

Stigmergy, Self-Organization, and Sorting in Collective Robotics

Abstract Many structures built by social insects are the outcome of a process of self-organization, in which the repeated actions of the insects interact over time with the changing physical environment to produce a characteristic end state. A major mediating factor is stigmergy, the elicitation of specific environment-changing behaviors by the sensory effects of local environmental changes produced by previous behavior. A typical task involving stigmergic self-organization is brood sorting: Many ant species sort their brood so that items at similar stages of development are grouped together and separated from items at different stages of development. This article examines the operation of stigmergy and self-organization in a homogeneous group of physical robots, in the context of the task of clustering and sorting Frisbees of two different types. Using a behavioral rule set simpler than any yet proposed for ant sorting, and having no capacity for spatial orientation or memory, the robots are able to achieve effective clustering and sorting showing all the signs of self-organization. It is argued that the success of this demonstration is crucially dependent on the exploitation of real-world physics, and that the use of simulation alone to investigate stigmergy may fail to reveal its power as an evolutionary option for collective life forms.

Owen Holland*
Chris Melhuish
Intelligent Autonomous
Systems Engineering
Laboratory
Faculty of Engineering
University of the West of
England
Coldharbour Lane
Frenchay
Bristol BS16 1QY
United Kingdom
oholland@cyberlife-institute.
co.uk
chris.melhuish@uwe.ac.uk

Keywords
collective robotics, stigmergy, self-organization

1 Introduction: Stigmergy

Stigmergy is a concept occasionally used in biology to describe the influence on behavior of the persisting environmental effects of previous behavior. It was originally proposed by Grassé [12] to explain some of his observations on termite building behavior. 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 tâches, la régulation des constructions ne dépendent pas directement des ouvriers, mais des constructions elles-mêmes. *L'ouvrier ne dirige pas son travail, il est guidé par lui.* C'est à cette stimulation d'un type

* Now at The CyberLife Institute, CyberLife Technology Ltd., Quayside, Bridge Street, Cambridge CB5 8AB, UK.

particulier que nous donnons le nom du STIGMERGIE (*stigma*, piqûre; *ergon*, travail, oeuvre = oeuvre stimulante). (p. 65)

[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*, wound from a pointed object; *ergon*, work, product of labor = stimulating product of labor).]

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*” ([12], p. 79).

If stigmergy is indeed at the root of the building behavior 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 organized and coordinated behaviors and behavioral 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 behavior produced (and the work done) at that location in the future. If a drop of pheromone was deposited instead, it could also change the behavior 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 behavior 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 behavior of the relevant agent to be appropriately affected by previous environmental changes, and this can happen in a number of distinct ways:

1. The agent’s choice of action may be affected (a qualitative effect).
2. 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).

Method 1 captures Grassé’s sense of action being guided, and Method 2 also includes the element of intensity of activity. The qualitative effect in Method 1 may of course be internally controlled by some threshold mechanism acting on a quantitatively varying input. However, there is also a third possibility that is not included in Grassé’s formulation:

3. 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 that 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, the driver 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 1 and 2 may be thought of as active [13] in that they affect the agent itself. Passive stigmergy is very close to purely physical situations in which some constantly acting force—often a fluid—changes the environment in such a way as to change its future effect on the environment; 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 behaviors. Stigmergy is essentially a mechanism that 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. [4], any structure emerging from this repeated interaction develops by a process of self-organization (SO) (see also [7, 9, 16]). Bonabeau and his colleagues [3] have provided a useful brief summary of the nature and properties of SO. They define and describe SO as

a set of dynamical mechanisms whereby 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. (p. 189)

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 that 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 behavior-mediated changes that has been exploited by evolution to produce the striking phenomena found in social insect colonies; Bonabeau et al. [4] have pointed out some of the possible ways in which evolution may favor the emergence of some aspects of self-organization.

