

Stigmergy, self-organisation, and sorting in collective robotics

Owen Holland and Chris Melhuish
Intelligent Autonomous Systems Engineering Laboratory
Faculty of Engineering
University of the West of England
Coldharbour Lane
Frenchay
Bristol BS16 1QY

Submitted to Artificial Life

1. Introduction: Stigmergy

Stigmergy is a concept occasionally used in biology to describe the influence on behaviour of the persisting environmental effects of previous behaviour. It was originally proposed by Grassé [12] to explain some of his observations on termite building behaviour. Grassé had observed that worker termites in the presence of particular configurations of a construction (and of other workers) would be stimulated to a high degree of activity, and would tend to add building material to specific parts of the construction. As the construction was changed by these additions, the site of addition of further material would be modified, leading to the progressive growth and completion of the feature; the termites would then switch to constructing another such feature, or would begin a new task apparently triggered by the presence of the completed feature. The phrasing of his introduction of the term is worth noting:

“La coordination des taches, la regulation des constructions ne dependent pas directement des ouvriers, mais des constructions elles-memes. *L’ouvrier ne dirige pas son travail, il est guidé par lui*. C’est à cette stimulation d’un type particulier que nous donnons le nom du STIGMERGIE (*stigma*, piqueure; *ergon*, travail, oeuvre = oeuvre stimulante).”

[“The coordination of tasks and the regulation of constructions does not depend directly on the workers, but on the constructions themselves. *The worker does not direct his work, but is guided by it*. It is to this special form of stimulation that we give the name STIGMERGY (*stigma*, goad; *ergon*, work, product of labour = stimulating product of labour).”]

In the English summary, the concept is expressed more directly:

“The stimulation of the workers by the very performances they have achieved is a significant one inducing accurate and adaptable response, and has been named *stigmergy*.”

If stigmergy is indeed at the root of the building behaviour of termites, ants, bees, wasps, and other social insects, then it is certainly a powerful principle, as social insect constructions are remarkable for their complexity, size, and adaptive value. However, it is possible to extend the idea easily to other domains; it can then be seen as an even more impressive and general account of how simple systems can produce a wide range of apparently highly organised and coordinated behaviours and behavioural outcomes, simply by exploiting the influence of the environment. In Grassé's vision, a worker deposits a piece of building material (does 'work') in a particular location; this changes the sensory input subsequently obtained at that location, and hence may change the behaviour produced (and the work done) at that location in the future. If a drop of pheromone was deposited instead, it could also change the behaviour at that location in the future, at least until such time as it had completely evaporated. The laying and sensing of pheromones, especially in the form of chemical trails, underlies many of the spectacular abilities of ants, especially in the control of foraging, and is clearly an instance of Grassé's concept; the modern practice is to extend the definition of stigmergy by replacing the sense of 'work' with the sense of 'any environmental change produced by the animal'. (In fact the regulation of termite building behaviour is now understood to involve pheromones as well as constructional features.)

Although there have been several modern treatments of stigmergy as a general phenomenon, there is still room for more precision in its definition. For instance, Grassé's original sense of 'stimulation' should formally be refined. All that is necessary for stigmergy to occur is for the *outcome* of the behaviour of the relevant agent to be appropriately affected by previous environmental changes, and this can happen in a number of distinct ways:

- (i) the agent's choice of action may be affected (a qualitative effect)
- (ii) the selected action may be unchanged, but the exact position, strength, frequency, latency, duration, or any other parameter of the action may be affected (a quantitative effect)

(i) captures Grassé's sense of action being guided, and (ii) also includes the element of intensity of activity. The qualitative effect in (i) may of course be internally controlled by some threshold mechanism acting on a quantitatively varying input. However, there is also a third possibility which is not included in Grassé's formulation:

- (iii) a previous action at a location might affect neither the choice nor the parameters of a subsequent action, but only the outcome (a qualitative and/or quantitative effect)

This requires some explanation. Consider a car being driven along a muddy track. Although the driver might try to steer a particular course, the wheels may settle into deep ruts which take the car along another course. The actions taken by previous drivers have affected the outcome of the actions taken by the present driver. (Incidentally, he will have further deepened the ruts, and will have a still harder time of it the next time.) This

influence may be thought of as passive stigmergy [13] whereas (i) and (ii) may be thought of as active [13] in that they affect the agent itself. Passive stigmergy is very close to purely physical situations where some constantly acting force - often a fluid - changes the environment in such a way as to change its future effect on the environment; for example, the formation of sand dunes, river deltas, and meandering rivers are all instances of this.

We are now in a position to ask how stigmergy can actually produce complex patterns, whether of material or behaviours. Stigmergy is essentially a mechanism which allows an environment to structure itself through the activities of agents within the environment: the state of the environment, and the current distribution of agents within it, determine how the environment and the distribution of agents will change in the future. As has been made clear by Bonabeau et al [3], any structure emerging from this repeated interaction develops by a process of self organisation (SO). (See also [7, 16]). Bonabeau and his colleagues have provided a useful brief summary of the nature and properties of SO. They define and describe SO as "...a set of dynamical mechanisms where by structures appear at the global level of a system from interactions among its lower-level components. The rules specifying the interactions among the system's constituent units are executed on the basis of purely local information, without reference to the global pattern, which is an emergent property of the system rather than a property imposed upon the system by an external ordering influence." They go on to identify four basic ingredients of SO, and three characteristic signatures. The ingredients are positive feedback, negative feedback, the amplification of fluctuations, and the presence of multiple interactions; the signatures are the creation of spatiotemporal structures in an initially homogeneous medium, the possible attainability of different stable states (multistability), and the existence of parametrically-determined bifurcations. The mechanism of stigmergy, combined with environmental physics, provides the basic ingredients in social insects; the resultant SO produces outcomes which display the characteristic signatures. Stigmergic SO is distinguished from the purely physical SO mentioned in the last paragraph because it involves mobile agents. Agents can sense the local environment, and act on it, in ways determined by their physical and computational constituents. The possibilities for producing spatiotemporal structures both in the environment and in the distribution of agents within the environment are therefore infinitely greater than those arising directly from the environmental physics. It is this potential richness of behaviour-mediated changes which has been exploited by evolution to produce the striking phenomena found in social insect colonies; Bonabeau et al [3] have pointed out some of the possible ways in which evolution may favour the emergence of some aspects of self organisation.

There are several other comments which may usefully be made about stigmergy, and which can increase our understanding. One approach to stigmergy is simply to consider the minimal qualities of agent and environment which are necessary to support it. An agent has two key abilities: it can move through the environment, and it can act on the environment. The environment must be able to be changed locally by agents; and such changes must persist long enough to affect the choice, parameters, or consequences of agents' behaviour. (This effectively rules out stigmergy in empty or highly dynamic

environments, such as space, air, and water.) Change can be reduced to a small number of categories: material can be taken from the environment, or added to it, or some local quality of the environment can be altered. The scope of stigmergy is thus defined: the three types of environmental change may produce the three types of stigmergic action on the two agent abilities. It should be clear that some form of stigmergy must inevitably be in operation in many biological systems, and can be expected to occur in many artificial systems when they are widely deployed in the real world. However, only those instances of stigmergy which give rise to SO will produce noticeable or useful effects

Some additional clues to the origins and underlying principles of stigmergy can be gathered from the observation that, as Grassé pointed out in his original paper [12], there are two ways of structuring the generation of behavioural sequences in insects (and, by extension, in simple agents of any type). In the first, found in solitary species such as the digger wasp, the execution of the first movement in a sequence sets some internal state which then, often in conjunction with some appropriate external cue, initiates the second movement, and so on. In the second, found in both solitary and social insects, there is no such internal state; the external cue alone is sufficient. The second method often requires that the external cue is correlated with the successful completion of the first movement.

