Autonomous Robots 17, 93–105, 2004

Hormone-Inspired Self-Organization and Distributed Control of Robotic Swarms

WEI-MIN SHEN, PETER WILL AND ARAM GALSTYAN
Information Sciences Institute, University of Southern California, 4676 Admiralty Way, Marina del Rey, CA 90292, USA
shen@isi.usc.edu
will@isi.usc.edu
galstyan@isi.usc.edu

CHENG-MING CHUONG
Department of Pathology, School of Medicine, University of Southern California, 4676 Admiralty Way, Marina del Rey, CA 90292, USA
chuong@pathfinder.usc.edu

Abstract. The control of robot swarming in a distributed manner is a difficult problem because global behaviors must emerge as a result of many local actions. This paper uses a bio-inspired control method called the Digital Hormone Model (DHM) to control the tasking and executing of robot swarms based on local communication, signal propagation, and stochastic reactions. The DHM model is probabilistic, dynamic, fault-tolerant, computationally efficient, and can be easily tasked to change global behavior. Different from most existing distributed control and learning mechanisms, DHM considers the topological structure of the organization, supports dynamic reconfiguration and self-organization, and requires no globally unique identifiers for individual robots. The paper describes the DHM and presents the experimental results on simulating biological observations in the forming of feathers, and simulating wireless communicated swarm behavior at a large scale for attacking target, forming sensor networks, self-repairing, and avoiding pitfalls in mission execution.

Keywords: self organization, self reconfiguration, modular robots, distributed control, robot swarms, Digital Hormones

1. Introduction

The term robot swarm has been introduced recently to imply teams of autonomous robots that can collaboratively accomplish global missions. One special type of such swarms that is particularly interesting to us is a system that contains a great number of small and simple robots that are mobile, agile, and affordable, have local communication, and collaborate towards common goals. Just like Army ants foraging in the rainforest, once triggered and driven by some given task signals, such robot swarms will pursue their goals relentlessly. They surmount all difficulties, obstacles, destructions, and pitfalls in achieving their goals. Their individual courses may be non-deterministic but their overall behavior is organized and targeted. They do not have fixed leaders but coordinate their actions via a totally distributed control mechanism. They can self-repair damage to their organization and self-adjust their tactics and strategies.

This paper presents the Digital Hormone Model (DHM) as a bio-inspired distributed control method for robot swarms and self-organization. In this model, robots are viewed as biological cells that communicate and collaborate via hormones, and execute local actions via receptors. This model can be formalized as a

mathematical system that has three basic components: a dynamic self-reconfigurable network of autonomous robots that have "connectors" for physical or communication links; a set of probabilistic "receptor" functions that allow individual robots to select actions based on their local topology, states, sensors, and received hormones; and a set of equations for hormone diffusions and reactions. The diffusion-reaction "radius" can be interpreted physically in robots as the cell radius in a mobile phone system or a sensor network, or the range of a walkie-talkie device. Robots interact inside the cell radius directly and by relayed messages across cells.

This model allows many simple robots in large-scale systems to communicate and react to each other, and self-organize into global patterns that are suitable for the given task and environment. All communication and reactions are local and use signals that are similar to hormones that regulate cell activities in biological organisms, and require no unique identifiers for individual robots. The model combines advantages from Turing's reaction-diffusion model, stochastic reasoning and action, dynamic network reconfiguration, distributed control, self-organization, and adaptive and learning techniques.

2. Related Work

Throughout the history of computer science, there have been many computational models for coordination and self-organization among a large number of autonomous entities. Turing's reaction-diffusion model (Turing, 1952) is perhaps one of the earliest. Turing used differential equations to model the periodic pattern formation in a ring of discrete cells or continuous tissues that interact with each other through a set of chemicals he called "morphogens." The morphogens can diffuse from cell to cell and can react with each other. Turing analyzed the interplay between the reactions and diffusions of morphogens and concluded that their nonlinear interactions could lead to the formation of spatial patterns in their concentrations. Turing's model was startlingly novel, and since the publication of this reaction-diffusion model, it has been supported both mathematically (Murray, 1989) and experimentally (Ouyang and Swinney, 1991), and many applications have been described (Meinhardt, 1982). Models have been developed not only for local interactions, but also for incorporating long-range interaction such as those in large-scale spatially organized neural nets (Murray, 1989). In computer science, Witkin and Kass (1991) extended the original reaction-diffusion model by allowing anisotropic and spatially non-uniform diffusion, as well as multiple competing directions of diffusion. The extended model has been successfully applied to synthesis of textures with different patterns.

