

# SWARM-BOT: Pattern Formation in a Swarm of Self-Assembling Mobile Robots

Erol Şahin<sup>1</sup>, Thomas H. Labella<sup>1</sup>, Vito Trianni<sup>1</sup>, Jean-Louis Deneubourg<sup>2</sup>, Philip Rasse<sup>2</sup>, Dario Floreano<sup>3</sup>, Luca Gambardella<sup>4</sup>, Francesco Mondada<sup>3</sup>, Stefano Nolfi<sup>5</sup>, Marco Dorigo<sup>1</sup>

<sup>1</sup>IRIDIA - Université Libre de Bruxelles, Belgium

<sup>2</sup>CENOLI - Université Libre de Bruxelles, Belgium

<sup>3</sup>ASL - Swiss Federal Institute of Technology, Lausanne, Switzerland

<sup>4</sup>IDSIA, Manno-Lugano, Switzerland

<sup>5</sup>Institute of Cognitive Sciences and Technologies - CNR, Roma, Italy

**Abstract**—In this paper we introduce a new robotic system, called *swarm-bot*. The system consists of a swarm of mobile robots with the ability to connect to/disconnect from each other to self-assemble into different kinds of structures. First, we describe our vision and the goals of the project. Then we present preliminary results on the formation of patterns obtained from a grid-world simulation of the system.

**Keywords**—swarm intelligence, ant algorithms, self-assembling, mobile robot, pattern formation

## I. INTRODUCTION

In this paper we present some of the initial results obtained within the SWARM-BOTS project, a Future and Emerging Technologies project funded by the CEC.

A *swarm-bot* is a self-assembling and self-organizing artifact composed of a number of smaller devices, called *s-bots*. S-bots can aggregate to form a swarm-bot, and a swarm-bot can split up into its s-bots components, as necessary. Beside aggregating and disaggregating, a swarm-bot can change its shape on-line and has adaptive capabilities to match environmental variability.

Self-assembly, that is, the self-organized creation of structures composed of independent entities that are autonomous in their control, occurs in a wide range of natural systems ranging from chemistry to biology (for a review see [1], [2]). A particularly interesting form of self-assembly is observed in social insects, which create different types of structures by physically attaching to each other: curtains, festoons, ovens, thermoregulatory clusters, swarms, bivouacs, bridges, doorways, pulling chains and rafts are all examples of self-assembling in social insects.

Self-assembly represents a challenging class of problems for *swarm intelligence* [3]: the study of how collectively intelligent systems can be created by a number of simple autonomous agents. One of the main interests of the SWARM-BOTS project is the design of swarm intelligent control algorithms that let a group of autonomous mobile robots self-assemble. In this paper, we focus on one of the basic components of any self-assembling system: spatial pattern formation.

## II. SWARM-BOT

As we mentioned above, a swarm-bot is a new robotic system composed of a swarm of s-bots and capable of self-assemble and self-organize to adapt to its environ-

ment. The s-bots have connectors around their body and can connect to other s-bots to create physical structures. The system as a whole can dynamically self-assemble into different structures to perform certain tasks, and then split into s-bots to perform other tasks.

The swarm-bot concept lies between two main streams of robotics research: collective robotics [4], [5], [6] and metamorphic robotics [7], [8]. In collective robotics, a group of autonomous mobile robots interact with each other to accomplish a particular task. However, unlike s-bots, they do not have the ability to attach to each other by making physical connections. In metamorphic robotics, the system consists of connected self-contained modules that, although autonomous in their movements, remain attached to each other, lacking the full mobility of s-bots.

## III. AN IMAGINARY SCENARIO

As a help to describe our long-term vision, here we provide an imaginary scenario. Four phases of this imaginary scenario where a swarm of s-bots transports a heavy object to a goal are depicted in Fig.1. The object (large hexagon) is placed on the left side. The goal is marked with a large circle on the right side of the sketch, and is not visible from the initial location of the swarm of s-bots. In the swarm there are two types of s-bots: the carrier s-bots are specialized in carrying objects and have more powerful motors than the explorer s-bots (distinguishable by their lighter bodies in Fig.1), which are equipped with additional sensors that allow them to see the goal (in Fig.1a, the view of the goal is obstructed by the walls), the obstacles and each other.