There are several other comments that may usefully be made about stigmergy and that can increase our understanding. One approach to stigmergy is simply to consider the minimal qualities of agent and environment that 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' behavior. (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 that 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 behavioral 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 that 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 subtasks 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 subtask itself, it may be necessary to augment the subtask to provide some additional external cue, or sign.) In particular, where there are many similar cues for a certain subtask at a given location, the rate of performance of the subtask will be a function of the number of agents at that location [14]. (This would not necessarily be the case if an agent had to be in a particular internal state to be able to respond to the cue.) If there are many locations with such cues, the subtask will be performed fastest at the locations that 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 that do not control agent distribution were identified in [15], a computer simulation of a task inspired by the building behavior of wasps.

Some of the most useful insights into stigmergy have been provided by simulations. For example, in their paper titled "The dynamics of collective sorting: Robot-like ants and ant-like robots," Deneubourg et al. [8] presented a simulation showing that simple agents, specified in terms that 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 [6]) 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 they 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 that were unable to detect whether or not they were moving any objects, that had no memory, and that 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-organization. Small clusters were formed at first through the action of the threshold mechanism; by random accretion and loss, 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 article describes and discusses several collective robot experiments aimed at extending the approach of Beckers et al. [2] to examine the possible role of stigmergy and self-organization in the simple task of sorting two types of objects, as studied in simulation in [8]. The article is organized 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-organization.

2 Spatial Sorting

Spatial sorting is an informal notion; to structure this investigation, some formalization 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 that 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 recognize that they may arise but find it more appropriate in these early investigations to defer considerations until the distinctions that 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 summarized by Deneubourg et al. [8]:

[T]he eggs are arranged in a pile next to a pile of larvae and a further pile of cocoons, or else the three 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 different piles as before. (p. 356)

However, sorting need not be into piles; Franks and Sendova-Franks [10] describe the brood sorting behavior 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 et al.'s model [8] certainly captures many features of the behavior 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" (p. 120) and speculate 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 Boekhorst [15]. Most accounts in the literature deal with foraging tasks where the objects must be brought back to a designated location that is found by some means of navigation; see Cao et al. [5] 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 behavior, 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 15) in cracks in rock; this constrained environment means that they are behaviorally adapted to life in two dimensions. It is therefore possible to study the behavior of a colony by providing it with a particular two-dimensional habitat—the space between two glass slides—which allows unrivaled opportunities for observation and recording.

Leptothorax is also peculiarly suitable as a model for robotic investigations of some collective behaviors, 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 behaviors tend to involve the movement of single lumps of material (carborundum grit of regular dimensions is usually provided by the experimenters) that 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.

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 (23 cm 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)




Figure 1. A U-bot.

- *Maneuverability*: 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 infrared proximity sensors (three facing forward, one backward) set to around 20 cm. 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 8 cm beyond the robot's circular chassis. The Frisbees, which are 23 cm in diameter, are placed on the floor concave side up. If a robot moves toward a Frisbee that 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 that are free to swing fore and aft but cannot move sideways. As the gripper continues to move, the barbels ride over the edge of 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 op-




Figure 2. A U-bot gripping a Frisbee.

tical sensor detects its type. Several types are in use: plain yellow, plain black, black ring with white center, and white ring with black center—see, for example, Figure 3.

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 forward by the rim.) However, if it is in the lowered position, it will engage the rim, and the Frisbee will be pulled backward.




Figure 3. U-bots with Frisbees.




Figure 4. The robot arena.

The second mechanism is similar to that used by Beckers et al. [2]: The whole gripper is suspended so that a backward-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) that is octagonal in shape, with 4-m sides; the area is 1,760 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 m above the center of the arena is linked to a PC-based video capture system that 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 type; 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 article show this distortion, which overemphasizes the central area of the arena, and care must be exercised in their interpretation.)