Other techniques that are closely related to the Digital Hormone Model include the Stochastic Cellular Automata (SCA) (Gutowitz, 1991; Toffoli, 2000; Lee et al., 1991) and Amorphous Computing (AC) (Abelson et al., 1999; Nagpal, 1999; Wolpert, 1969). In Artificial Life, many mathematical models have been proposed for pattern development based on distributed cells (Takagi and Kaneko, 2002). In fact, the DHM has been proposed as a potential mechanism for development and differentiation in Artificial Life systems (Shen et al., 2002). Swarm robotics (Bonabeau et al., 1999) is a very active research area and has many proposed approaches. In comparison with our hormone-inspired approach (Shen et al., 2002), the most related approach is the pheromone-based control (Parunak and Brueckner, 2001; Payton et al., 2002), which show that a set of autonomous agents can use pheromones to form interesting and complex global behaviors and exhibit swarming behaviors (Parunak, 2003).

The concept of biological hormone (Kravitz, 1988) has inspired many researchers to build computational systems. These include Autonomous Decentralized Systems (ADS) (Ihara and Mori, 1984; Mori et al., 1985), homeostatic (different from hormones) robot navigation (Arkin, 1992), and integration of behaviors (Brooks, 1991). The ADS are probably the earliest attempt to build systems that are robust, flexible, and capable of doing on-line repair. The ADS technology has been applied to industrial problems (Mori, 1999), and has the properties of on-line expansion, on-line maintenance, and fault-tolerance.

The DHM presented in this paper is different from the above approaches for self-organization and swarming control. In particular, the DHM extends Turning's reaction-diffusion model by considering not only the interplay between reactions and diffusions, but also the network topological structure around each robot, the local sensory and actuator states, and the movements of individual robots. The consideration of topology information also distinguishes the DHM from the models in Amorphous Computing where the primarily concerns are the positional information of individual entities. Compared to Cellular Automata, the equations in DHM deal with continuous space, thus more suitable for modeling spatial behaviors of mobile robots in

real-world environments. Different from the ADS, the DHM is also applicable to self-reconfigurable systems and robots where new configurations can be planned and executed based on the environment and the tasks in hand (Shen et al., 2002). Different from pheromone-based approaches, the DHM does not use residue-like mechanism for propagating signals but relying on diffusion and reaction among many different signals. One advantage is that DHM can establish distributed control without assuming any "place agents" as in Parunak et al. (2002). In addition, the DHM has explicit representations for network links (physically or virtual), and supports dynamic changes of these links. Finally, it is interesting to notice that most existing approaches for robot swarms assume that robots have globally unique identifiers for communication and cooperation, where the DHM can do without this assumption just as in biological systems cells are not known to have any globally unique identifiers (Jiang et al., 1999; Chuong et al., 2000; Yu et al., 2002).

3. The Digital Hormone Model

The DHM is inspired by four factors: (1) biological discoveries about how cells self-organize into global patterns, (2) the existing self-organization models, such as Turing's reaction-diffusion model, (3) the stochastic cellular automata, and (4) the distributed control systems for self-reconfigurable robots. In biological systems, different cells respond to different hormones because different cells have different receptors designed to bind with particular hormones. The different types of hormones and target cells present in vertebrates are so great that virtually every cell either processes or responds to one hormone or another. Hormones provide the common mechanism that makes it possible for cells to communicate without identifiers and addresses, and they support a broad spectrum of seemingly diverse biological effects.

The basic idea of the Digital Hormone Model is that a swarm is a network of robots that can dynamically change their links in the network. Through the links in the network, robots use hormone-like messages to communicate, collaborate, and accomplish global behaviors. The hormone-like messages are similar but not identical to content-based messages. They do not have addresses but propagate through the swarm. All robots have the same decision-making protocol, but they will react to hormones according to their local topology and state information so that a single hormone may cause different robots in the network to perform different actions. Note that hormone propagation is different from message broadcasting. There is no guarantee that every robot in the network will receive the same copy of the original message because a hormone may be modified during its propagation.