The environment consists of a rough terrain and contains obstacles. These consist of walls (dark rectangles) that neither the s-bots nor the swarm-bot can go over, and small/large holes (the small and large patches), which can be passed over by a swarm-bot, but not by a single s-bot.

In the first phase, Fig.1a, some of the carriers discover the object and grip it by its edges. After some pulling and pushing they find out that they cannot move the object in any direction and therefore they call for help, so that other carriers connect to them. At the same time, the explorers disperse in the environment in search of the goal. In the second phase, Fig.1b, the explorers discover the goal, and a path from the current location of the swarm-bot to the goal is chosen. In determin-

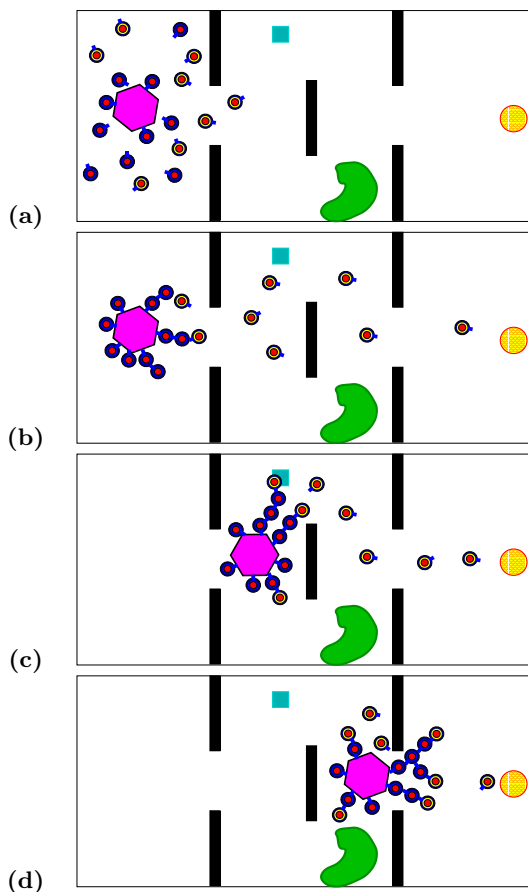


Fig. 1. Four phases of an imaginary transportation task by a swarm-bot.

ing the path, the s-bots prefer the upper route with the small hole on it to the lower one. By this time, the carriers connect to each other creating a swarm-bot that can move the object in a desired direction. Some of the explorers connect to the periphery of the swarm-bot to lead it towards the goal. Through physical connections, their pull is ‘felt’ by the swarm-bot and all the carriers align themselves to provide maximum force on the object towards the desired direction. In the third phase, Fig.1c, the swarm-bot passes across the first opening between the walls, and reaches the small hole. Note that the shape of the swarm-bot and the placement of the peripheral explorers are different from those of the second phase, due to the central wall and to the change in the direction of the movement. The swarm-bot can go over the small hole, by providing support to the s-bots who are over it. The last phase of the scenario, Fig.1d, shows how the carriers and the explorers at the periphery dynamically re-configure the swarm-bot, by disconnecting and connecting again, to be able to pass between the walls.

#### IV. GENERAL OBJECTIVES

The tasks embedded in the imaginary scenario described above can be roughly categorized into three:

- **Dynamic shape/pattern formation/change:** The s-bots should be able to self-assemble into a number

of different planar geometric configurations, like those found in ant colonies and in patterns of differential adhesion by developing cells.

The swarm-bot should have the ability to move as a coherent structure. For instance, when the swarm of s-bots is following a light, the s-bots can self-assemble into a swarm-bot, where the s-bots with light sensors form the periphery. The s-bots who lack light sensors or can’t see the light directly can align themselves with the current direction of the swarm-bot through the use of local interactions. The swarm-bot should be able to dynamically re-configure itself while passing through narrow passages, through the disconnection and re-connection of s-bots. Emphasis is given to trigger the self-assembly of these patterns/shapes by changing the environment, rather than via issuing explicit commands. For instance, the shape of a swarm-bot should be changed by changing the positioning of the lights in the environment, rather than explicitly commanding the swarm-bot to mold itself into an arbitrary shape.