5 Experiments and Results

5.1 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. [2] 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

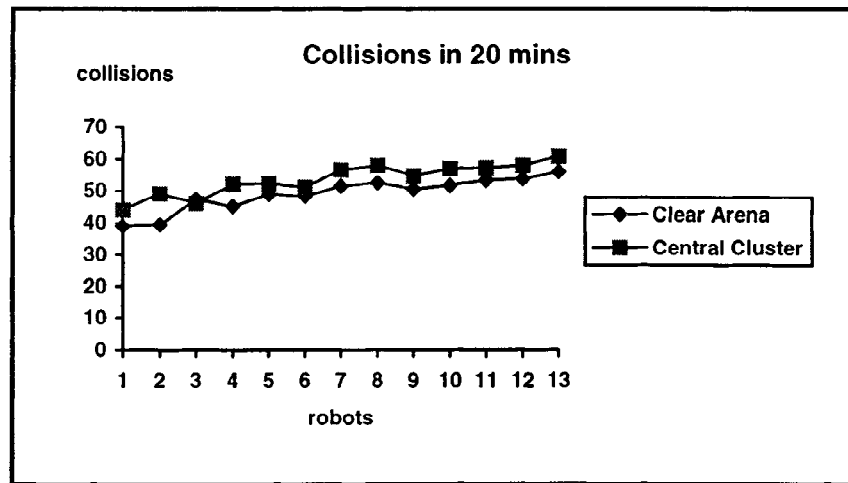


Figure 5. Average number of collisions between robots in 20 minutes with a clear arena, and with a central cluster of Frisbees.

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.

5.2 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 that 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 infrared sensors to the exact distance to enable Frisbees left on the boundary to be retrieved. The modified rule set, in decreasing order of priority, is shown in Figure 6. 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;

<p>Rule 1: if (gripper pressed & object ahead) then make random turn away from object</p> <p>Rule 2: if (gripper pressed & no object ahead) then reverse small distance (dropping the Frisbee) make random turn left or right</p> <p>Rule 3: go forward</p>

Figure 6. Adaptation of Beckers et al. [2] rule set for clustering.

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 forward 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 that is in turn touching the boundary, a U-bot that 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.

Forty-four Frisbees were placed in the arena at a uniform and regular spacing. In anticipation of the sorting work reported later, half of the Frisbees (plains) were painted black, and half were painted white with a black center (rings); at this stage, the robots were programmed to treat both alike. Ten 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. [2], first aggregating in small clusters, then forming larger ones, and eventually, after 8 h 25 min, 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. [2] 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 that 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. (p. 5)

We therefore decided to carry out a series of experiments to see how behavior at the boundary could affect the clustering process.

5.3 Experiment 3: Algorithmic Boundary Effects

While modifying the Beckers et al. algorithm to cope with the rigid boundary, we had noticed that setting the infrared detection distance too short would occasionally cause the robots to deposit Frisbees at the boundary. This guided us toward 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 that 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.

<p>Rule 1: if (gripper pressed & object ahead) then</p> <p style="padding-left: 40px;"><i>either</i> (with prob p) make random turn away from object</p> <p style="padding-left: 40px;"><i>or</i> (with prob $[1-p]$) reverse small distance (dropping the Frisbee) make random turn left or right</p> <p>Rule 2: if (gripper pressed & no object ahead) then reverse small distance (dropping the Frisbee) make random turn left or right</p> <p>Rule 3: go forward</p>

Figure 8. Rule set for probabilistic action selection.

Forty-four Frisbees were set out in a regular pattern. The overhead camera captured a frame every 5 minutes. Ten 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 Frisbees, and was stable up to 11 h 20 min. Probabilities of $p = 0.9$ and 0.95 produced two main central clusters after 5 h and 2.5 h, 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. Of the extra trials, three also formed a single cluster at the boundary (after 9 h 25 min, 10 h 35 min and 13 h 10 min) and two formed central clusters, one with 40 Frisbees (6 h 0 min) and the other with 35 Frisbees, with 9 strewn around the periphery (7 h 35 min).

This series of experiments is particularly interesting, for two reasons. First, it displays all the characteristic signatures of self-organization identified in [4]: 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

Table 1. Final states of Experiment 3 as a function of p values.

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

(a) $p = 1.0$ (b) $p = 0.95$ (c) $p = 0.9$ (d) $p = 0.88$ Figure 9. Final states of Experiment 3 as a function of p values.




Figure 9. Continued.

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 that 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-organization.

