Recipes for Collective Movement

S. Goss[1], J.L. Deneubourg[1], R. Beckers[1] and J.-L. Henrotte[2]

[1]Unit of Behavioural Ecology
Center for Non-Linear Phenomena and Complex Systems
sgoss@ulb.ac.be

[2]Pragma

CP 231, Université Libre de Bruxelles
1050 Bruxelles, Belgique

Abstract

This paper explores the use of decentralised algorithms to coordinate the movement of a group of robots, as simulated on a computer. Each robot obeys the same rule which essentially determines at each step how many degrees it turns as a function of the sum of the signals it receives from the other robots. Each robot emits an omni-directional signal and has four receptors (ahead, behind, left right). Its position and movement thus affects the position and movement of others and vice-versa. While no group behaviour is programmed explicitly, a number of collective movement patterns readily emerge from these interactions, including moving in a circle, a spiral, an advancing figure of eight, a line, a loose bunch and different types of fly-like milling. Finally, where the group moves in a heterogenous environment containing, for example, targets to be collected, these distributed algorithms can be used to elicit long-range group exploration in regions of low target density and tighter short-range group exploration in regions of higher target density, without requiring the robots to measure local target density or communicate the presence of targets.

Introduction.

Autonomous robots that are relatively cheap and simple are now being produced in a number of laboratories. This article proposes a mechanism to coordinate the movement of a group of such robots. As in previous work (Deneubourg and Goss 1989; Deneubourg et al 1990; Theraulaz et al 1990; Goss and Deneubourg 1992), we favour a decentralised approach, such that each robot makes its own decisions based on local sensory input. The collective movement pattern is not programmed explicitly but results from the interactions between the robots.

The logic we propose is not very different from that governing the movement of a flock of birds or a school of fish (Parr 1927; Breder 1959, 1976; Sakai 1973; Suzuki and Sakai 1973; Huth and Wissel 1989), being that each robot chooses its direction as a function of the position of other nearby robots (see also the pioneering work of Walter 1950, 1951; Jacobson 1958; Braitenberg 1984). We exclude the possibility that each member has a more widespread knowledge of the postions of the other members, or of the group's centre of gravity (e.g. Reynolds 1987). The feed-back in such systems makes the link between the individual rules and the overall pattern highly non-intuitive, and there was an element of serendipity in the creation of the algorithms. This does not, however, mean that the patterns produced are necessarily unstable.

Given the much-discussed reality gap between simulation and implementation, we have chosen to make our simulations as simple as reasonable, and then a bit more, and intend this paper as an illustration of a potential mechanism rather than a technical description.

The simulated robots

Each robot emits an omni-directional signal and has four receivers placed at right angles to each other (ahead, behind, left, right) that can receive signals uniformly within a 90° sector. The strength of the signals from other robots detected by any one receiver are additive and are inversely proportional to the square of the distance between the emitter and the receiver, with a 10% random component.

At each time step, each robot rotates by a variable angle ÎÈ as as a function (see below) of the input from its four receivers. It then moves in a straight line at a fixed speed for one time unit. Different robots have slightly different speeds, with a 10% variation. The robots move in a 200 x 200 unit square arena. In case of collision, whether with a wall or another robot, the robot selects a new direction at random.

The rules and the results

The following parameters intervene in the decisions:

S_0: the threshold signal strength, equivalent to no signal detected.
S_{low}: a low signal strength.
S_{high}: a high signal strength.
ÎÈ : the change in direction (degrees).

A signal strength of 0.01 is equivalent to a signal from a robot 10 length units away. Unless otherwise stated, the simulations were performed with 10 robots, initially bunched at the centre of the arena.

Rule 1

> if (signal ahead > S_{low}) then maintain direction.
> Otherwise, if (signal left > S_0) then turn left ÎÈ,
> else if (signal right > S_0) then turn right ÎÈ.

Rule 1 says that as long as there are robots sufficiently nearby ahead of you then carry on. Otherwise, if you can detect something at all to the left or to the right then turn towards it. Intuitively this is a follow-my-leader process, and one of the stable configurations it gives is an anti-clockwise roundabout movement, with a wide range of parameter values centred around S_0=0.0001, S_{low}=0.01, ÎÈ=20-60° (fig. 1a). This configuration is reached very quickly, and while it can be perturbed, it rapidly reforms.