Mathematically speaking, the Digital Hormone Model consists of three components: a specification of a dynamic network, a probabilistic function for individual robot behavior, and a set of equations for hormone reaction, diffusion, and dissipation.

A *Dynamic Network of Swarm Robots* (DNSR) is specified as a network of N autonomous robots. Each robot has a set of *connectors* through which the robot can dynamically connect to other robots to form *edges* for communication or physical (mechanical) coupling. The concept of connector is theoretically new but it has been used in many engineered systems. For example, in a wireless network, the connectors of a robot are the channels it can use to communicate with others. A channel of a robot must be "connected" to a channel of another robot to form an edge of communication. In self-reconfigurable robots, the connectors are physical so that an edge is a physical coupling and a network of robots can form physical structures with different shapes and sizes. The connectors are valuable and finite resources for robots. Because connectors can be joined and disjoined, they make the edges in a network dynamic, and the reconfiguration of network possible. Let N_t and E_t denote all the robots and edges that exist in a dynamic network at time t, then a DNSRat time t can be defined as:

$$DNSR_t \equiv (N_t, E_t) \tag{1}$$

Note that both N_t and E_t can change dynamically because robots can autonomously join, leave, or be damaged, and edges can be formed and disconnected by the connectors of the robots. Different from classical models, robots do not have unique IDs, the number of robots and edges in the network is not known, and there is no global broadcast. A robot can only communicate with its current neighbors through its current edges. This local communication assumption is realistic and necessary for large-scale DNSR systems, for two arbitrary robots can be so far away that direct communication is not possible, especially when robots only have limited resources. Through local communication, robots can either generate hormones or propagate hormones. By default, a generated hormone will be sent to all the current edges of its generator, and a received hormone

will be propagated to all the current edges except the one through which the hormone is received.

The second component of the Digital Hormone Model is a specification of individual robot behavior, which is similar to the concept of receptors in biological cells. A robot in the network can select its actions, *B,* based on a probability function, *P,* that is conditioned on four local factors: the connector information, *C*; the sensor information, *S*; the values of local variables, *V*; and the received hormones, *H*:

$$P(B \mid C, S, V, H) \tag{2}$$

The actions, *B*, of a robot include the commands to the local sensors and actuators, as well as actions that change connectors and generate or propagate hormones. Different from most existing probabilistic models, such as hidden Markov models, partially observable Markov decision processes, and reinforcement and Q-Learning models, the *P* function here considers not only sensor and state information *S* and *V*, but also topological information, *C*, and communication information, *H*. These allow the Digital Hormone Model to support dynamic reconfigurations and self-organization in network structures. The function *P* is local and homogenous for all robots, but can greatly influence the global behaviors of the network and predict and analyze the global network performance in the large. For example, in the simulation of forming of feathers to be discussed later, the characteristics of *P* can influence whether or not any global patterns can be formed. Biologically speaking, we believe that the function *P* partially simulates the hormone receptors and the control mechanisms found in the biological cells. The *P* function is programmed by the system designers initially, but can be dynamically changed by the robots themselves through learning techniques.

The third part of the Digital Hormone Model is the specification for hormone reaction, diffusion, and dissipation. Following Turing (1952) and Witkin (1991), we assume in the mathematical description that hormone reaction and diffusion occur through a two-dimensional medium, although analogous results can be derived for arbitrary dimensions and some higher dimensions are indeed used in applications of self-reconfigurable robots. The concentration of each hormone is a function of position and of time. We denote the concentration function for a particular hormone by $C(x, y)$, where x and y are 2D space dimensions. The reaction-diffusion-dissipation equation governing the hormone is then given by:

$$\frac{\partial C}{\partial t} = \left(a_1 \frac{\partial^2 C}{\partial x^2} + a_2 \frac{\partial^2 C}{\partial y^2} \right) + R - bC \tag{3}$$

The first term on the right is for diffusion, and a_1 and a_2 are constants that represent the rate of diffusion in x and y directions respectively. The function R is the reaction function governing C, which depends on all the other concentrations of hormones. The constant b is the rate for dissipation. The Eq. (3) is usually considered to be a part of an environmental function G responsible for the implementation of the dynamics of communication or other effects of actions. For example, if two robots send out radio signals with the same frequency at the same time, then G will be responsible for simulating the interference between the two signals. Although the G function is in principle a part of the environment, it can be simulated by the actions of the robots as described later.