This second strategy is more appropriate for social insects, for many reasons; more importantly, it sets the scene for stigmergy. Because there are many identical agents available, there is no longer any requirement that a connected sequence of movements (or sub-tasks making up a task) must be carried out by a single agent. The presence of the cues alone will ensure that a complete sequence is executed, even if each movement is performed by a different agent. (Where there is no suitable cue available from the end-state of the sub-task itself, it may be necessary to augment the sub-task to provide some additional external cue, or sign.) In particular, where there are many similar cues for a certain sub-task at a given location, the rate of performance of the sub-task will be a function of the number of agents at that location. (This would not necessarily be the case if an agent had to be in a particular internal state in order to be able to respond to the cue.) If there are many locations with such cues, the sub-task will be performed fastest at the locations which have the greatest numbers of agents present. Stigmergy can thus control the morphogenetic development of a construction or other spatial pattern by controlling the distribution of agents within the environment rather than just by controlling the elicitation of building actions at particular sites. Some constraints placed on stigmergic construction algorithms which do not control agent distribution were identified in [15], a computer simulation of a task inspired by the building behaviour of wasps.

Some of the most useful insights into stigmergy have been provided by simulations. For example, in their paper on 'The dynamics of collective sorting: ant-like robots and robot-like ants', Deneubourg et al [8] presented a simulation showing that simple agents, specified in terms which could equally well apply to ants or robots, could use stigmergy to achieve two generic tasks known to be performed by ants, and to be of fundamental importance to them: the clustering of scattered objects of a single type (see also [5]) and the grouping and sorting of objects of two different types. For sorting, the

agents needed to be able to sense the local densities of the different types of brood items, which was achieved by using a short-term memory, and also needed to know the type of any brood item they were carrying. Clustering was the result of the mechanism operating on only a single type of item.

However, studies using artificial physical agents (robots) may be able to yield deeper insights, perhaps because they are embedded in real-world physics, and share its constraints and opportunities with stigmergic social insect systems. Beckers, Holland, and Deneubourg [2] were able to achieve clustering with an even simpler algorithm, using physical robots which were unable to detect whether or not they were moving any objects, which had no memory, and which could sense the local density of objects only as being below or above a fixed threshold. The mechanism was thought to be a form of stigmergy, acting to produce self-organisation. Small clusters were formed at first through the action of the threshold mechanism; by random accretion, some became larger than others; as larger clusters were less likely to lose objects and more likely to gain them than smaller clusters, the eventual outcome was a single large cluster.

This paper describes and discusses several collective robot experiments aimed at extending the approach of Beckers et al to examine the possible role of stigmergy and self-organisation in the simple task of sorting two types of objects, as studied in simulation in [8]. The paper is organised as follows. Section 2 proposes a simple taxonomy of types of spatial sorting. Section 3 gives some examples of clustering and sorting in social insects and robots. In section 4, the robots and the environment used in this work are described. Section 5 describes the algorithms being investigated, and presents the experimental results, which are discussed in Section 6 in the context of stigmergy and self-organisation.

2. Spatial sorting

Spatial sorting is an informal notion; in order to structure this investigation, some formalisation is required. We propose dividing spatial sorting into four basic types:

- Clustering: grouping a class of objects within a continuous area that is a small fraction of the area of the available environment
- Segregation: grouping two or more classes of objects so that each occupies a continuous area of the environment which is not occupied by members of any of the others
- Patch sorting: grouping two or more classes of objects so that each is both clustered and segregated, and each lies outside the boundary of the other
- Annular sorting: forming a cluster of one class of objects, and surrounding it with annular bands of the other classes, each band containing objects of only one type

These are still quite informal definitions, but serve to distinguish between at least some basic types of sorting. It is still possible to see an example of something that is clearly sorting, but to be unable to assign it unambiguously to one of these categories. For

example, if the two types of objects in an environment are sorted so that one type is tightly clustered, but the other is uniformly distributed within the remaining area, should this be classed as mere clustering, as segregation, or as annular sorting with a broad outer ring? We do not propose to answer these questions here; we recognise that they may arise, but find it more appropriate in these early investigations to defer considerations until the distinctions which need to be made become apparent in the context of research findings.

3. Sorting in social insects and robots

The type of sorting carried out by some social insects, which corresponds to patch sorting, was summarised by Deneubourg [8]: "...the eggs are arranged in a pile next to a pile of larvae and a further pile of cocoons, or else categories are placed in entirely different parts of the nest...if you tip the contents of a nest out onto a surface, very rapidly the workers will gather the brood into a place of shelter and then sort it into piles as before." However, sorting need not be into piles; Franks and Sendova-Franks [10] describe the brood sorting behaviour of *Leptothorax unifasciatus*, the outcome of which is concentric rings of progressively more widely spaced brood items at different stages of development. This corresponds to annular sorting. It should be noted that social insect sorting is typically imperfect and noisy, and that frequent or continuous rearrangement is often seen.

The mechanisms of sorting in ants are not known with any degree of certainty, but Deneubourg's model [8] certainly captures many features of the behaviour of some ants. However, Franks and Sendova-Franks [10] note that the model "...is not, in its present form, sufficient to create concentric brood clusters", and speculated that other factors may be involved, such as self-sorting by size [1], or the presence of special pheromones determining the space given to each type of brood item.

There do not appear to be any accounts of sorting analogous to brood sorting being carried out by physical robots. As noted above, the related task of clustering using minimally capable robots was demonstrated by Beckers, Holland, and Deneubourg [2], and Maris and te Boekhurst [15]. Most accounts in the literature deal with foraging tasks where the objects must be brought back to a designated location which is found by some means of navigation; see Cao *et al* [6] for a useful bibliography and review.

4. Materials and methods

The robots and experimental environment used in this study were designed to investigate a range of social insect behaviour, with particular emphasis being placed on building tasks. The design of the robot system was mainly inspired by the work of Franks and his collaborators on *Leptothorax tubero-interruptus*, *Leptothorax unifasciatus*, and related species [10, 11]. These ants live in small colonies (typically with a few hundred members, but occasionally with as few as fifteen) in cracks in rock; this constrained environment means that they are behaviourally adapted to life in two dimensions. It is therefore

possible to study the behaviour of a colony by providing it with a particular two dimensional habitat - the space between two glass slides - which allows unrivalled opportunities for observation and recording.

Leptothorax is also peculiarly suitable as a model for robotic investigations of some collective behaviours, because current robots, which are wheeled, operate well in two dimensions but are extremely limited in their abilities to operate in the third dimension. *Leptothorax* building behaviours tend to involve the movement of single lumps of material (carborundum grit of regular dimensions is usually provided by the experimenters) which are placed next to other lumps rather than piled on top of them. A similar level of functionality can be achieved in robots by simply pushing and pulling building blocks around the floor, rather than lifting them and piling them up.




Figure 1: A U-bot

The robots (Figure 1), known as U-bots, were designed and built in our laboratory to provide a flexible and capable platform for a range of collective robot experiments. They have the following features:

- size: the robots are small enough (23cm diameter) to be easily portable, yet are large enough to operate on most floor surfaces, and to carry most sensor types (ultrasonic, IR, CCD video cameras)
- manoeuvrability: differential drive, powerful motors, and high-resolution optical quadrature encoders enable turning on the spot, reversing, and tight control of speed and position
- endurance: each robot will run for around 3 hours under conditions of frequent acceleration and deceleration, with all electronics operational
- computational power: a Motorola 68332 processor with up to 16 Mb of memory provides ample capacity

- flexibility: extra power rails, multiplexed I/O ports, and A/D conversion provide for the addition of a range of sensors and effectors at a later date
- reliability: the precision machined aluminium chassis bears all structural loads, and provides a secure fixing for PCBs, sensors, and effectors; the gearboxes are protected from shock loads by a preset torque-limiting clutch; all input, output, and power lines are protected or fused.

For this set of experiments, the robots were equipped with four infra red proximity sensors (three facing forwards, one backwards) set to around 20cm. The main distinguishing feature, however, is the gripper fitted to the front of the robots, which was designed to sense, grip, retain, and release frisbees. The gripper projects a maximum of 8cm beyond the robot's circular chassis. The frisbees, which are either red or yellow, and are 23cm in diameter, are placed on the floor concave side up. If a robot moves towards a frisbee which is directly in its path, the frisbee will fit neatly inside the semicircular part of the gripper (Figure 2). As the gripper moves round the frisbee, the raised edge of the frisbee contacts the lower parts of two small weighted 'barbels' mounted above the gripper which are free to swing fore and aft, but cannot move sideways. As the gripper continues to move, the barbels ride over the edge of




Figure 2: a U-bot gripping a frisbee

the frisbee and dangle freely inside the concave area. The function of the barbels is to retain the frisbee when the robot turns on the spot, which they do because they are unable to swing sideways. When the frisbee is right inside the gripper, an optical sensor detects its colour. (In order to ensure correct identification under all conditions of lighting, the red frisbees have been painted with black paint, with a white central portion to enable accurate visual location - see for example Figure 3).