Fig. 1b shows how this works. Robot 1 detects a signal in its ahead receiver, and so maintains its direction. Robots 2 and 3 detect no signal in their ahead receivers, but do detect a signal in their left receivers, and so turn left, thereby maintaining the roundabout movement.

For low values of ÎÈ, around 20°, and S_{low}<0.001, the radius of the circle increases as S_{low} decreases. The robots starting at the same point move away from each other in an anti-clockwise spiral pattern (fig. 1c), which either settles into a circle or reaches the boundary limits, in which case the robots move around the edge of the room. It is interesting to note that the dominant anti-clockwise direction is due to the test for turning left in rule 1 coming before the test for turning right. Reversing the order makes clockwise movement dominate.

While the roundabout pattern dominates the (S_{low}, ÎÈ) space, three other regions have stable movement patterns. With parameters around S_0=0.0001, S_{low}=0.1, ÎÈ=5°, each robot makes anti-clockwise looping movements, with the group staying together, and

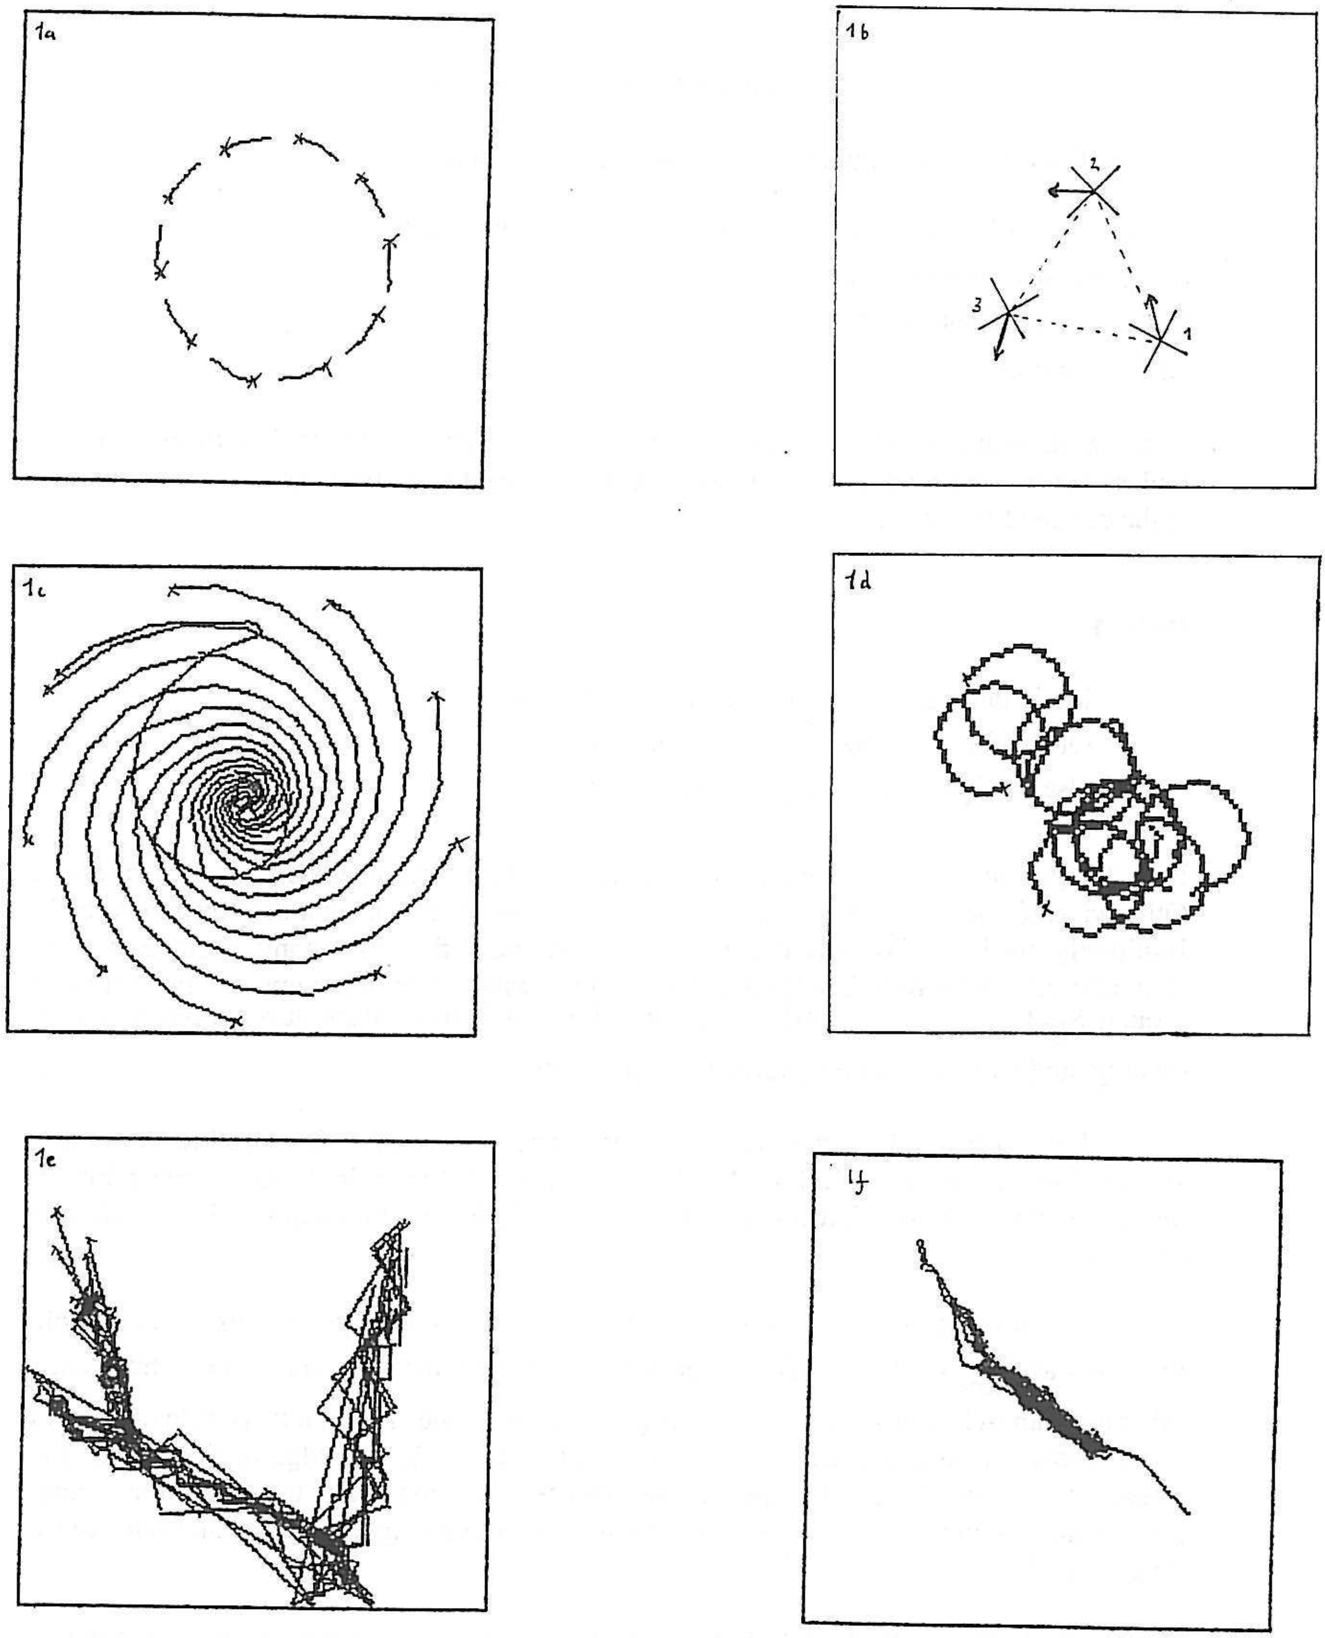

Fig. 1. Different patterns generated by rule 1. When present, x marks a robot, the lines trace their previous positions. **1a**. Roundabout. **1b**. Schema of 3 robots interacting. The arrow indicates their orientation, the cross indicates the four receivers each covering 90°, the dotted lines indicate the direction of signals they receive. **1c**. Spiral (one robot began in a clockwise direction). **1d**. Circling lek (enlarged 2x). **1e**. Travelling lek. **1f**. Elastic line pattern.

remaining at the same point for long periods of time (fig. 1d). The overall impression is that of a group of flies, termed a lek, circling around a fixed point, with individuals occasionally breaking away and buzzing around the room before being reabsorbed in the lek.