As we can see from the above definitions, the Digital Hormone Model is an integration of *dynamic network* (Eq. (1)), *topological stochastic action selection* (Eq. (2)), and *distributed control by hormone reaction-diffusion* (Eq. (3)). This integration provides a very powerful coordination mechanism for dynamic networks of swarm robots. The execution of DHM is very simple. All robots in the swarm asynchronously execute the basic control loop in Fig. 1.

To demonstrate the DHM, let us define a simple DHM_0 shown in Fig. 2. In this simple model, cells are shown as black dots and can move in a space of discrete grids. Each cell occupies one grid at a time and can secrete hormones (shown as the gray areas around a cell) to the neighboring grids to influence other cells' behaviors. For simplicity, we assume for now that all cells synchronize their actions and the grids carry out the reaction and diffusion of hormones. A cell at a grid (a, b) can secrete two types of hormones, the activator A and the inhibitor I. The *diffusion* of A and I at a

1. Select actions by $P(B|C,S,V,H)$;
2. Execute the selected actions in B;
3. Perform hormone generations and propagations;
4. Simulate hormones diffusion, reaction, and dissipation;
5. Go to Step 1.

Figure 1. The basic control loop in DHM.

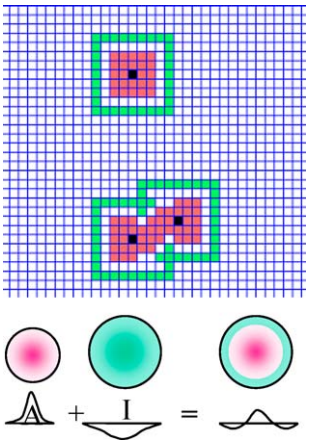

Figure 2. The simple DHM_0.

surrounding grid (x, y) are given by the standard distribution functions:

$$C_A(x, y) = \frac{a_A}{2\pi\sigma^2} e^{\frac{(x-a)^2+(y-b)^2}{2\sigma^2}} + R \tag{4}$$

$$C_I(x, y) = -\frac{a_I}{2\pi\rho^2} e^{\frac{(x-a)^2+(y-b)^2}{2\rho^2}} + R \tag{5}$$

where a_A, a_I, σs and ρ are constants, and $\sigma < \rho$ in order to satisfy the Turing stability condition that the diffusion rate of the inhibitor must be greater than that of the activator. Note that because $\sigma < \rho$, A has a sharper and narrower distribution than I, and these characteristics are similar to those observed in the biological experiments (Jiang et al., 1999; Chuong et al., 2000; Yu et al., 2002). We assume that the hormone A has the positive value and the hormone I has the negative value. For a single isolated cell, the hormone concentration in its neighboring grids looks like three "colored rings" (see the lower-right corner in Fig. 2). The activator hormone dominates the inner ring; the inhibitor hormone dominates the outer ring; and the middle ring is neutral where the hormones of A and I have canceled each other. The *reaction* between two hormones in a grid is computed by summing up all present concentrations of "A"s and "I"s in the grid:

$$R = \sum_N (C_A + C_I) \tag{6}$$

When two or more cells are near each other, the hormones in the surrounding grids are summed up to compute the combined hormone concentration. In the upper part of Fig. 2, we have illustrated the combined hormone concentrations around a single cell and around two nearby cells. Since the grids are discrete, the rings around the cells are shown as squares instead of circles.

When all cells are moving in synchronization, there may be a chance that multiple cells will "collide" in the same grid. The collision of cells is solved in a simple manner. All cells first "virtually" move to the grids they selected. If there are multiple cells in the same grid, then the extra cells will be randomly distributed to those immediate neighboring grids that are empty. This is an environmental function, not a cellular action. But this action will ensure that no grid is hosting more than one cell at any time.