- **Navigation on rough terrain:** The swarm-bot should be capable of autonomously moving on rough terrain guided by sensory information gathered by individual s-bots. The terrain might contain various types of obstacles that the swarm needs to traverse. This might require the s-bots to self-assemble into a swarm-bot, or a swarm-bot to split into smaller swarm-bots or even into s-bots. For instance, when going over holes that are larger than the body of a single s-bot, the s-bots can connect to each other to provide physical support to those s-bots suspended over the hole.

The swarm-bot should be able to move efficiently from a point to another on a rough terrain. For this task, the swarm-bot approach has the advantage, when compared to metamorphic robotic systems that cannot disband at will, that its s-bots can break-up from the swarm-bot to explore the environment for finding the optimal (or a good) route to the goal. Finding an optimal route is a difficult task since the length of the route needs to be combined with other factors such as the roughness of the terrain and the presence of obstacles in the environment. For this purpose, individual s-bots can form a dynamic network to the goal, leading the swarm-bot along the optimal route.

- **Pulling/pushing objects:** The swarm-bot should be able to pull/push objects that are too heavy for a single s-bot to carry. Such an ability is observed in the prey retrieval behavior of ants. The s-bots that want to retrieve a heavy object will call for help and, by getting connected to other s-bots, they can create a stronger swarm-bot that can accomplish the task.

#### V. SIMULATION

We have developed a simulator to study our control algorithms for the swarm-bot. Figure 2 shows three groups of s-bots in the simulation environment. Each circle represents an s-bot. Two different types of s-bots are shown on the left and right side of the figure. The color of the ring shows the type of the s-bot. The rectangular tentacle placed on the body represents the gripper. The inner circle shown on the body denotes a controllable light placed on top of the s-bot for signalling. The movement of the s-bot is modeled similarly to that of a differential

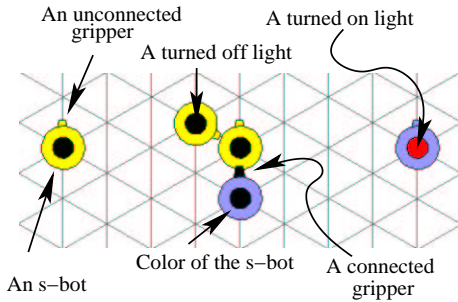


Fig. 2. Left and right side: two different types of disconnected s-bots. At the center, the lower s-bot grips the upper one. This is indicated by a gripper that is longer and darker. The s-bot on the left remains unconnected to the two, although it is within gripping distance to the upper s-bot. A black inner circle represents a turned off light (such as the one on the leftmost s-bot), and a lighter one represents a light that is on. The light of the rightmost s-bot in the figure is on, whereas the lights of the other s-bots are off.

drive mobile robot. The control of the movement is discretized to guarantee that the s-bot is always positioned on the nodes of a hexagonal grid. This simplification made it easier to implement the connection and disconnection of the s-bots.

When not gripped, the s-bots can move freely in the environment guided by their sensors. However, when gripped, an s-bot loses its mobility. For instance, the upper s-bot, shown in a group of three at the center of Fig.2, is immobilized since it is gripped by the lower s-bot. All the other s-bots in the figure are free to move since they are not gripped by other s-bots. Physical pushes are not considered in the current simulation.

The s-bots are equipped with different short-range and long-range sensing and signaling modalities. All the signalling and sensing is modeled as a light source/light sensor pair. The signal emitted by a source decreases in intensity with the square of the distance. The sources are characterized by the intensity, color, and the beam width of the signal. The top light<sup>1</sup> of the s-bot, described above, is an omnidirectional source. An s-bot type is signalled by the color of its light. To make the orientation of an s-bot visible, three other colored sources are placed uniformly around the s-bot. These are called the left, right and the rear lights of the s-bot.