With parameters around S_0=0.0001, S_{low}=0.001, ÎÈ=80°, the robots form a more or less triangular group that moves in a straight line. At the apex of the triangle, one of the robots acts as a leader. Receiving no signal from ahead or laterally, it moves in straight line. Those behind it are content to follow it, and at the rear of the triangle, receiving more complex signals, the robots tend to zig-zag back and forth, elongating the triangle. The group alternates this pattern with brief periods of tight milling about a point, and with short periods when it fragments into independent sub-groups with similar behaviour, and that merge when they pass near each other. Fig. 1e shows how the group moves about the room, starting at the top right corner, and changes in direction being preceded by a a period of milling.

Finally, as S_{low} increases from 0.001 towards 0.1, with S_0=0.0001, ÎÈ=80° as for the moving triangular pattern, the robots spontaneously spread themselves out in line, moving backwards and forwards along the length of the line, with occasional milling at its centre (fig. 1f). The line stretches and shrinks (elastic line pattern), and its orientation drifts, but the robots do not stray far from their common axis.

Depending on the parameter values, rule 1 can thus generate four stable patterns, a roundabout, a looping stationary lek, a travelling triangular group, and an elastic line.

The roundabout pattern which dominated the parameter space in rule 1 can for certain purposes be considered as a dead-end. The following modification to line 1 is sufficient to eliminate the roundabout pattern. With ÎÈ<25°, the robots form the looping, stationary lek found in rule 1 (fig. 1d). With ÎÈ>30°, the robots form the elastic line found in rule 1 (fig 1f).

randomly, every 5 steps, turn by a random value between -ÎÈ/2 and +ÎÈ/2.
Otherwise, if (signal left > S_0) then turn left ÎÈ,
else if (signal right > S_0) then turn right ÎÈ.

Rule 2

if (signal ahead > S_{low}) then maintain direction.
Otherwise, if (signal left > signal right) then turn left ÎÈ,
else turn right ÎÈ.

The first part of rule 3 is as rule 1, but the second part differs in that it compares signal strengths from the left and right receivers. As in rule 1, the roundabout pattern dominates the (S_{low}, ÎÈ) space, and with ÎÈ>40 becomes a tight milling pattern. With values around S_{low}=0.1, ÎÈ=10, the robots stay grouped together in a lek-like pattern, occasionally breaking off and then rejoining the group.

Within a small range of values around S_{low}=0.01, ÎÈ=15, the robots tend to form a pattern in which they follow each other in a figure of eight, that moves as a whole in a straight line perpendicular to its widest axis (figs. 2a,b). The robots can be "trapped" in a roundabout formation, but the pattern can be made more stable by allowing those at the edge of the 8 to speed up, without of course their having to know where they are. This can be achieved by adding the last three lines as below.

if (signal ahead > S_{low}) then maintain direction.
Otherwise, if (signal left > signal right) then turn left ÎÈ,
else turn right ÎÈ.
if ((signal left > S_0) and (signal right < S_0))
or ((signal left < S_0) and (signal right > S_0))
then double speed, else normal speed.

Rule 3

for each signal
  if (signal > S_{high}) then signal:= - signal
  else if (signal > S_{low}) then signal:= 0

calculate dX = signal ahead - signal behind
calculate dY = signal left - signal right
if (|dx|<S_0) and (|dy|<S_0)
  then turn randomly left or right between 0 and ÎÈ/2°
else turn to the same axis as the (dX, dY) vector

(:= reads "becomes equal to"). Roughly speaking this rule says that if the signals are of medium strength move randomly, if they are too strong move away, if too weak move nearer. The *a priori* intention was to space the robots out evenly. The algorithm does this, but moreover tends to string them out in a line (S_{high}=0.01, S_{low}=0.007) (fig. 3). The line is not very straight, and bends and rotates, but is very stable, each robot oscillating over a small part of the distance between its neighbors (line dispersion pattern).