For cell behaviors, DHM_0 is governed by a function $P_0(B \mid C, S, V, H)$ defined as follows:

B: Each cell has ten actions. B_0 for secreting the A and I hormones, and $B_1, \ldots, B_9$ for moving into the nine neighboring grids: north, south, west, east, northeast, northwest, southeast, southwest, and self (the occupying grid);
C: Each cell has eight connectors in this simple model, one for each neighboring grid;
S: Each cell has nine hormone sensors, one for each of the neighboring grids;
V: Cells have no local variables in this model;
H: The nine hormone values sensed by the sensors;

$$P_0(B \mid C, S, V, H) = \begin{cases} P_0(B_0 \mid C, S, V, H) = 1.0; \\ P_0(B_i \mid C, S, V, H) = \textit{BestNeighborFunction}, \text{ where } i = 1, \ldots, 9. \end{cases}$$

The function *BestNeighborFunction* is defined so that the probability of moving to a particular neighboring grid is proportional to C_A and inversely proportional to C_I in that grid, and the sum of these probabilities is 1. Every cell always executes B_0 to secret hormones. Note that the probability $P_0(B \mid C, S, V, H)$ is computed in two independent parts: one for B_0, and the other for B_1 through B_9.

Given Eqs. (4), (5), and (6) $P_0(B \mid C, S, V, H)$, DHM_0 can be used to investigate how hormones affect self-organization and whether they can enable locally interacted robots to form globally interesting patterns. We can also change the characteristics of these parameters, and observe and analyze the global effects in the large.

4. The DHM for Self-Organization

In order to apply the proposed DHM to robot swarm, it is important to realize the analogy between robot swarm and biological morphogenesis. Morphogenesis is a process in which many cells move and grow simultaneously to form organisms or body parts. Similar to individual robots in a swarm, each cell must interact and collaborate with other cells in order to perform the correct local action to achieve the desired global results. Interestingly enough, recent research in the forming of feathers has revealed that this process may be distributed among all the related cells. In this section, we will briefly introduce this biological process and then illustrate how DHM is used to simulate the "swarm" behaviors of cells. This exercise will provide a foundation for us to apply DHM to robot swarms.

In biological systems such as chicken embryos, feathers are developed as follows. First, homogeneous skin cells aggregate and form feather buds that have approximately the same size and space distribution. The feather buds then grow into different types of feathers depending on the region of the skin. Earlier theories believed that such a process was initiated and coordinated by some "key" cells on the skin. However, these theories have been challenged by the recent findings in biological experiments. Even if the original distribution of the homogenous skin cells is randomly altered, the cells can still grow into feather buds (Jiang et al., 1999; Chuong et al., 2000; Yu et al., 2002). These experiments suggest that there are no predetermined molecular addresses, and that the periodic patterning process of feather morphogenesis is likely a self-organizing process based on physical-chemical properties and reactions between homogeneous cells. During these biological experiments, biologists observed some interesting relations among the reaction and diffusion characteristics of the hormones secreted from the cells, the size and space distribution of the final feather buds, and the initial density of cell population. In particular, they observed that while the number of formed feather buds is proportional to the cell population density, the size of the feather buds remains approximately the same regardless of different population densities. The size of the feather buds, however, is related to the reaction and diffusion profiles of the activator and inhibitor hormones secreted from the cells. If the concentration ratio of the activator hormone to the inhibitor hormone is high, then the final size of the feather buds will be larger than usual. If the ratio is balanced, then the size of the formed feather buds will be normal. If the ratio is low, then the size of the formed pattern will be smaller than usual. These observations are most interesting to us because they can be used as basic criteria for validating or falsifying any simulation models for such distributed and organizational behaviors.

Using the DHM such as the one defined in the last section, we would like to investigate:

- Will DHM_0 enable cells to self-organize into patterns at all?
- Will the size of final patterns be invariant to the cell population density?
- Assuming that the hormone diffusion profiles are fixed, will the results match the observations made in the biological experiments?
- How do the hormone diffusion profiles affect the size and shape of the final patterns as shown in the biological experiments?
- Will an arbitrary profile enable self-organization and pattern formation?