The sensors are directional and can sense the color and the intensity of the signal within their viewing angle. On the s-bots, short-range sensing is achieved by six proximity sensors, uniformly distributed around the body of the s-bot. It allows the s-bots to sense the presence of other objects and robots in the immediate vicinity. The long-range sensing is achieved by six directional light sensors, placed uniformly around the s-bot. Through all these sensors an s-bot can sense the top lights when they are on, as well as the type and the orientation of other s-bots.

It is important to note that there is no explicit communication between the s-bots. The coordination of the s-bots takes place solely through their embodiment, the connections among them, and the signals emitted.

<sup>1</sup>In the rest of the article, this is referred as the *light* that is turned on or off by the s-bot.

## VI. PATTERN FORMATION

Pattern (or shape) formation is often used as a task to demonstrate the self-assembling ability of collective or metamorphic robotic systems. Below we present the preliminary results from self-assembly experiments of three patterns: chains, stripes, and center-periphery. In all the experiments, initially the s-bots are randomly dispersed in the environment. The experiments on the formation of chains is done using a single type of s-bot whereas two different types of s-bot are used in the other two experiments. In the experiments, all s-bots of a same type are controlled by the same control behavior. In the formation of patterns, the s-bots can connect to each other or stay disconnected. However, as will be shown in chain formation experiments, the connection of s-bots alter the dynamics of the pattern formation.

### A. Chain formation

Chaining requires the creation of linear connections and can be considered as one of the basic patterns of self-assembly. Although this pattern occurs in various species, a duck followed by its ducklings is the best known example. A collective robotics example of chain formation is given by [9], where a single robot chain starts and grows from a predefined and detectable location in the environment.

The chaining behavior comprises two components: a long range attraction to the light and a short range attraction to the rear light of s-bots. In the absence of any sensory stimulus an s-bot performs a random walk in the environment. In the presence of light sources, an s-bot chooses to move in a certain direction with a probability proportional to the light intensity perceived in that direction.

In the long range, the top lights are more visible and cause the unconnected s-bots to move in the direction of chains or other s-bots in the environment. As a consequence, longer chains become a more attractive target for unconnected s-bots. In the short range, the rear lights of s-bots become a stronger force of attraction. An s-bot can align itself behind another s-bot creating or lengthening a chain. However, due to the local interactions, a chain can also grow from the head. At any moment, each s-bot in a chain can disconnect with probability  $P_d$  from the chain. Note that unlike the disconnection probability, the connection probability of an unconnected s-bot is determined by a complex dynamic interplay of multiple factors, and is not easily controllable.

Three experiments are carried out to study chain formation.

- Gripped chain formation with disconnection probability  $P_d$ . In this experiment, the s-bots form chains by gripping other s-bots. Only the s-bots at the tail of the chains can disconnect with the probability given by  $P_d > 0$ .
- Un-gripped chain formation with fixed  $P_d$ . In this experiment, the s-bots form chains by aligning at the rear of another s-bot, but they remain un-gripped. The lack of gripping allows any s-bot in a chain to leave the chain with a same fixed disconnection probability  $P_d$ .
- Gripped chain formation with  $P_d$  inversely proportional to the length of the chain. This experiment is

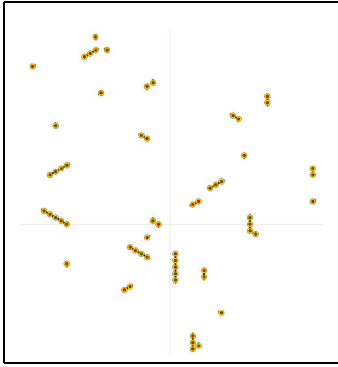


Fig. 3. A snapshot of chain formation.

the same as the first one, but the disconnection probability is modulated by the amount of light the last s-bot receives. As a result, longer chains become more stable. At the start of each experiment 60 s-bots are uniformly dispersed in a 21 by 21 grid. Each experiment is run for 3000 time steps.