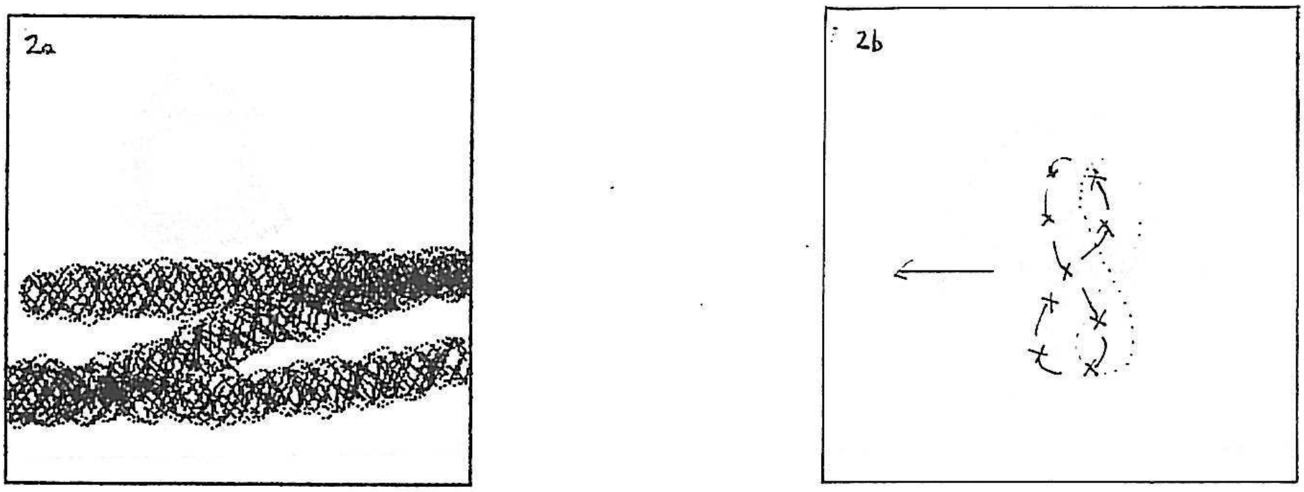

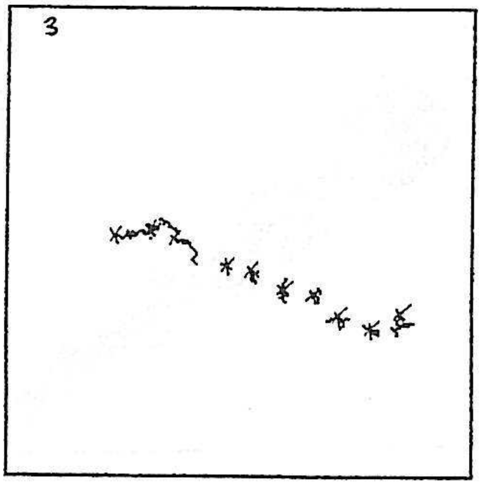

Fig. 2a. Figure of 8 pattern generated by rule 3. **Fig. 2b.** Schema of figure of 8 pattern. Dotted line represents position shortly before, arrow indicates overall direction of movement.

Fig. 3. Line dispersion pattern generated by rule 4.