Figure 3: U-bots with frisbees

The vital functionality of the gripper is provided by two further mechanisms. The first is a pin mounted at the rear of the gripper, which can be lowered by a small electric motor so that it projects down inside the concave rim of the frisbee. If the pin is in the raised position, and the robot reverses, the frisbee will be left in position. (The barbels are passively swung forwards by the rim.) However, if it is in the lowered position, it will engage the rim, and the frisbee will be pulled backwards.

The second mechanism is similar to that used by Beckers *et al* [2]: the whole gripper is suspended so that a backwards-acting force greater than some threshold value can move the gripper and trigger a microswitch. Most importantly, the force can act on the ends of the gripper jaws (when a misaligned frisbee is struck) as well as on the area within the jaws. The preset value is adjusted so that the microswitch is not triggered when the gripper strikes or pushes a single frisbee, but is always triggered by two or more.

The robots are run in a large arena (Figure 4) which is octagonal in shape, with 4 metre sides; the area is 1760 times the area of a robot. (This is the same order of magnitude as the ratio of the area of a *Leptothorax* nest to a single ant.) A video camera mounted 6 metres above the centre of the arena is linked to a PC-based video capture system which automatically records an image at preset intervals. The intention is that eventually the images will be automatically processed by an image analysis system to yield the coordinates of frisbees of each colour;



Figure 4: The robot arena

however, the distortion produced by the lens means that a calibration and preprocessing stage must first be undertaken, and this has not yet been completed. (The images reproduced in this paper show this distortion, which overemphasises the central area of the arena, and care must be exercised in their interpretation.)

5. Experiments and results

Experiment 1: Adequacy of arena size

Before carrying out any clustering experiments, we decided to assess the intrinsic characteristics of our robots when operating in our arena. Beckers et al had noted a steep and progressive rise in the number of collisions between robots as the number of robots in their arena was increased, and had identified this factor as being responsible for the large deterioration in performance when the number of robots was increased beyond a small limit. We wished to avoid this deterioration if at all possible; we therefore repeated their strategy of measuring the number of collisions between various numbers of robots both in an empty arena, and in an arena containing the movable objects. Because we were using greater numbers of robots, it would have been practically impossible to record collisions by observation as they had done, and so the process was automated by programming each robot to record the number of collisions it experienced with other robots and with the boundary. (The robots are unable to discriminate between the two conditions). From 1 to 13 robots were run for 20 minutes in an empty arena. The runs were then repeated with a centrally placed cluster of 22 frisbees. The average numbers of collisions for each robot under each condition are shown in Figure 5. As the rate of increase in number of collisions with increasing numbers of robots is low and roughly constant under both conditions, the system is operating under a linear rather than an exponential regime, and we can assume that our results from using up to 10 robots will be interpretable without considering any consequences of direct interactions between robots.

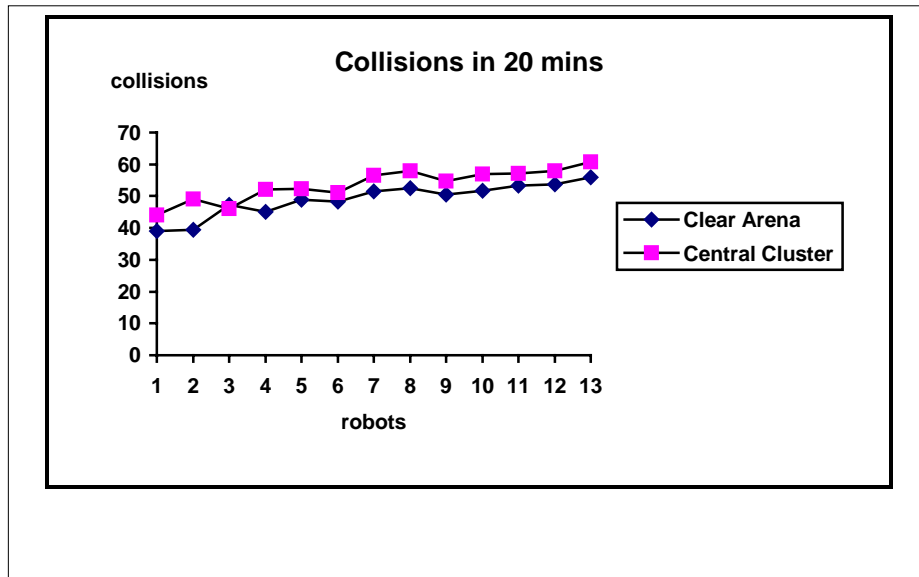


Figure 5: Average number of collisions between robots in 20 minutes with a clear arena, and with a central cluster of frisbees

Experiment 2: Basic clustering

Our first task was to see whether the new robot and gripper arrangement, combined with a suitable algorithm, could reproduce the clustering observed by Beckers *et al* [2]. The robots were programmed with a rule set corresponding to that used by Beckers, but with a slight difference imposed by the nature of the arena boundary. Beckers' arena had a deformable boundary wall which enabled robots to push their grippers past pucks trapped on the boundary, and to scoop them up. The new boundary is made of rigid plastic pipe, and we found that it was extremely difficult to adjust the infra red sensors to the exact distance to enable frisbees left on the boundary to be retrieved. The modified rule set, in decreasing order of priority, became (Figure 6):

```

Rule 1:
if (gripper pressed & Object ahead) then
  make random turn away from object

Rule 2:
if (gripper pressed & no Object ahead) then
  reverse small distance
  make random turn left or right

Rule 3:
  go forward
  
```

Figure 6: Adaptation of Beckers et al rule set for clustering

The rules work as follows. A U-bot approaching another U-bot or the boundary will detect it, but will continue moving until it hits it with its gripper. (The U-bots are designed to withstand frequent collisions). The combination of the gripper being pressed and an obstacle being detected will then trigger a random turn away from the side on which the obstacle was detected. If the U-bot was pushing a frisbee, it will retain the frisbee during the turn due to the action of the barbels. If a U-bot's gripper is pressed in the absence of an obstacle, as happens when it attempts to push more than one frisbee, the U-bot will reverse for a short distance (just over half a frisbee radius) and make a random turn; the effect of this is to leave the frisbee at the location where the gripper was triggered. If neither of these conditions is met, the U-bot moves forwards in a straight line.

If a frisbee is touching the boundary, a U-bot heading straight for it will receive it into its gripper, which will immediately be triggered; since the boundary wall will have been detected, the U-bot will obey Rule 1 and turn through a random angle away from the boundary, keeping the frisbee in its gripper. It will then set off in a straight line (Rule 3) in the new direction, taking the frisbee with it. The U-bots can thus remove single frisbees from the boundary. However, if a frisbee lies so that it is touching a frisbee which is in turn touching the boundary, a U-bot which runs into the first frisbee will have its gripper triggered, but will be too far away from the boundary to sense it. It will therefore obey Rule 2, and will back off the frisbee, leaving it in place.

44 frisbees were placed in the arena at a uniform and regular spacing. In anticipation of the sorting work reported later, half of the frisbees (blacks) were painted black, and half were painted black with a white centre (rings); at this stage, the robots were programmed to treat both alike. 10 robots were released, and the video system was set to record a frame every 5 minutes. The frisbees behaved very much as had the pucks in the experiments of Beckers et al, first aggregating in small clusters, then forming larger ones, and eventually, after 8 h 25 m, forming a cluster of 40 frisbees, each separated by less than one frisbee diameter from the others. (As the experiments ran so slowly, we had decided that using the same termination criterion as Beckers et al -100% of objects in a single cluster - would take so much time that we would be severely limited in our ability to investigate a wide range of experimental conditions. We therefore used a criterion of 90%; given that the natural end point is typically a dynamic equilibrium, this seems justifiable).

