self-assembling in a swarm of autonomous robots', in *Proc. of the 8th Int. Conf. on Simulation of Adaptive Behaviour*, pp.405–414, MIT Press, Cambridge, MA.
- Tuci, E., Ampatzis, C., Trianni, V., Christensen, A. L. and Dorigo, M. (2008) 'Self-assembly in physical autonomous robots: the evolutionary robotics approach', in *Proc. of the 11th Int. Conf. on the Simulation and Synthesis of Living Systems (Artificial Life XI)*, pp.616–623, MIT Press, Cambridge, MA.
- Wang, Z.D., Takano, Y., Hirata, Y. and Kosuge, K. (2004) 'A pushing leader based decentralized control method for cooperative object transportation', in *Proc. of the 2004 IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, Vol. 1, pp.1035–1040, IEEE Computer Society Press, Los Alamitos, CA.
- Weisstein, E.W. (2008) *Parallel Axis Theorem*, Eric Weisstein's World of Physics, retrieved on 7 September 2008, available at <http://scienceworld.wolfram.com/physics/ParallelAxisTheorem.html>.
- Whitesides, G.M. and Grzybowski, B. (2002) 'Self-assembly at all scales', *Science*, Vol. 295, No. 5564, pp.2418–2421.
- Yamada, S. and Saito, J. (2001) 'Adaptive action selection without explicit communication for multirobot box-pushing', *IEEE Trans. Syst., Man, Cybern. C*, Vol. 31, No. 3, pp.398–404.