5.4 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 infrared collision avoidance sensors. To reduce the acceptance angle

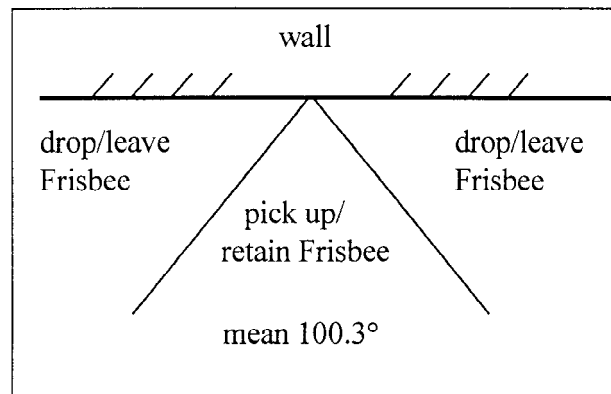


Figure 10. Arrangement of infrared collision sensor for Experiment 4.

(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 10 robots used this gave a mean sensing range of 1.08 Frisbee diameters (the distance from the center eye to the wall, with the robot orthogonal to the wall, was 0.88 Frisbee 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” behavior) 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” behavior) will be triggered instead.

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 Frisbees clustered against the wall in a single cluster was achieved; the times to completion were 10 h 15 min, 13 h 30 min, and 14 h 25 min. Figure 11 shows the final frame of the first trial.

It has therefore been possible to produce the same stigmergically self-organized 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 emphasizes a principle often articulated within robotics: Behavior is the outcome of the interaction between the sensors, the computation, the actuators, the body, and the environment. Since stigmergic self-organization typically involves the interaction of all these factors over extended time frames, it is probably even more sensitive than single behaviors 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-organized outcome—it is not constrained to act only on the computational or algorithmic aspect of the agent.

5.5 Experiment 5: The Pull-back 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




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

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 backward 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 are 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).

The U-bot's behavior 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 that is already on a cluster (perhaps at the boundary), the new version of Rule 2 will cause it to drag the plain backward and leave it the pull-back 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 infrared 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 forward under Rule 3. (Of course, the rear sensor cannot detect Frisbees, so any Frisbees behind the robot during a pull-back movement are simply pushed out of the way.)

Twenty-two ring and 22 plain Frisbees were uniformly spaced throughout the environment (Figure 13); six robots were used, with the pull-back distance being set at 2.6 Frisbee diameters.

Figure 14 shows the distribution of Frisbees after 7 h 35 min, 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 that can be regarded as sorting of some

```

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 (dropping the Frisbee)
  make random turn left or right

Rule 3:
go forward

```

Figure 12. Rule set for the pullback algorithm.

kind; within the classification proposed in the introduction, it is at least segregation, and inclines toward 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.

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 h 5 min




Figure 13. The starting configuration for Experiment 5.




Figure 14. Experiment 5 after 7 h 35 min.

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 h 45 min.

To check that this was not a fluke, we repeated the experiment four 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.

We decided to measure the time at which 20 rings were first in the same cluster. Table 2 shows these times for all five 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 h 20 min) being more than nine times the fastest (2 h




Figure 15. Experiment 5 after 8 h 5 min.




Figure 16. Experiment 5 after 1 h 45 min.

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

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

45 min). 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.

5.6 Experiment 6: Varying the Pull-back 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 pull-back 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 h 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 17 shows the (average) number of plains in the final cluster for each trial (including data from Experiment 3); Figure 18 shows the (average) times to completion.

The paucity of data points at some values makes interpretation difficult, but it may be reasonable to summarize the data as follows:

- for pull-back 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 pull-back distance to a minimum near 5.2, and then increases.

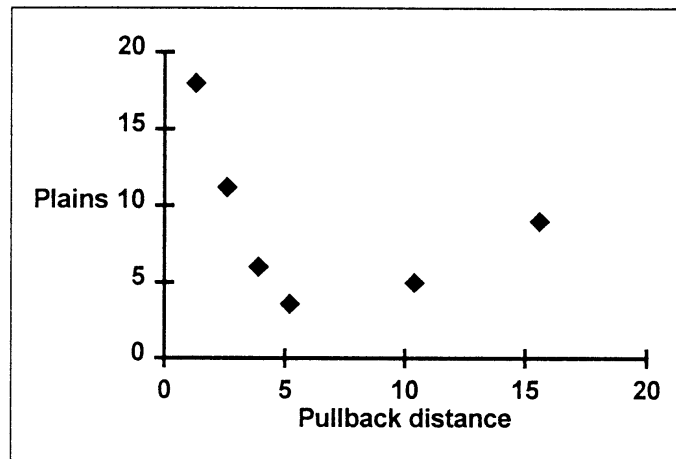


Figure 17. Average number of plains in cluster versus pull-back distance.