To answer these questions, we have conducted two experiments. In the first experiment, we use the same hormone diffusion profile and run a set of simulations on a space of 100×100 grids (using periodic boundary conditions) with different cell population densities ranging from 10% ($\sim$1000 cells) through 75% ($\sim$7500 cells). Starting with cells randomly distributed on the grids, each simulation runs up to 1,000 action steps, and records the configuration snapshots at steps of 0, 50, 500, and 1,000. As we can see from the results in Fig. 3, cells in all simulations indeed form patterns. We observe that for relatively small densities (up to 40%) the cells form isolated clusters as shown in the two bottom rows. Furthermore, it seems that the size of those clusters depends very weakly on the cell density. This results matches the observations made in the biological experiments. If one increases the cell density, on the other hand, the cells start to form stripe-like patterns (two top rows in Fig. 3). Note that orientation of the stripes can be both vertical and horizontal, depending on the initial cell distribution.

In the second set of experiments, we started with the same cell population density, but varied the hormone diffusion profiles by changing the parameters for Eqs. (4) and (5). We wanted to observe the effects of different hormone profiles on the results of pattern formation. As we can see in Fig. 4, when a balanced profile of activator and inhibitor is given (see the second row),

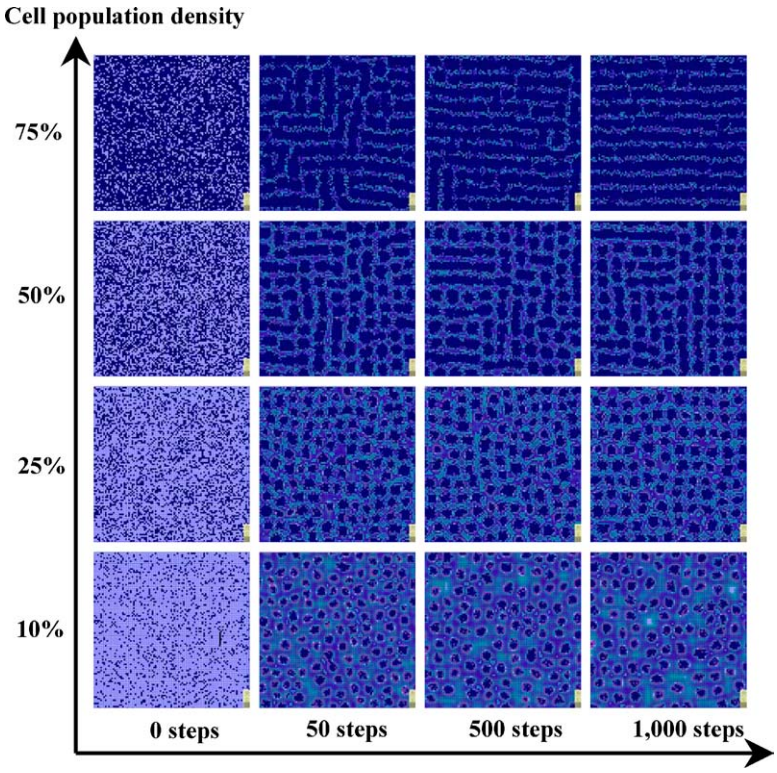

Figure 3. Pattern formation with different cell density but a common hormone profile.

the cells will form final patterns as in the first set of experiments. As the ratio of activator over inhibitor (σ/ρ) increases, the size of final clusters also increases (see the third row). These results are a qualitative match with the findings in the reported biological experiments (Jiang et al., 1999).

When the ratio of A/I becomes so high that there are only activators and no inhibitors (increases a_A/a_I), then the cells form larger and larger clusters through cluster aggregation (see the fourth row). On the other hand, when the ratio is so low that there is only inhibitor and no activator, then the cells will never form any patterns (see the first row), regardless of how long the simulation runs. This shows that not all hormone profiles enable self-organization. These results are yet to be seen in biological experiments, but they are consistent with the principles of hormone-regulated self-organization and thus qualified as meaningful predictions of cell self-organization by hormones.

5. The DHM for Robot Swarming

Although the above simulation results have shown that the DHM can indeed demonstrate the self-organization for cell-like development and differentiation in pattern formation, practical details of going from the simulation to physical robots are usually significant enough to overshadow the basic attraction/repulsion, reaction/diffusion concepts. Thus, it is important to move beyond the basic "grid world" simplifications of simulation to more realistic settings.