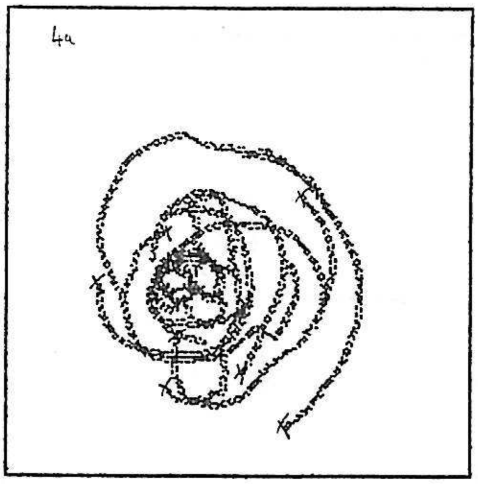

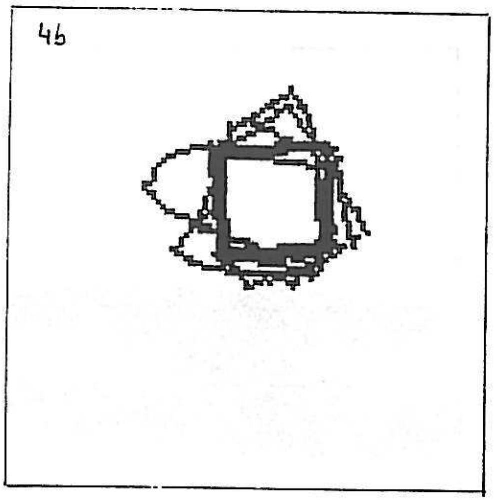

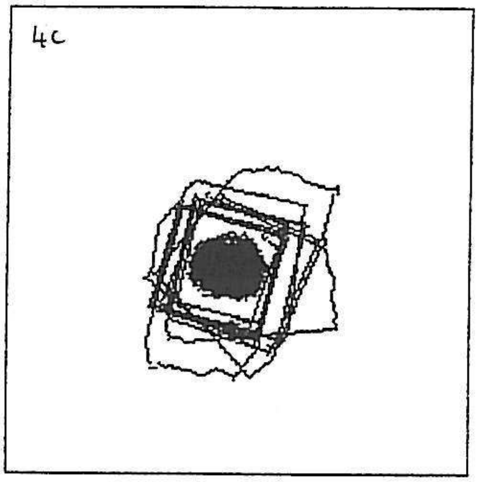

5

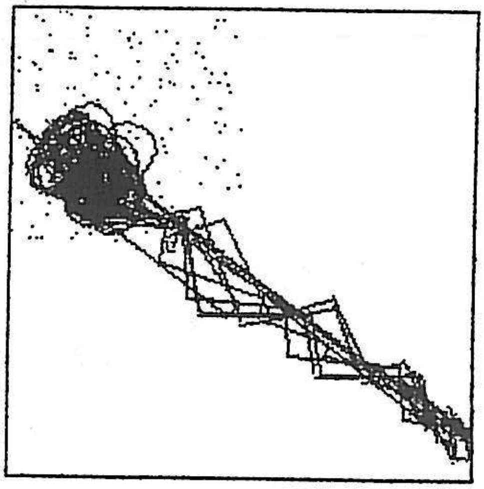

Fig. 4a. Autumn leaf pattern generated by rule 5. **Fig. 4b.** Square orbit pattern from rule 5, with 2 robots instead of 10 (enlarged 2x). **Fig. 4c.** Concentric square pattern from rule 5, with 3 robots. Two turn around each other at the centre, the third moves in a square orbit around them.

Fig. 5. Travelling lek pattern moves in straight line, pattern changes to circling lek pattern upon entering top left quadrant containing items.

Rule 4

calculate dY = signal ahead - signal behind
calculate dX = (signal left + signal right) / 10
if ($|dx|<S_0$) and ($|dy|<S_0$)
  then turn randomly left or right between 0 and ÎÈ/2°
else turn to the same axis as the (dX, dY) vector

This rather arbitrary adaptation of rule 4 generates some interesting patterns. With 10 in the group, the robots move in a clockwise vortex, like autumn leaves whirling about each other in a twisting wind, as one observer put it (fig. 4a). The group can split into sub-groups. In each sub-group the robots turn about each other, and the more robots the wider the orbits. The sub-groups drift, enter each other's influence, and merge together.

With only two robots in the group, they turn about each other in a roughly square orbit, one being opposite the other (fig. 4b). With three robots in the group (fig. 4c), two form a tight central sub-group turning about each other, while the third performs a square orbit about the first two!

Adapting the pattern to fit the circumstances.

This final section explores the idea that one pattern might be appropriate in one situation and another pattern in a second situation. How then could one arrange for the group to switch from one pattern to the other without explicitly telling it to do so, and without requiring the robots to recognise the different situations? For example, consider a heterogenous area, one part of which is comparitively empty and requires a wide-ranging exploratory movement such as the moving lek of rule 1. When the group comes across the second richer zone filled with "interesting items", however these might be defined, a pattern of movement that keeps them in the rich zone, such as the circcling lek movement of rule 1 would be more appropriate.