Figure 7 shows the starting, intermediate, and final stages of this experiment. They correspond so clearly to those reported by Beckers et al that there is little to add to their description; we feel it reasonable to conclude that the two situations are illustrations of essentially the same process. However, we noticed that in the middle phase of each trial the robots had great difficulty 'stealing' frisbees from the intermediate sized clusters. The problem seemed to be geometric in origin: even in a small cluster, unless the form of the cluster was extremely ragged, it was rare for a straight line trajectory to strike the cluster at a point which would allow a frisbee to be removed but would not trigger the gripper.

With the Beckers et al setup, a robot could strike a cluster almost tangentially and remove one or two pucks without triggering the gripper







Figure 7: Progress of Experiment 2

However, we had become interested in the question of the extent to which either set of experiments corresponded to the clustering of ant corpses by ants, as described in [8]. Our concern was aroused by a note in Bonabeau et al [4]: "It must be emphasised that a very large arena was necessary in Deneubourg et al's experiments to obtain "bulk" clusters: in effect, ants are attracted towards the edges of the experimental arena if these are too close to the nest, resulting in clusters almost exclusively along the edges." We therefore decided to carry out a series of experiments to see how behaviour at the boundary could affect the clustering process.

Experiment 3: Algorithmic boundary effects

While modifying the Beckers et al algorithm to cope with the rigid boundary, we had noticed that setting the infra red detection distance too short would occasionally cause the robots to deposit frisbees at the boundary. This guided us towards a strategy of systematically varying the probability that robots would leave frisbees at the boundary, and evaluating the effects. This could most easily be controlled by varying the algorithm. Figure 8 shows the format which was adopted. Rule 1, which is triggered by the detection of the boundary condition (gripper pressed & Object ahead) has two possible actions, giving two possible outcomes if the robot is carrying a frisbee (retaining the frisbee, or leaving it on the boundary). On each boundary detection, the action to be performed is decided by selecting the action probabilistically, with p representing the probability of retention.

Rule 1:
if (gripper pressed & Object ahead) **then**
 with probability p
 make random turn away from object
 else with probability $(1-p)$
 reverse small distance (dropping the frisbee)
 make random turn left or right

Rule 2:
if (gripper pressed & no Object ahead) **then**
 reverse small distance (dropping the frisbee)
 make random turn left or right

Rule 3:
 go forward

Figure 8: Rule set for probabilistic action selection

44 pucks were set out in a regular pattern. The overhead camera captured a frame every 5 minutes. 10 robots were used, and the following values of p were investigated: 0.0, 0.5, 0.8, 0.88, 0.9, 0.95 and 1.0. As the time for this series of experiments was severely limited by external factors, only one trial was conducted for each value of p , and in some cases a trial was terminated when two large central clusters had formed. (From

previous observations it was known that such a configuration could remain for a very long period before one single cluster would eventually form; such a configuration was therefore considered to indicate that a single cluster would be formed eventually.)

Table 1 sets out the results for the experiment; the corresponding configurations of frisbees at termination are shown in Figure 9. With $p = 0$, after 3 h 15 min all the frisbees were strewn around the periphery with no perceptible grouping; no qualitative change was seen up to 11 h when the experiment was terminated. With $p = 0.5$, all frisbees were at the periphery after 40 min, but again no grouping was seen up to termination after 11 h. However, with $p = 0.8$ and 0.85 a single major cluster was formed at the boundary, with some 15 singletons strewn around the periphery; this seemed to be stable, and termination was again at 11 h. With $p = 0.88$, the size of the major cluster against the wall reached the 90% criterion of 40 pucks, and was stable up to 11 h 20 min. Probabilities of $p = 0.9$ and 0.95 produced two main central clusters after 5 hours and 2.5 hours respectively. From earlier experience and observation it was concluded that these would eventually form one large central cluster and the runs were halted. With $p = 1$, the expected central cluster was allowed to emerge completely, after 6 h 35 min.

There was clearly a systematic progression as p was increased, with a qualitative transition at around $p = 0.88$. We therefore carried out five further trials with this p value.


probability of retention p	RESULTS
1.0	leads to a central cluster after 6 hours 35 minutes
0.95	leads to a central cluster, stopped when 2 main central clusters formed. Stopped ~2.5hours
0.9	leads to a central cluster, stopped when 2 main central clusters formed. Stopped ~5hours
0.88	1 cluster formed at edge. 40/44 at 9hrs 5m continued to be stable up to 11hrs 20min.
0.85	1 major cluster formed at edge and approx. 15 singletons around the periphery. stopped after 11h
0.8	1 major cluster formed at edge and approx. 15 singletons around the periphery. Stopped after 11hrs
0.5	All pucks taken to periphery (frame 8, 0hr40mins)but no single cluster formed Stopped at 11 hrs
0.0	All pucks taken to periphery (frame 3 0hr15m) but no single cluster formed. Stopped at 3hrs.

Table 1: Final states of experiment 3 as a function of p values


Of the extra trials, three also formed a single cluster at the boundary (after 9hrs 25m, 10hrs 35m and 13hrs 10m) and two formed central clusters, one with 40 frisbees (6hrs 0m) and the other with 35 frisbees, with 9 strewn around the periphery (7hrs 35m).




9 (a) $p = 1.0$




9 (b) $p = 0.95$



9 (c) $p = 0.9$




9 (d) $p = 0.88$




9 (e) $p = 0.85$



9 (f) $p = 0.8$



9 (g) $p = 0.5$



9 (h) $p = 0.0$

Figure 9: Final states of experiment 3 as a function of p values

This series of experiments is particularly interesting, for two reasons. First, it displays all the characteristic signatures of self organisation identified in [3]: the creation of spatiotemporal structures in an initially homogeneous medium (the arena with the initial regular distribution of frisbees); the possible attainability of different stable states (the different stable outcomes at $p = 0.88$); and the existence of parametrically-determined bifurcations (the transition from peripheral to central clustering at around $p = 0.88$). Second, it shows that a single rule set can produce both of the qualitative possibilities (peripheral and central clustering) that are observed in the natural system which the system is attempting to model. Taken together, these observations may provide some indication that the robot system and the ants are indeed using the same underlying principle of stigmergically mediated self-organisation.

Experiment 4: Sensor-mediated boundary effects

As a final exploration of clustering, we attempted to induce peripheral clustering not by varying the algorithm, but by varying the sensor characteristics. As noted above, we had observed that certain sensor settings could induce the robots to leave frisbees at the boundary. It was also clear that peripheral clustering was produced only by a certain range of probabilities of leaving frisbees at the boundary. Since the robots change course by making turns through random angles, it is reasonable to assume that the angle at which they approach and strike the boundary is randomly drawn from some (unknown) distribution. By arranging for the robots to sense the boundary correctly only when striking it within some limited range of angles of incidence, it should therefore be possible to produce the equivalent of the probabilistic action selection used in the previous set of experiments. We did this by adjusting both the range and angle of acceptance of the infra red collision avoidance sensors. In order to reduce the acceptance angle (180° for the unmodified sensor array) the left and right sensors were disconnected and only the central sensor was used. The range of this sensor was then reduced by fixing a single layer of translucent tape over the sensor. For the ten robots used this gave a mean sensing range of 1.08 puck diameter (the mean distance of the centre eye to wall for all robots was found to be 0.88 puck diameters) and a mean acceptance angle of 100.3° . (This arrangement was arrived at by a series of successive modifications interspersed with brief trials.)

Figure 10 illustrates the principle. With only the central sensor functioning, Rule 1 (controlling the 'retain or pick up frisbee at boundary' behaviour) will be triggered only when the robot approaches the boundary at an angle within the central region. In a robot colliding with the wall on a trajectory within either of the two blind regions, Rule 2 (controlling the 'drop frisbee' behaviour) will be triggered instead.