The density of s-bots, their speed, the probability to grip and to stay together affect the probability of collision and hence the mean size of the chains, as well as the distribution of their size. The pattern formation process is dynamic, due to the probabilistic nature of the control, and we observed that it gets out of the initial transitory phase after approximately 2000 time steps. Figure 3 shows a snapshot of chain patterns formed in the simulation. Each experiment is run 15 times for each  $P_d$  value, and the snapshots between the 2500 and 3000 time steps are used in the analysis. In each snapshot the number  $c_i$  of chains of size  $i$  is counted ( $c_1$  corresponds to the number of chains of size 1, that is, unconnected s-bots in the environment). The mean length of chains in the environment is then computed as  $\mu_c = \sum_i i c_i / \sum_i c_i$ .

Figure 4a plots the mean length of the chains with respect to  $P_d$ , the disconnection probability for the first experiment. Note that when the probability of disconnection is 0, then each s-bot remains in the first chain it gets connected to, and the system converges into a static pattern. The average of  $\mu_c$  and its standard deviation over 15 runs are respectively approximately 3 and 0.5. As the disconnection probability increases, initially an increase in  $\mu_c$  is observed, followed by a slow decrease. This is consistent with the intuition that as  $P_d$  increases, the chains will be less likely to form, bringing down  $\mu_c$  to the minimal value of 1.

Figure 4b plots the distribution of chains at different lengths for  $P_d = 0.01$ . It can be seen that over 40% of the s-bots remain unconnected, and the number of chains formed decreases with the length of the chain.

Figure 4c plots  $\mu_c$  for the second experiment, where the chains are formed without connections. It can be seen that although the results are similar when  $P_d = 0$ , pattern formation is less stable in this case. It should also be noticed that there is no increase in  $\mu_c$  due to the lack of connections.

In the third experiment, the disconnection probability of an s-bot is defined to be inversely proportional to the light the s-bot perceives. This relation is set as  $P_d = e^{-\tau s}$ , where  $s$  is the light intensity measured. Since

only the last s-bot in the chain can disconnect, the light intensity measured by the front light sensor is used. Figure 4d plots  $\mu_c$  for different values of  $\tau$ . It can be seen that this greatly enhances the mean chain length.

### B. Stripe formation

A stripe is a 2-D pattern often observed in the collective movement of animals (a typical example are social caterpillars [10]).

The control behavior for stripe formation comprises three components which can be described as: (i) grip the rear only of s-bots of different type and disconnect with a probability proportional to the intensity of the light perceived, (ii) turn on the light when gripped by other s-bots and when the neighboring s-bots are in a “good” position for a stripe, and (iii) make one of the following moves choosing via a roulette wheel selection with a fixed probability associated to each move: (a) move towards the light, (b) move towards other s-bots, (c) align with the neighbours, and (d) make a random move.

This behavior generates a clustering mechanism, as a result of the attraction among s-bots; a chaining mechanism, due to the probabilistic gripping of s-bots of different type; and an orientation mechanism, that promotes the creation of chains parallel to each other to form stripes.

One experiment is conducted with 37 s-bots (20 of one type and 17 of the other type). Initially, the s-bots are uniformly random dispersed on a 21 by 21 grid, as shown on the left side of Fig.5. The final snapshot after simulating 20,000 time steps is shown on the right side of Fig.5. Although the process of stripe formation is dynamic, it can be seen that the system converges into a fairly good stripe pattern.

Figure 6 plots the evolution of order in the system with respect to time. Two different metrics of order are plotted. The first one, denoted as  $e_1$ , measures local order. An s-bot positioned inside a perfect stripe formation measures a local order of 6. The local order metric decreases by one for each wrongly placed neighbor.

In order to have a better evaluation, we compute the value of an additional metric that gives a more global view of the process and is related to the observable order of the structure. It is computed as the average of three components: the mean distance from the center of mass, the average orientation difference from the mean orientation of the group of s-bots, and the number of unconnected s-bots. Each component is normalized. This metric is plotted as  $e_2$  in Fig.6. The peaks and troughs in the plots of these metrics indicate that the pattern goes through structural modifications through which it becomes more ordered in time.