Suppose that the interesting items appear in zone 1 but not in zone 2. Start the robots with the moving lek rule. If a robot meets an item it "collects" it and switches to the circling lek rule for a limited time period of 200 steps before reverting to the moving lek rule (unless it meets another item in which case it stays for a new 200 steps in the circling lek rule).

Initially the group would adopt the moving lek pattern, keeping in it as they move around the empty or low density zone. Meanwhile items are accumulating in the richer zone. When the group by chance enters the rich zone, many of them quickly encounter an item and switch rules. If the density is high, enough of them do so for the group to adopt the second pattern which keeps them in the same spot. Also, the robots will probably meet another item before the 200 step limit is reached or shortly after, and the second pattern will be maintained. The robots "consume" the items until the density has dropped, or they stray out of the rich area, one by one switching back to the moving lek rule. The group reforms with the first pattern and starts exploring again.

Fig. 5 illustrates this, showing a group travelling in straightish lines in the empty zone, and being "captured" by the richer zone in the top left corner, the circling lek movement keeping them there as long as there are enough items.

Discussion

The approach developed here towards the control of collective movement is intended as a complement to a more centralised or hierarchical control. As on-board visual analysis, learning capacity and inter-robot communication become more powerful and more reliable it will be possible to generate more regular patterns of movement using a more direct approach. Nevertheless such self-organised movements have the general advantage of readily adapting to environmental heterogeneities and different forms of perturbation, such as the breakdown of one or more members of the group. They are conceived to function without supervision and without prior knowledge as to the lay-out of the area in which they are placed. Furthermore, they only require the robots to have local sensors and relatively simple short-range communication systems. There are a number applications for which such advantages could be decisive.

Acknowledgements

This work is supported in part by the Belgian programme on inter-university attraction poles. Jean Louis Deneubourg is a fellow the Belgian Fonds National de Recherche Scientifique.

References

Braitenberg V. 1984. Vehicles. Experiments in synthetic Psychology. MIT press, Cambridge MA.

Breder C.M. 1959. Studies on social groupings in fishes. *Bulletin of the American Museum of Natural History* **98**: 397-481.

Breder C.M. 1976. Fish schools as operational structures. *Fishery Bulletin* **74**: 471-502.

Deneubourg J.L. and Goss S. 1989. Collective patterns and decision-making. *Ecology, Ethology & Evolution*, **1**, 295-311.

Deneubourg J.L., Goss S., Franks N.R., Sendova-Franks A., Detrain C. and Chrétien L. 1990. The dynamics of collective sorting: robot-like ants and ant-like robots. In: Simulation of Adaptive Behaviour: from Animals to Animats. Meyer J.-A. and Wilson S. (eds.), MIT Press, 356-365.

Goss S. and Deneubourg J.L. 1992. Harvesting by a team of robots. In: Proceedings of the 1st European Conference on Artificial Life, Paris 1991. Bourgine P. and Varela E. (eds.), MIT Press, Cambridge Mass.(in press).

Huth A. and Wissel C. 1989. How can fish school? In: Biological motions. Alt W. and Hoffmann G. (eds.). Berlin: Springer-Verlag (in press).

Jacobson H.J. 1958. On models of reproduction. *American Scientist* **46** (3), 255-284.

Parr A.E. 1927. A contribution to the theoretical analysis of the schooling behaviour of fishes. *Occasional Papers Bingham Oceanographical Colloque* **1**, 1-32.

Reynolds C.W. Flocks, herds and schools: A distributed behaviour model (Proceedings of SIGGRAPH '87). *Computer Graphics* **21**, 25-34.

Sakai S. 1973. A model for group structure and its behaviour. *Biophysics* **13**, 82-90.

Suzuki O.G. and Sakai S. 1973. Movement of a group of animals. *Biophysics* **13**, 281-281.

Theraulaz G., Goss S., Gervet J. and Deneubourg J.L. 1990. Task differentiation in Polistes wasp colonies: a model for self-organising groups of robots. In: Simulation of Adaptive Behaviour: from Animals to Animats. Meyer J.-A. and Wilson S., (eds.). MIT Press, 346-355.

Walter W.G. 1950. An imitation of life. *Scientific American* **182** (5), 42-45.

Walter W.G. 1950. A machine that learns. *Scientific American* **51**, 60-63.