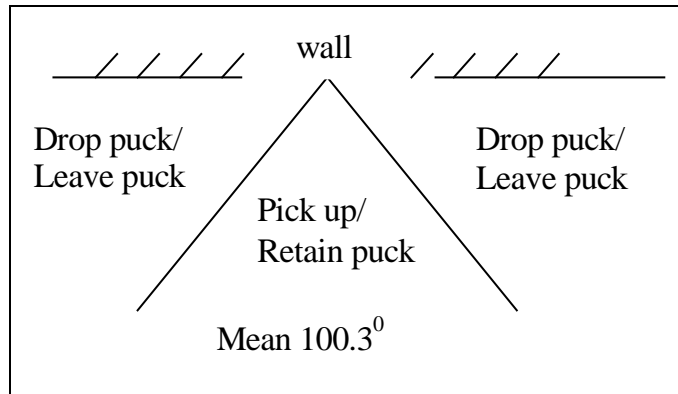


Figure 10: Arrangement of infra-red collision sensor for experiment 4

We recorded three trials using this arrangement, with 10 robots, 44 frisbees initially set out in a regular pattern, and with the overhead camera set to grab a frame every 5 minutes. In all three trials a similar end condition of 40 pucks clustered against the wall in a single cluster was achieved; the times to completion were 10hrs 15m, 13hrs 30m and 14hrs 25m. Figure 11 shows the final frame of the first trial.

It has therefore been possible to produce the same stigmergically self-organised outcome by tuning either a computational parameter or the physical aspect of a sensory parameter. This experimental result is particularly useful in the present context because it emphasises a principle often articulated within robotics: behaviour is the outcome of the interaction between the sensors, the computation, the actuators, the body, and the environment. Since stigmergic self-organisation typically involves the interaction of all these factors over extended time frames, it is probably even more sensitive than single behaviours to the alteration of any aspect of any of the participating factors. This point is easily missed in many discrete grid-based simulations, where environmental features are somehow known directly, and actions always have precise and invariant effects on the environment. This demonstration also serves to show that evolution has many possible sites of action for creating and modulating a stigmergically mediated self-organised outcome - it is not constrained to act only on the computational or algorithmic aspect of the agent.



Figure 11: Final frame at end of first trial of Experiment 4

Experiment 5: The pullback algorithm

Following on from our observation that the slow progress in Experiment 2 might be due to the stability of the clusters, we decided to try to speed up the clustering by modifying the algorithm so that frisbees were not deposited hard up against one another, but some little way off. To do this, it was necessary to employ the pin-dropping mechanism to enable robots to pull the frisbees backwards for some distance before releasing them. We did not want to introduce too much noise into the situation, so we decided to use a mixture of plain yellow frisbees (plains) and black-and-white ring frisbees (rings) but to apply the pull-back tactic only to the plains. The two types of frisbees were identified by combining the inputs from two optical reflective sensors, one positioned to register the reflectance of the central part of a frisbee in the gripper, and one positioned to inspect the outer part.

The rule set now becomes more complex (Figure 12):

```
Rule 1:  
if (gripper pressed & Object ahead) then  
    make random turn away from object  
  
Rule 2:  
if (gripper pressed & no Object ahead) then  
    if plain then  
        lower pin and reverse for pull-back distance  
        raise pin  
    endif  
    reverse small distance  
    make random turn left or right  
  
Rule 3:  
    go forward
```

Figure 12: Rule set for the pullback algorithm

The U-bot's behaviour with respect to rings is unchanged. However, if it is pushing a plain and hits another frisbee, or if it is not pushing a frisbee and collides with a plain which is already on a cluster (perhaps at the boundary), the new version of rule 2 will cause it to drag the plain backwards and leave it the pullback distance away from the contact point. One situation that emerged in early trials was that a reversing robot would frequently strike the boundary. Since the robots have powerful motors with closed loop control, and the boundary is a lightweight structure, the boundary would be moved each time this happened. We therefore activated the rear infra-red sensor to detect obstacles while the robot was reversing; obstacle detection would cause the robot to stop, make a random turn, raise the retaining pin, and move off forwards under Rule 3. (Of course, the rear sensor cannot detect frisbees, so any frisbees behind the robot during a pullback movement are simply pushed out of the way).




Figure 13: The starting configuration for Experiment 5

22 ring and 22 plain frisbees were uniformly spaced throughout the environment (Figure 13); 6 robots were used, with the pull-back distance being set at 2.6 frisbee diameters.

Figure 14 shows the distribution of frisbees after 7 hours 35 minutes, which was the earliest time at which at least 20 of the 22 ring frisbees were in a single cluster. (We defined a cluster as a group of frisbees in which any member was within a frisbee radius of at least one other member.) There is a central dense core of 17 rings, with 11 plains and 4 rings being packed around this core, and the other ring and plains scattered more loosely nearby. This is clearly an outcome which can be regarded as sorting of some kind; within the classification proposed in the introduction, it is at least segregation, and inclines towards annular sorting. We had conceived the algorithm as a means of improving the rate of clustering, and had not thought that it might achieve sorting on its own. However, since the algorithm was probably simpler than anything we might have devised, we decided to explore it further; the remainder of the experiments reported here are the results of that investigation.




Figure 14: Experiment 5 after 7 hours 35 minutes




Figure 15: Experiment 5 after 8 hours 5 minutes

The experiment was allowed to run on for half an hour; the cluster became more compact (Figure 15) but the number of rings and plains in the cluster at 8 hours 5 minutes was unchanged. The progress of the experiment over time reflected the final state in miniature, with a small number of tight ring clusters each surrounded by plains gradually giving way to the eventual pattern. Figure 16 shows the distribution of frisbees after 1 hour 45 minutes.




Figure 16: Experiment 5 after 1 hour 45 minutes

In order to check that this was not a fluke, we repeated the experiment 4 times. In every case, the outcome was similar, in that a single large cluster was formed, with the rings on the inside and the plains on the outside, surrounded by a halo of more distant plains. However, in the replications, the distribution of plains within the cluster was generally less good, in that some plains were trapped within the body of rings. Again, the intermediate stages consisted of smaller clusters with similar morphology.

Trial	1	2	3	4	5
Time in hours	7.58	2.75	25.3	11.7	4.50
Number of plains	11	12	11	10	12

Table 2: Experiment 5: Times to completion for 5 trials with pullback distance of 2.6 diameters

We decided to measure the time at which 20 rings were first in the same cluster. Table 2 shows these times for all 5 trials, and also the number of plains included in the cluster. The number of plains is almost constant; however the times are strikingly variable, with the slowest time (25 hours 20 minutes) being more than nine times the fastest (2 hours 45 minutes). Examination of the video records revealed that the very long times were associated with the formation of two clusters of the same size at opposite sides of the arena; these clusters could remain quite stable for many hours. Other than this, we could see no particular reason for the differences in rates of progress, and can only conclude that, with these parameter settings, the variance of the time to completion is very high.

Experiment 6: Varying the pullback distance

The obvious next step was to vary the pull-back distance to establish any effects on the outcome, the time to completion (defined as 20 of the rings in the same cluster), and the quality of sorting. Single trials were run using pullback distances of 1.3, 3.9, 5.2, 10.4, 15.6, and 26.6 frisbee diameters. In every case a single cluster of rings was eventually formed, except for the trial using 26.6 diameters, which appeared to be making no progress in any clustering whatsoever after 6 hours, and was terminated. The number of plains in the single cluster varied systematically with pull-back distance, decreasing from 18 (at 1.3 diameters) to a minimum of 3 (at 5.2 diameters) and increasing again to 10 (at 15.6 diameters). We added a further 4 trials at the 'best' value (5.2). Figure 10 shows the (average) number of plains in the final cluster for each trial (including data from Experiment 3); Figure 11 shows the (average) times to completion.