### C. Center-periphery pattern formation

The formation of center/periphery patterns requires all s-bots of a certain type to remain inside the swarm, whereas those of the other type form a periphery around them. The ability to form such a pattern may allow specialization and a better use of the swarm’s resources. For example, inside the body of a carrier s-bot, larger motors that provide more power and larger batteries that last longer could fill the space created by the removal of sensors. Formation of such a pattern compensates the

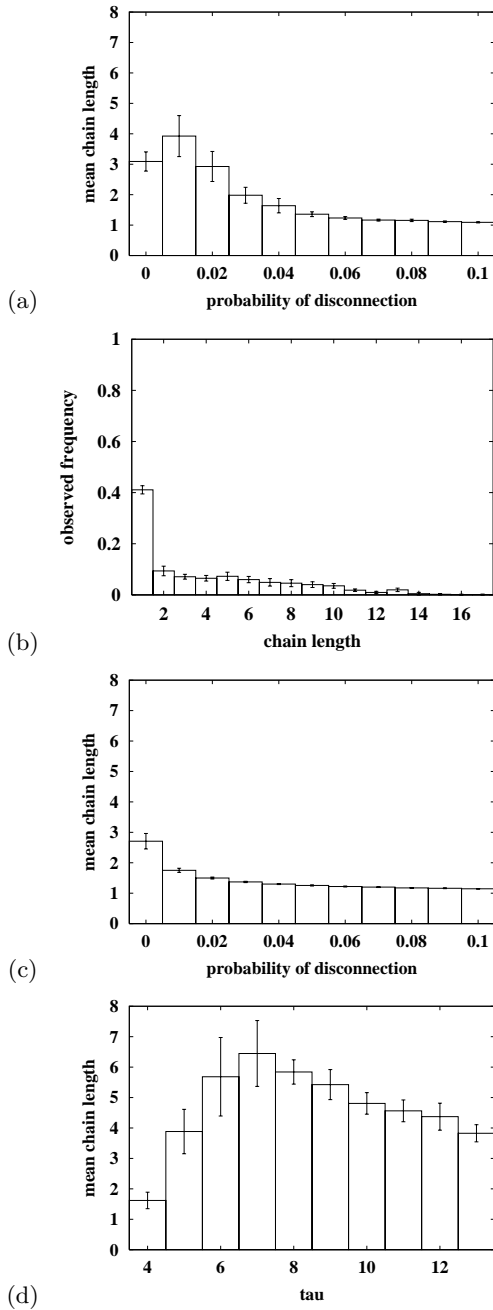


Fig. 4. (a,c,d) The mean chain length,  $\mu_c$ , is plotted with respect to the probability of disconnection,  $P_d$ . (b) The chain length distribution for  $P_d = 0.01$ .

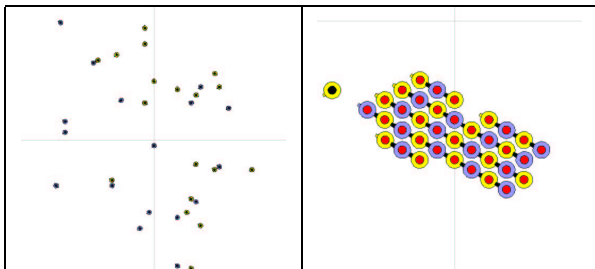


Fig. 5. Initial and final snapshots in the stripe formation experiment.