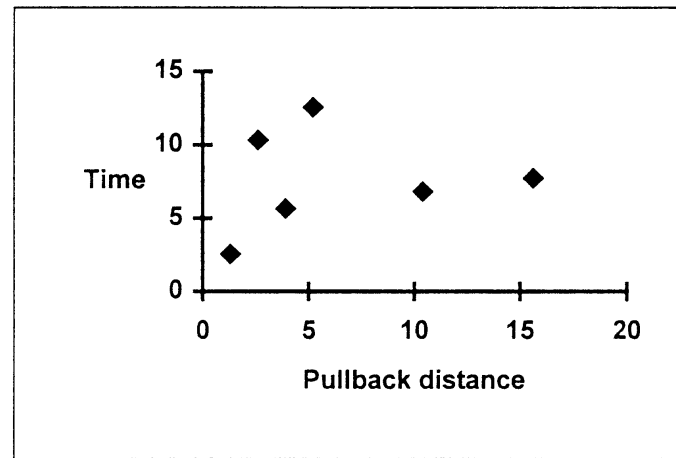


Figure 18. Average time to completion versus pull-back distance.

- for the pull-back distance at the minimum, the time to completion can vary considerably.

A key additional factor is the spatial distribution of the plains that 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 pull-back 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.




Figure 19. Good segregation at pull-back distance of 5.2 Frisbee diameters.

How should we judge the ability of the simple pull-back algorithm to produce sorting and segregation? It is clear that the algorithm can produce good segregation when the pull-back 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 center 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.

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 pull-back distance by the type of Frisbee, it was determined randomly? All these were examined in a series of experiments.

5.7 Experiment 7: Does the Clustering of Plains Improve with Time?

Twenty-two ring Frisbees were placed in a tight cluster in the center of the arena, and 22 plains were uniformly distributed over the remaining space. Six robots with a pull-back 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 center of the arena were measured. After 5 h 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 h due to external factors. The results are shown in Figure 20.

Although there are too few data points to make a firm decision, it is clear that any reduction in mean measured distance from the center 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.

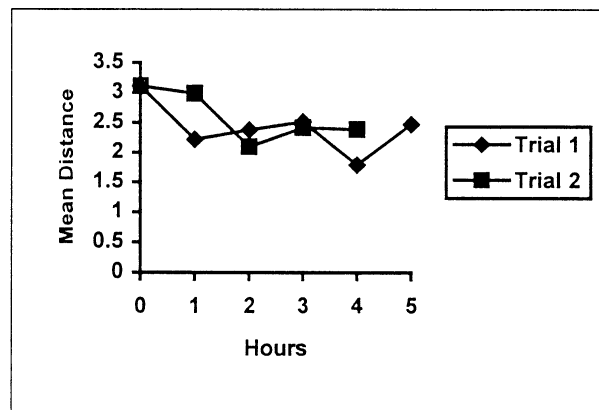


Figure 20. Mean distance of plains from cluster center versus time.

5.8 Experiment 8: Will Plains Cluster in the Absence of Rings?

Twenty-two plain Frisbees were uniformly spaced throughout the arena, and six robots with the pull-back distance set to 2.6 Frisbee diameters were released. After 6 h 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. 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 h, and for the remainder of the time it behaved exactly like the first experiment.

Since the pull-back 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 anchor spatially the dynamic cluster of plain Frisbees.
- for pull-back distances greater than 2.6, the arena may be too small to enable any clustering of plain Frisbees to be seen.

5.9 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.