The paucity of data points at some values makes interpretation difficult, but it may

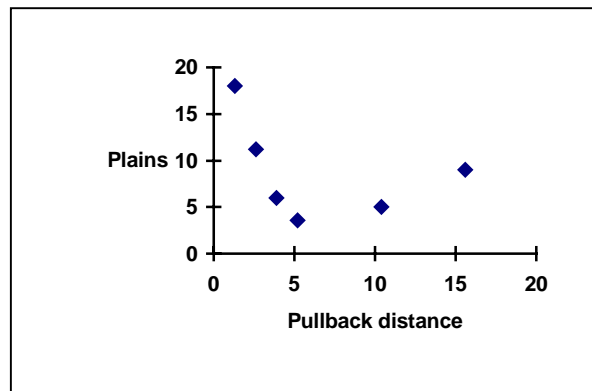


Figure 17: Plains in cluster versus pullback distance

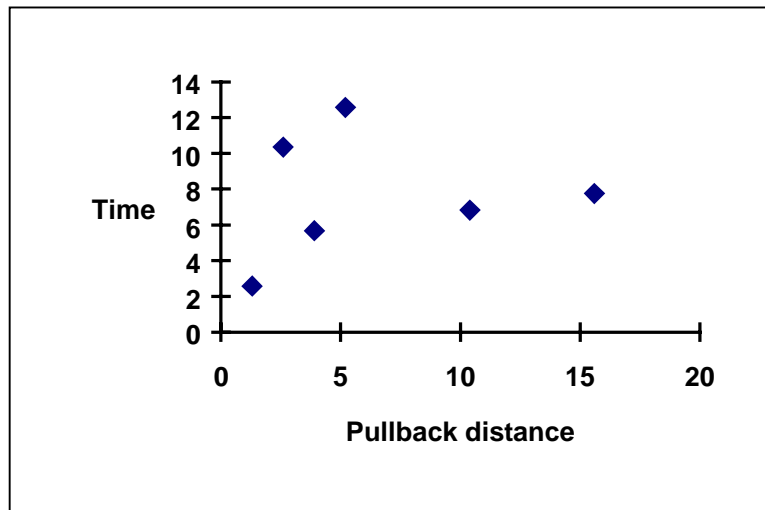


Figure 18: Time to completion versus pullback distance

be reasonable to summarise the data as follows:

- for pullback distances of 15.6 or less, a single cluster containing at least 20 rings is eventually formed
- the number of plains forming part of the cluster decreases with increasing pullback distance to a minimum near 5.2, and then increases.
- for the pullback distance at the minimum, the time to completion can vary considerably

A key additional factor is the spatial distribution of the plains which are not members of the central cluster. (Let us call them detached plains). From observation three things are clear:

- detached plains tend to be found near the final cluster
- as the pullback distance increases from 2.6, the detached plains tend to become more widespread in the arena. (This is difficult to judge for values up to 10.4, simply because there are more detached plains).
- detached plains are often found near the two detached rings at termination

How should we judge the ability of the simple pullback algorithm to produce sorting and segregation? It is clear that the algorithm can produce good segregation when the pullback distance is around 5.2 frisbee diameters (Figure 19). It is also clear that a degree of annular sorting is also occurring in some trials, with rings being concentrated in the centre of the cluster, and plains being found at and beyond the edges of the cluster. However, the values giving good segregation are not the values giving the annular sorting.




Figure 19: Good segregation at pullback distance 5.2 frisbee diameters

At this stage several questions arose. If we ran the experiments for longer, would the plain frisbees tend to form a better grouping? Would plain frisbees show any degree of spatial association if there were no clusters of ring frisbees to act as foci? Could not segregation alone be achieved easily and economically simply by ignoring the plain frisbees, and using the basic clustering algorithm on the rings? And what would happen if, instead of determining the pullback distance by the colour of frisbee, it was determined randomly? All these were examined in a series of experiments.

Experiment 7: Does the clustering of plains improve with time?

22 rings frisbees were placed in a tight cluster in the centre of the arena, and 22 plains were uniformly distributed over the remaining space. 6 robots with a pullback distance of 5.2 frisbee diameters were released. At the end of each hour of elapsed time, the distances of the plain frisbees from the centre of the arena were measured. After 5 hours the experiment was terminated, the frisbees were replaced in their original positions, and a second trial was begun. This had to be terminated after 4 hours due to external factors. The results are shown in Figure 20.

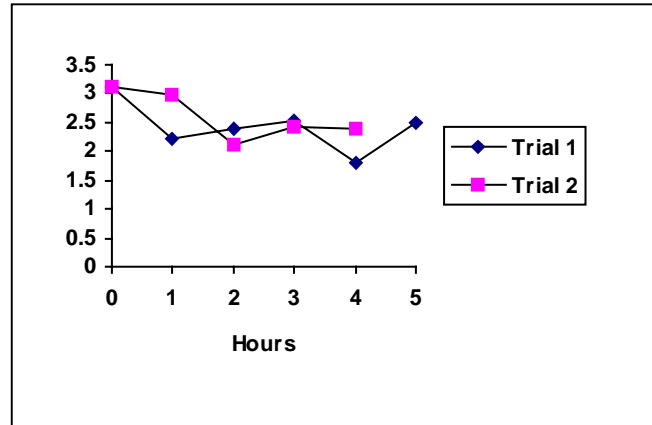


Figure 20: Mean distance of plains from cluster centre versus time

Although there are too few data points to make a firm decision, it is clear that any reduction in mean measured distance from the centre of the cluster is slight, and occurs within a couple of hours. (The initial decrease is of course simply a function of our uniform initial placement.) It does not look as if the termination of the experiments when 20 ring frisbees are first found in a cluster is leading to any significant undervaluation of the ability of the algorithm to cluster the plains.

Experiment 8: Will plains cluster in the absence of rings?

22 plain frisbees were uniformly spaced throughout the arena, and 6 robots with the pullback distance set to 2.6 frisbee diameters were released. After 6 hours there was no sign of any static cluster whatsoever, and the experiment was terminated. However, for much of the time most of the frisbees tended to be loosely grouped together, with a majority of them within an area of perhaps 30% or 40% of the arena. Such an arrangement would dissolve on a time scale of some tens of minutes, only to reform in another location. In order to check that continuing the experiment would not have led eventually to a tight cluster of plains, we performed an additional experiment, using the same robots and frisbees, but starting the frisbees off in a tight central cluster. Within a short time, the central cluster had been broken up; the experiment was continued for 6 hours, and for the remainder of the time it behaved exactly like the first experiment.

Since the pullback distance used was relatively short, and had produced some of the more convincing sorting outcomes in Experiment 5, we concluded that:

- a static focus of ring frisbees is necessary to spatially anchor the dynamic cluster of plain frisbees
- for pullback distances greater than 2.6, the arena may be too small to enable any clustering of plain frisbees to be seen

Experiment 9: Can segregation be achieved by ignoring the plains completely?

Within the confines of our paradigm, the plain frisbees cannot be ignored passively, because frisbees are acquired by the robots in a passive way, by colliding with them. Instead, a plain frisbee that has been collected in the gripper must be released immediately. This is achieved by modifying the rule set - paradoxically, by making it more complex. (Figure 21) This does not mean that plain frisbees will never be moved. In all the experiments described here, both ring and plain frisbees can be moved for short distances when randomly struck by robots. This background of essentially random movement can probably be safely ignored in most experiments; however, in this experiment it became a determining factor.

<p>Rule 1: if (gripper pressed & Object ahead) then make random turn away from object</p> <p>Rule 2: if (plain OR [gripper pressed & no Object ahead]) then reverse small distance make random turn left or right</p> <p>Rule 3: go forward</p>
--

Figure 21: Rule set for Experiment 9

Once again, 22 frisbees of each type were uniformly spaced throughout the arena, and 6 robots were released. To begin with, matters seemed to be progressing as in most of the other experiments, but after a time it became clear that the plain frisbees, which were being moved around a small distance at a time by effectively random contacts, could not easily be removed from the boundary once they had arrived there. In turn, ring frisbees dropped next to these plains also proved difficult to remove. This experiment was terminated after 18 hours; no single cluster was ever formed, but mixed aggregations of plains and rings constantly formed and dispersed close to the periphery.