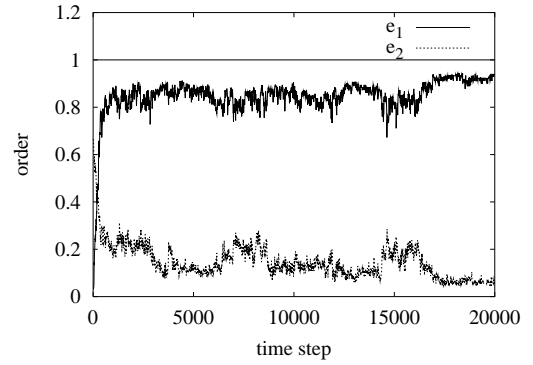


Fig. 6. Evolution of stripe formation in time as measured by two metrics:  $e_1$  represents the number of s-bots in good position as seen from a local point of view. It is normalized by the maximum value of  $e_1$  which corresponds to a striped hexagon pattern obtainable by the 37 s-bots.  $e_2$  is the global evaluation metric related to the observable order in the structure.

drop in the autonomous navigation ability (such as the ability to detect/move towards a light) of the individual carriers. Center/periphery patterns are observed in ant colonies and animal societies. For example, soldier ants create a periphery around the trail protecting the other ants that are bringing food to the nest. Inside the nests of social insects, the specialists are segregated (two or more concentric rings) or form a cluster at the center of the swarm. In some species of birds and mammals, cubs are usually kept at the center of the group, protected by a ring of adults.

The two types of s-bots, called *center* and *periphery* s-bots, have different control behaviors. For both types of s-bots the movement direction is determined by the vectorial sum of three components: attraction towards the light, attraction to/repulsion from other s-bots and a random movement vector. Center s-bots are attracted by each other whereas periphery s-bots are repulsed by both types. Both types grip and disconnect with a probability that is function of the intensity of light. Periphery s-bots turn their lights on only when they are gripped by at least two other s-bots whereas center s-bots turn their lights on also when they are surrounded by other s-bots.

This behavior accomplishes center-periphery pattern formation in two phases. In the first phase, the center s-bots cluster and connect to each other while the periphery s-bots are repelled out. In time, smaller clusters of center s-bots get dissolved into a single, big cluster, which creates a light source powerful enough for the periphery s-bots to overcome the repulsion and create the periphery.

Forty s-bots (20 of each type) are used in this experiment. Figure 7 shows the final snapshot of the center-periphery pattern obtained. Figure 8 shows the performance of the center periphery pattern formation. Different evaluation metrics are used for different types of s-bots. The first metric,  $m_b$ , defined for the center s-bots, is defined as the mean square distance of the center s-bots from their own center of mass. The second metric,  $m_y$ , computes the mean distance of periphery s-bots from center s-bots center of mass. A third evaluation metric,

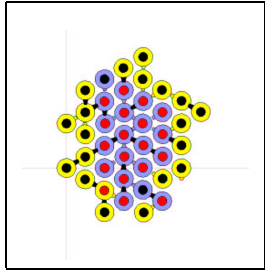


Fig. 7. The center-periphery pattern formed by the swarm-bot.

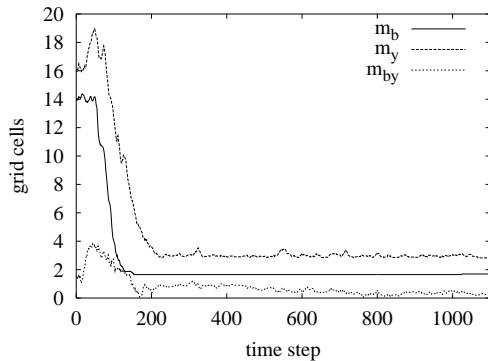


Fig. 8. Evaluation of the center/periphery formation.  $m_b$  is the mean distance of center s-bots from their center of mass.  $m_y$  represents the mean distance of periphery s-bots from center s-bots center of mass.  $m_{by}$  is the distance between the centers of mass of the two types of s-bots

$m_{by}$ , is defined as the distance between the centers of mass of the two types of s-bots.

In the first part of the experiment, center s-bots cluster since they are attracted to each other: at time step 50, two s-bots grip a third one, which turns on the light and cause a state transition for both s-bot types, which are attracted by light and start to aggregate. At time step 120, the center cluster is formed and center s-bots do not move any more, corresponding to the flat landscape of the metric  $m_b$ . At this point, the increased intensity of light of the center cluster attracts the periphery s-bots, which connect to the cluster forming the periphery of the pattern.