To apply the DHM to realistic mobile robots, the first question we face is how to implement the diffusion and reaction of hormones in a robot swarm. To solve this problem, we assume that all robots have short-range wireless communication (either RF or Infrared) and can talk to robots that are in proximity. Different from pure biological experiments that require geographic proximity between cells, the DHM requires only topological proximity in which a neighbor robot is defined as one directly reachable in a single communication hop. To implement the secretion of a hormone, each robot broadcasts a signal that carries the type information of that hormone. To implement the diffusion of a hormone, each receiver robot determines the direction (e.g., via a directional antenna) of the incoming signal and the distance of the signal source (e.g., by measuring the strength of the signal), and then

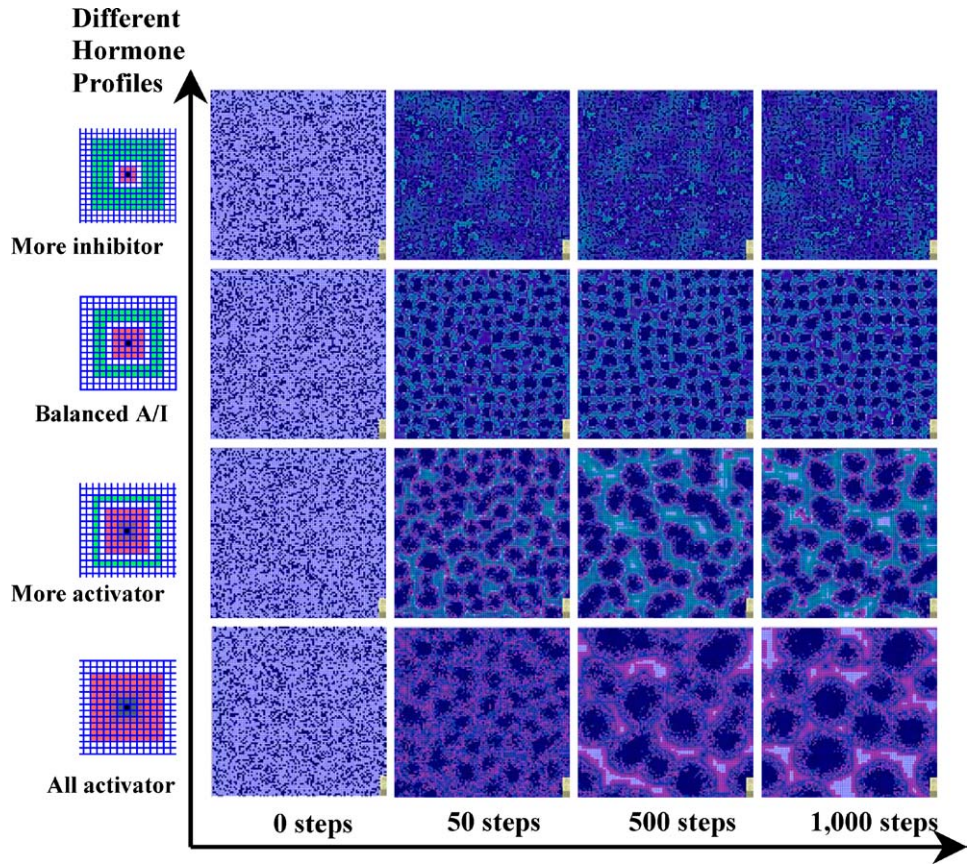

Figure 4. Pattern formation of different individual hormone profiles.

applies the relevant diffusion functions (e.g., Eqs. (4) and (5)) to compute the "concentration" values of that particular hormone at the current and nearby locations. To implement the reaction of hormones, each robot collects all hormonal signals in a period of time, and then computes the reaction of the collected hormones using the reaction function (e.g., Eq. (6)). Using this solution, the control loop of each robot is the same as in Fig. 1, except that in Step 4 each robot must collect wireless hormonal signals, compute concentration value for each received hormone, and then compute the reaction of the collected hormones. To avoid obstacles and collision, each robot must now check the selected moving direction before moving. If that direction is blocked, then the robot must switch to the next best direction. In this solution, robots are not in a discrete grid world but in continuous space. The DHM supports the movements in continuous space because the equations (e.g., 2, 3, 5, 6) are all continuous. In situations where orientation-special patterns are to be formed, the implementation of DHM would require a means for all robots to maintain a common orientation reference (e.g., a compass).