Experiment 10: Applying the pullback algorithm randomly to rings and plains

In order to see the effects of applying the pullback strategy randomly to both rings and plains rather than always to plains alone, we modified the robot program to determine the variable 'frisbee type' at random when the gripper was triggered, rather than by sensing. The starting arrangement was as in Experiment 3, with 22 frisbees of each type spaced uniformly over the arena; 6 robots were used, with the pullback distance set to 5.6 frisbee diameters. This experiment corresponds exactly to what we had first thought of as the solution to the cluster stability problem; however, our experience with the previous experiments now led us to expect that a single tight cluster would not be formed, and that instead a rather looser and mobile assembly would be seen. We were mistaken: the

normal termination criterion of 40 frisbees in a single tight cluster was reached after 15 hours 15 minutes. This is much slower than the recorded times to completion using the original Beckers et al strategy (8 h 25 m and 6 h 35 m) although this may not be a representative observation, given the typical variability seen in these experiments.

On close examination of the recordings, the reason for obtaining a single tight cluster became clear. It is simply a function of the relative ease of leaving a frisbee on the cluster, compared to the difficulty of taking one off. A frisbee left hard against a cluster by the 'ring' behaviour has a relatively small probability of being removed by a robot under Rule 3, as has been noted previously, but may be removed to the pullback distance by a 'plain' behaviour under Rule 2. A frisbee left at the pullback distance from a cluster by the 'plain' behaviour is equally likely subsequently to be pushed onto the cluster by a robot using the 'ring' behaviour, or pushed onto the cluster and then pulled back by a 'plain' behaviour. Since the probabilities of a 'ring' or 'plain' behaviour being applied are equal at 0.5, all that is necessary to achieve tight clustering is that the probability of being removed from a cluster by a 'plain' behaviour should be less than the probability of being added to a cluster by a 'ring' behaviour. This is clearly the case, and so tight clustering is to be expected.

6. Discussion

The work reported here has shown that both the segregation and also the crude annular sorting of two types of object differing only in colour can be achieved by a system of simple mobile robots which can sense only the colour of the object they are carrying, and have no capacity for spatial orientation or memory. This shows that this sorting problem can be solved by agents simpler than was previously supposed possible, and raises some questions:

- What is the role of stigmergy and self-organisation in the segregation and sorting tasks?
- Could the results obtained here have any significance for the understanding of stigmergy in social insect behaviour?
- Why do robotic experiments involving stigmergy appear to yield simpler solutions than abstract computer simulations?

These questions will be dealt with in turn.

6.1 What is the role of stigmergy and self-organisation in the segregation and sorting tasks?

The segregation and sorting seen in these experiments can be divided into two aspects:

- (i) The progression from small clusters to larger clusters, and eventually to a single cluster
- (ii) The differentiation of object types within clusters

We believe that these proceed more or less independently, and involve different processes of self-organisation. However, the stigmergic constituents underlying both processes are similar. The aspect of acting on the environment is combinatorially quite complex; it is worth running through a simplified version of the main possibilities to appreciate this. A robot may be in any of three states: pushing a plain frisbee, pushing a ring frisbee, and pushing no frisbee. At a given position, it may encounter any of seven basic environmental states: the boundary, another robot, an isolated plain frisbee, a plain frisbee in contact with the boundary or another frisbee, a ring frisbee, a ring frisbee in contact with the boundary or another frisbee, or no frisbee. (By 'in contact' we mean that the frisbee is in contact with the specified obstacle so that it pushes against it, and therefore triggers the gripper, when struck by the robot moving in the current direction. It is also assumed that, in all encounters with frisbees when the gripper is empty, the encountered frisbee lodges fully in the gripper.) The outcome of each of the twenty-one interactions depends on the algorithm; the table below shows the outcomes for the basic pullback algorithm, with CD indicating a change of direction.

type of frisbee carried	object in path	outcome
ring	boundary	no change/CD
ring	robot	no change/CD
ring	ring (isolated)	ring dropped/CD
ring	ring (in contact)	ring dropped/CD
ring	plain (isolated)	ring dropped/CD
ring	plain (in contact)	ring dropped/CD
ring	none	no change
plain	boundary	no change/CD
plain	robot	no change/CD
plain	ring (isolated)	plain dropped (pullback)/CD
plain	ring (in contact)	plain dropped (pullback)/CD
plain	plain (isolated)	plain dropped (pullback)/CD
plain	plain (in contact)	plain dropped (pullback)/CD
plain	none	no change
none	boundary	no change/CD
none	robot	no change/CD
none	ring (isolated)	ring picked up
none	ring (in contact)	no change/CD
none	plain (isolated)	plain picked up
none	plain (in contact)	plain dropped (pullback)/CD
none	none	no change

The aspect of movement through the environment is rather simpler. After every encounter marked CD above, a robot makes a random turn (between some limits) and sets off along a roughly straight course. In an empty arena of a known shape it would be possible to calculate the distribution of trajectories of a single robot; however, with multiple robots in an arena containing many frisbees, it would probably be difficult to improve on the easy first assumption that trajectories are an independent succession of straight lines from random points in random directions.

It is comparatively easy to see how clustering emerges from the interaction of these two aspects. Any isolated ring frisbee will eventually be picked up and deposited against another frisbee; no ring frisbee in contact with another frisbee can be picked up. Isolated plain frisbees will be dropped near any frisbee which is struck, and can be picked up and moved the pullback distance from any frisbee with which they are in contact. We can thus expect to see the formation of small clusters, which may be diffuse because of the presence of plain frisbees; the diffuseness will increase with increasing pullback distance. However, a robot which is pushing a frisbee towards a cluster will leave it on or near the cluster unless the straight-line path as wide as the robot passes cleanly through the cluster; this is impossible if the cluster is dense, and increasingly unlikely as cluster size increases, even if it is quite diffuse. Other things being equal, the probability that a robot pushing a frisbee on a succession of random straight-line courses collides with a given cluster will be a monotonically increasing function of the diameter of the cluster. Since such a collision will result in the frisbee being left on or near the cluster, the expected rate of acquiring such frisbees will also be a monotonically increasing function of the diameter.

A robot which is not pushing a frisbee is potentially capable of removing a frisbee from the cluster. For this to be possible, the robot's random straight-line trajectory must collide with a frisbee which is not 'in contact'. It is clear that, in general, such trajectories will have to intersect the periphery of the cluster; it is also clear that, other things being equal, smaller or more diffuse clusters will be more vulnerable than larger or more dense ones. It is therefore likely that the expected rate of losing such frisbees will be a monotonically decreasing function of the cluster diameter.

For a given cluster, the expected net rate of change in the number of frisbees of which it is composed will be ([expected rate of acquisition] - [expected rate of loss]). This function can be expected to increase monotonically with cluster diameter. However, since the number of frisbees is constant, the sum of all growth rates over all clusters must be zero, and so the largest cluster must have a positive expected rate of growth, and the smallest cluster must have a negative expected rate of growth. The smallest cluster will therefore always tend to lose frisbees, and will eventually disappear; the new smallest cluster will also then have a negative expected rate of growth, and will also tend to disappear, and this process will continue until there is a single cluster remaining.

This simple argument explains why the Beckers et al algorithm achieves clustering without requiring the modulation of the probabilities of dropping and picking up objects as a function of the sensed local density of the objects: the random distribution of straight-line trajectories, which is itself produced by interactions, interacts with the cluster diameter, which is an expression of local density, to modulate the rate of collisions, which directly determines the rates of dropping and picking up objects. If, instead of pushing objects against each other on a planar surface, the robots stacked and destacked them in vertical piles, then this process would not operate, and some means of sensing local density (height of stacks) would then be necessary.

The mechanism of brood sorting in ants was suspected by Franks to be variant of self-sorting or self-sieving: "(In sand)...self-sorting occurs under the influence of shaking and gravity as the smaller particles are able to move down the crevices that lie between the larger particles (see Barker and Grimson 1990)...A similar process may occur when the ants sort their brood. Essentially the tendency of the ants to cluster their brood provides a centripetal force that serves instead of gravity." [10]. The shaking is clearly provided by the random movements of the ants. Franks sees the brood items already in place as exerting some influence on the space around them (possibly a pheromone) which affects the distance at which other brood items will be dropped. We believe that a similar type of self-sorting mechanism is responsible for the distribution of the different types of objects in the experiments reported here, but that it differs from Franks' proposal in important ways.