The interplay between attraction to light, repulsion from s-bots and random movements make the periphery a quite dynamic structure, which allows to reach a certain symmetry of the structure, as highlighted by the third metric displayed in Fig.8.

## VII. CONCLUSIONS

In this paper we have introduced the concept of *swarm-bot* and presented preliminary results about pattern formation in simulation. Inspired from social insect studies [1], we employed a probabilistic approach to control the s-bots. Three points are worth mentioning about the formation of these patterns which arised from the use of probabilistic control of the s-bots. First, the s-bots are able to create patterns, such as a single stripe pattern, that are beyond the perceptions of individual s-bots. Second, the patterns created, although

not moving, are dynamic structures to/from which s-bots constantly connect/disconnect. Third, although the pattern formation has a complex dynamic, some aspects of the resulting patterns, such as the mean length of chains in chain formation, can be controlled through some parameters such as the disconnection probability in chain formation behavior.

The swarm-bot concept is currently being implemented at different levels by the groups involved in the project: the control algorithms presented in this paper will in the near future be tested both on 3-D simulations and hardware implementations.

## VIII. ACKNOWLEDGEMENTS

Marco Dorigo acknowledges support from the Belgian FNRS, of which he is a Senior Research Associate.

The SWARM-BOTS project is funded by the Future and Emerging Technologies programme (IST-FET) of the European Community, under grant IST-2000-31010. The information provided is the sole responsibility of the authors and does not reflect the Community's opinion. The Community is not responsible for any use that might be made of data appearing in this publication.

The Swiss participants to the project are supported under grant 01.0012 by the Swiss Government.

## REFERENCES

- [1] S. Camazine, J.-L. Deneubourg, N.R. Franks, J. Sneyd, G. Theraulaz, and E. Bonabeau, *Self-Organisation in Biological Systems*, Princeton University Press, NJ, 2001.
- [2] C. Anderson, G. Theraulaz, and J.-L. Deneubourg, "Self-assemblages in insect societies," *Insectes Sociaux*, vol. 49, pp. 1–12, 2002.
- [3] E. Bonabeau, M. Dorigo, and G. Theraulaz, *Swarm Intelligence: From Natural to Artificial Systems*, Oxford University Press, New York, NY, 1999.
- [4] A. Martinoli, A. J. Ijspeert, and L. M. Gambardella, "A probabilistic model for understanding and comparing collective aggregation mechanisms," in *Advances in Artificial Life (ECAL99)*, D. Floreano, J.-D. Nicoud, and F. Mondada, Eds., Berlin, Germany, 1999, pp. 575–584, Springer Verlag.
- [5] A. Martinoli and F. Mondada, "Collective and cooperative group behaviours: Biologically inspired experiments in robotics," in *Proceedings of the Fourth International Symposium on Experimental Robotics*, Stanford, CA, 1995, pp. 3–10, Springer Verlag.
- [6] C. Melhuish, "Exploiting domain physics: Using stigmergy to control cluster building with real robots," in *Advances in Artificial Life (ECAL99)*, D. Floreano, J.-D. Nicoud, and F. Mondada, Eds., Berlin, Germany, 1999, pp. 585–595, Springer Verlag.
- [7] A. Pamecha, C. Chiang, D. Stein, and G. Chirikjian, "Design and implementation of metamorphic robots," in *Proceeding of the ASME Design Engineering Technical Conference and Computers in Engineering Conference*, Irvine, CA, 1996, pp. 1–10.
- [8] M. Yim, Y. Zhang, J. Lamping, and E. Mao, "Distributed Control for 3D Metamorphosis," *Autonomous Robots*, vol. 10, pp. 41–56, 2001.
- [9] B. Werger and M. Matarić, "Robotic food chains: Externalization of state and program for minimal-agent foraging," in *Proceedings of the 4th International Conference on Simulation of Adaptive Behavior (SAB96)*, MA, 1996, pp. 625–634, MIT Press/Bradford Books.
- [10] T.D. Fitzgerald, *The Tent Caterpillars*, Cornell University Press, Ithaca, NY, 1995.