Once again, 22 Frisbees of each type were uniformly spaced throughout the arena, and six robots were released. To begin with, matters seemed to be progressing as in

<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 (dropping the Frisbee) make random turn left or right</p> <p>Rule 3: go forward</p>

Figure 21. Rule set for Experiment 9.

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 h; no single cluster was ever formed, but mixed aggregations of plains and rings constantly formed and dispersed close to the periphery.

5.10 Experiment 10: Applying the Pull-back Algorithm Randomly to Rings and Plains

To see the effects of applying the pull-back 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; six robots were used, with the pull-back 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 h 15 min. This is much slower than the recorded times to completion using the original Beckers et al. [2] strategy (8 h 25 min and 6 h 35 min) 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” behavior has a very small probability of being removed by a robot under Rule 3, as has been noted previously, but may be removed to the pull-back distance by a “plain” behavior under Rule 2, provided it is perfectly aligned with the gripper of the incoming robot, which will happen only occasionally. A Frisbee left at the pull-back distance from a cluster by the “plain” behavior is equally likely subsequently to be pushed onto the cluster by a robot using the “ring” behavior, or pushed onto the cluster and then pulled back by a “plain” behavior. Since the probabilities of a “ring” or “plain” behavior 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” behavior

should be less than the probability of being added to a cluster by a “ring” behavior. 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 visual appearance can be achieved by a system of simple mobile robots that can sense only the type of 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-organization in the clustering and sorting tasks?
- Could the results obtained here have any significance for the understanding of stigmergy and self-organization in social insect behavior?
- Why do robotic experiments involving stigmergy and self-organization 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-Organization in the Clustering and Sorting Tasks?

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

1. the progression from small clusters to larger clusters, and eventually to a single cluster
2. the differentiation of object types within clusters

We believe that these proceed more or less independently and involve different processes of self-organization. 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, which are mainly determined by previous behavior in the environment: 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. Let us also assume 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 21 interactions depends on the algorithm; Table 3 below shows the outcomes for the basic pull-back algorithm, with CD indicating a change of direction. Such tables define the stigmergic substrate of self-organization.

The aspect of movement through the environment is rather simpler. After every encounter marked CD in Table 3, 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

Table 3. Outcomes for the basic pull-back algorithm.

Type of Frisbee carried	Object in path	Outcome
ring	boundary	no change/CD ^a
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 (pull-back)/CD
plain	ring (in contact)	plain dropped (pull-back)/CD
plain	plain (isolated)	plain dropped (pull-back)/CD
plain	plain (in contact)	plain dropped (pull-back)/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 (pull-back)/CD
none	none	no change

^aCD indicates change of direction.

to improve on the easy first assumption that trajectories are an independent succession of straight lines from random points in random directions. (Note that this amounts to ignoring any stigmergic influence on movement.)

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 that is struck, and can be picked up and moved the pull-back 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 pull-back distance. However, a robot that is pushing a Frisbee toward a cluster will leave it on or near the cluster unless there is a gap in the cluster wide enough for the robot and Frisbee to pass right through. This will be impossible if the cluster is dense and becomes increasingly unlikely for diffuse clusters as cluster size increases. 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 that 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 that 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. [2] 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 [10] to be a 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. (p. 120)

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 that 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 toward the center 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 center of the cluster, and so any Frisbee brought or moved by the robot will necessarily be moving toward the center of the cluster. If, because of the outcome of previous actions, 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. If there is a dense central core, they can penetrate no further. Ring Frisbees will be left at the site of this collision; however, plain Frisbees will be pulled back, retracing the incoming trajectory for the pull-back distance before being dropped. Ring Frisbees will therefore tend to be deposited closer to the center of the cluster than will plain Frisbees and may be added to any dense core.

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 that has a positive resolved component away from the center of the cluster—for the pull-back distance before being dropped; it can end up further out than where it started from. Ring Frisbees can thus only be moved toward the center of the cluster, but no further than the outer limit of any dense core, whereas plain Frisbees are moved inward by the same distance, and then outward by the pull-back distance.