In the algorithm presented here, we can identify the factors which produce self-sorting. The question of the nature of the centripetal force is particularly interesting, because it can be argued that no such force is in fact necessary; all that is required is that there should be a tendency for items to be moved towards the centre of the cluster. In our experiments, this tendency is provided by the geometry of the situation and by the movement characteristics of the robots. If we consider a circular cluster, then any robot approaching the cluster must do so from outside the cluster; since robots travel in approximately straight lines, there will always be a positive resolved component in the direction of the centre of the cluster, and so any frisbee brought or moved by the robot will necessarily be moving towards the centre of the cluster. If the outside of the cluster consists of diffusely spaced plain frisbees, then robots pushing frisbees are likely to be able to penetrate a certain distance within the cluster before striking another frisbee. Ring frisbees will be left at the site of this collision; however, plain frisbees will be pulled back, retracing the incoming trajectory for the pullback distance before being dropped. Ring frisbees will therefore tend to be deposited closer to the centre of the cluster than will plain frisbees.

If the incoming robot is not pushing a frisbee, then it may strike either a ring frisbee or a plain frisbee. Any such frisbee not in contact with another frisbee will be picked up and pushed further into the cluster (since the robot does not change direction). If it is a ring frisbee, it will be left against the next frisbee to be encountered. However, if it is a plain frisbee, it will be pulled back along the incoming trajectory - in other words,

along a path which has a positive resolved component away from the centre of the cluster - for the pullback distance before being dropped; it can end up further out than where it started from. Ring frisbees can thus only be moved towards the centre of the cluster, whereas plain frisbees are moved inwards by the same distance, and then outwards by the pullback distance.

The root of the self-sorting mechanism in the pullback algorithm is clearly nothing to do with ring frisbees being able to penetrate into spaces inaccessible to plain frisbees; both types of frisbee can penetrate equally close to the centre of a cluster, but only the ring frisbees can remain there, while the plain frisbees are forced to move away. Under the constant bombardment of incoming robots on randomly determined trajectories, the frisbees naturally self-organise into a core of rings and an outer halo of plains.

We therefore conclude that the sorting and segregation produced by the pullback algorithm are the outcome of two distinct processes of self-organisation, both of which are mediated by stigmergy.

6.2 Could the results obtained here have any significance for the understanding of stigmergy in social insect behaviour?

The elements of the technique used here can certainly be found in the social insect literature. For example, an account of *Leptothorax* building behaviour [10] mentions the possible use of an increased resistance to pushing a building block forwards against other building blocks as the cue to drop it: "...if the ants drop their granule only if they meet sufficient resistance...". He also reports a behaviour which is reminiscent of the pullback algorithm: "...workers individually carry granules into the nest. They walk head first towards the cluster of their nestmates, who are already installed in the nest, forming a fairly tight group. After coming close to the group of ants, the builder then turns through 180° to face outwards from the nest. The worker then actively pushes the granule it is carrying into other granules already in the nest or, after a short time, if no other granules are encountered it simply drops its load."

It is therefore fair to say that the pullback algorithm contains elements which are found in ant behaviours, that it implements a natural sorting process of a type which has previously been thought to be used by ants for brood sorting, and that it requires less sensory information than any suggested alternative mechanism. We cannot ourselves establish whether it is in fact used by ants, but perhaps it should be added to the list of candidates for study. It must be admitted that it does not always produce very good segregation or sorting; on the other hand, the adaptive advantages of brood sorting as set out by Franks and Sendova-Franks [10] are so strong that a variant of this simple method may well have been evolved at some stage, to be later superseded by a more complex but better solution.

6.3 Why do robotic experiments involving stigmergy appear to yield simpler solutions than abstract computer simulations?

If we are to understand stigmergy, we should attempt to establish not just how it works, but why it works. It is essentially an exploitation of physics via behaviour: and it seems very likely that the richer the physics, the simpler can be the behaviour. Abstract simulations, especially of gridworlds, are attractively simple from the conceptual point of view, but they have two severe disadvantages: they can skate over serious problems of sensing and actuation, and they have extremely impoverished physical characteristics. Using robots to investigate stigmergy certainly solves the problems of impoverished physics - it gives access to all the physics of the real world, for good or ill - but its practical difficulties should not be underestimated. Very slight variations in the sensors, actuators, or environmental conditions can produce gross changes in the outcome of a stigmergically mediated experiment, to the extent that it is difficult to identify the factors active in determining the result without conducting large numbers of different trials, which soon becomes prohibitively expensive in time and other resources. Nevertheless we believe that the nature and sensitivity of stigmergic operation, and the huge number of potential access points for evolution, can only really be appreciated by using robots, and that studying even the simplest of tasks, such as those described in this paper, has much to contribute towards our knowledge of stigmergy.

Acknowledgments

The authors are grateful to Nigel Franks and Ana Sendova-Franks for their advice and assistance, and to Ian Horsfield and Jason Welsby for hundreds of hours of robot-minding.

References

- [1] Barker G. and Grimson M. (1990), 'The physics of muesli', *New Scientist* 126 (1718): 37-40
- [2] Beckers R. Holland O.E. and Deneubourg J.L. (1994), 'From local actions to global tasks: Stigmergy and collective robotics', in Brooks R. and Maes P. *Artificial Life IV*, MIT Press
- [3] Bonabeau E., Theraulaz G., Fourcassié V., Deneubourg J-L. (1998), 'The phase-ordering kinetics of cemetery organisation in ants', SFI Working Paper #98-01-008
- [4] Bonabeau E., Theraulaz G., Deneubourg J-L., and Camazine S. (1997) 'Self-organisation in social insects', *TREE* 12:5
- [5] Chrétien L., (1996) 'Organisation spatiale du matériel provenant de l'excavation du nid chez *Messor barbarus* et des cadavres d'ouvrières chez *Lasius niger* (Hymenopterae: Formicidae)' Ph.D. Thesis, Département de Biologie Animale, Université Libre de Bruxelles
- [6] Cao Y.U. Fukunaga A.S. Kahng A.B. and Meng F. (1995) 'Cooperative mobile robotics: Antecedents and directions', *IEEE International Conference on Intelligent Robots and Systems*, IEEE

- [7] Deneubourg J-L. and Franks N.R. (1995) 'Collective control without explicit coding: the case of communal nest excavation', *Journal of Insect Behaviour*, Vol 8, No. 4
- [8] Deneubourg J.L. Goss S. Franks N. Sendova-Franks A. Detrain C. and Chrétien L. (1991), 'The dynamics of collective sorting: Robot-like ants and ant-like robots', in Meyer J.A. and Wilson S.W. (eds): *Simulation of Adaptive Behaviour; From Animals to Animats*, MIT Press, Cambridge
- [9] Franks N.R. and Deneubourg J-L. (1997) 'Self-organising nest construction in ants: individual worker behaviour and the nest's dynamics', *Anim. Behav.* 1997, 54, 779-796
- [10] Franks N.R. and Sendova-Franks A.B. (1992), 'Brood sorting by ants: distributing the workload over the work-surface', *Behav. Ecol. Sociobiol.* 30: 109-123
- [11] Franks N.R. Wilby A. Silverman B.W. and Tofts C., (1992), 'Self-organising nest construction in ants: sophisticated building by blind bulldozing' *Animal Behaviour*, 44, 357-375
- [12] Grassé, P.P. (1959) *La reconstruction du nid et les coordinations inter-individuelles chez Bellicositermes natalensis et Cubitermes sp. La theorie de la stigmergie: Essai d'interpretation des termites constructeurs.* *Ins. Soc.*, 6, 41-83
- [13] Holland O.E. (1996a), 'Multi-agent systems: lessons from social insects and collective robotics', *AAAI Spring Symposium on Adaptation, Coevolution and Learning in Multiagent Systems*, Stanford
- [14] Holland O.E. (1996b) 'The use of social insect control methods with multiple mobile robots.' *IEE Conference on Control '96*, Exeter
- [15] Maris M. and te Boekhurst R. (1996), 'Exploiting physical constraints: heap formation through behavioural error in a group of robots' (publisher TBA)
- [16] Theraulaz G., Bonabeau E., and Deneubourg J-L. (1998), 'The origin of nest complexity in social insects', *Complexity* 3:6