The root of the self-sorting mechanism in the pull-back 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 center 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-organize into a dense core of rings and an outer halo of plains. We therefore conclude that the clustering and sorting produced by the pull-back algorithm are the outcome of two distinct processes of self-organization, both of which are mediated by stigmergy.

6.2 Could the Results Obtained Here Have Any Significance for the Understanding of Stigmergy and Self-Organization in Social Insect Behavior?

The elements of the technique used here can certainly be found in the social insect literature. For example, an account of *Leptothorax* building behavior [11] mentions the possible use of an increased resistance to pushing a building block forward against other building blocks as the cue to drop it: “[I]f the ants drop their granule only if they meet sufficient resistance . . .” (p. 374). Franks et al. also report a behavior that is reminiscent of the pull-back algorithm:

[W]orkers 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. (p. 361)

It is therefore fair to say that the pull-back algorithm contains elements that are found in ant behaviors, that it implements a natural sorting process of a type that 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 and Self-Organization 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 behavior, and it seems

very likely that the richer the physics, the simpler can be the behavior. 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 and self-organization 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 article, has much to contribute toward 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., & Grimson, M. (1990). The physics of muesli. *New Scientist*, 126 (1718), 37–40.
2. Beckers, R., Holland, O. E., & Deneubourg, J.-L. (1994). From local actions to global tasks: Stigmergy and collective robotics. In R. Brooks & P. Maes (Eds.), *Artificial life IV* (pp. 181–189). Cambridge, MA: MIT Press.
3. Bonabeau, E., Theraulaz, G., Deneubourg, J.-L., & Camazine, S. (1997). Self-organisation in social insects. *Trends in Ecology and Evolution*, 12(5), 188–193.
4. 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). Santa Fe, NM.
5. Cao, Y. U., Fukunaga, A. S., Kahng, A. B., & Meng, F. (1995). Cooperative mobile robotics: Antecedents and directions. In *Proceedings of IROS95 IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 226–234). New York: IEEE.
6. 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)*. Unpublished doctoral dissertation, Département de Biologie Animale, Université Libre de Bruxelles.
7. Deneubourg, J.-L. & Franks, N R. (1995). Collective control without explicit coding: The case of communal nest excavation. *Journal of Insect Behaviour*, 8(4), 417–432.
8. Deneubourg, J.-L., Goss, S., Franks, N., Sendova-Franks, A., Detrain, C., & Chrétien, L. (1991). The dynamics of collective sorting: Robot-like ants and ant-like robots. In J. A. Meyer & S. W. Wilson (Eds.), *Proceedings of the First International Conference on the Simulation of Adaptive Behavior: From Animals to Animats* (pp. 356–363). Cambridge, MA: MIT Press.
9. Franks, N. R. & Deneubourg, J.-L. (1997). Self-organising nest construction in ants: Individual worker behaviour and the nest's dynamics. *Animal Behaviour*, 54, 779–796.
10. Franks, N. R. & Sendova-Franks, A. B. (1992). Brood sorting by ants: Distributing the workload over the work-surface. *Behavioral Ecology and Sociobiology*, 30, 109–123.
11. Franks, N. R., Wilby, A., Silverman, B. W., & 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 théorie de la stigmergie: Essai d'interprétation des termites constructeurs. *Insectes Sociaux*, 6, 41–83.
13. Holland, O. E. (1996). Multi-agent systems: Lessons from social insects and collective robotics. In *Adaptation, Coevolution and Learning in Multiagent Systems: Papers from the 1996 AAAI Spring Symposium* (pp. 57–62). Tech. Rep. 55-96-01). Menlo Park, CA: AAAI Press.
14. Holland, O. E. (1996). The use of social insect control methods with multiple mobile robots. In *UKACC International Conference on Control* (pp. 1260–1265). (Conf. Pub. No. 427). London: Institution of Electrical Engineers.
15. Maris, M., & te Boekhorst, R. (1996). Exploiting physical constraints: Heap formation through behavioural error in a group of robots. In *Proceedings of IROS96* (pp. 1655–1660).
16. Theraulaz, G., Bonabeau, E., & Deneubourg, J.-L. (1998). The origin of nest complexity in social insects. *Complexity*, 3(6), 15–25.