With this new implementation of hormone diffusion and reaction for mobile robots in swarms, we have conducted a set of experiments in simulation to test the swarming behaviors of the DHM in spaces of closed boundaries. The experiments are (1) searching and seizing targets; (2) distributing and monitoring a given area or building; (3) self-repairing damages to the global patterns; and (4) avoid pitfalls by detouring. We now present the details of these experiments.

5.1. Searching and Seizing Targets

In this first experiment, we assume that there are targets in the environment that can be sensed by the robots in short distance. The task for a robot battalion is to search and seize such a target. The first row in Fig. 5 shows the conceptual idea of this task. Driven by the repulsive hormone, the robots first disperse uniformly from their initial location, and then some robots find

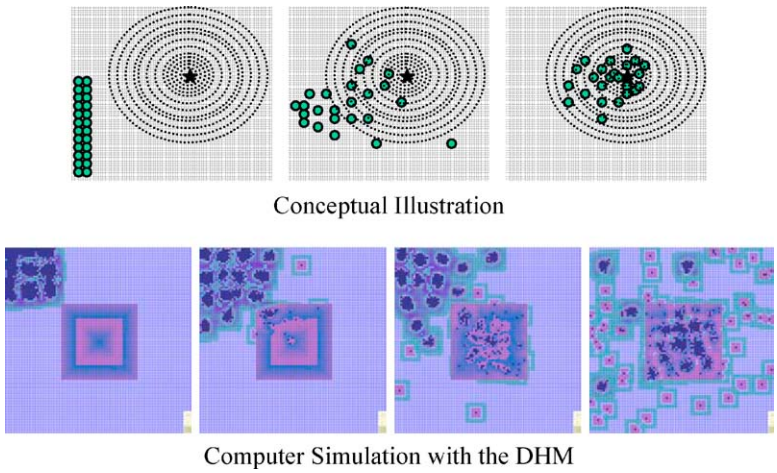

Figure 5. A swarm searching and seizing a target.

the target and are attracted to and aggregated around the target (due to the attractive hormone) and their aggregated hormone signals will create a large field to attract other robots. Such attraction will be propagated throughout the robot swarm and a gradient field for the target will be created. All robots will probabilistically follow the gradient and eventually surround the target. The location of the target may be static (such as a geographic location) or dynamic (such as an moving vehicle). Since robot's actions are stochastic, they will wander in the field to find other targets. To implement this behavior, we introduced a "target signal source" that will continuously generate activator hormone signals into its surrounding space and creates a hormone field to attract nearby robots.

The simulation results are shown in the second row in Fig. 5, where the space has closed boundaries (robots cannot go through them) and the robots are initially concentrated at the up-left corner. Once start running, the robots first wander around as before dispersing uniformly from the corner, but soon some of them are attracted to the target signal field. As time goes by, a sufficient number of robots are attracted by the signal and form an enclosure around the target. Notice that not all robots are devoted to the same target, and there are sufficient robots still searching for targets in the open space. This automatic task balancing is due to the non-deterministic robot behavior function in the DHM.

In reality, the target signal field can be created in many different ways. One obvious implementation is to launch a signal source at the target location or attach a signal source to the target object. The other way is to use GPS (Global Position System) to specify a target location and simulate the hormone field in the robot's communication and sensing systems. In other words, all robots will receive an attractive hormonal signal as if it was broadcasted from a specific GPS location.

5.2. *Spread and Monitor in a Building*

The DHM can also enable a swarm of robots automatically cover an area without any complicated or centralized control strategy. This problem of "area coverage" has also been addressed by many other approaches, including particle-based and potential-field-based approaches such as Howard and Mataric (2002). The similarity is that our repulsive and attractive hormones are analogous to their repelling and dissipative viscous forces; but the difference is that hormones are driven by diffusion and reactions while forces are governed by field properties and DHM is distributed and requires no unique identifiers for robots in the swarm.

The first row in Fig. 6 illustrates the concept of area coverage task, where a swarm of robots are to cover a floor of a building without knowing the layout of the rooms. When a large number of robots enter the floor, they are driven deeper into the empty space because the repulsive hormone is pushing them apart. Each robot has a higher probability to move away from a place where the repulsive hormone is strong (i.e., there are too many robots). The robots will not spread too thin because the attractive hormone is pulling them together. The balance between repulsive and attractive hormones ensures the formation of the desired global patterns.

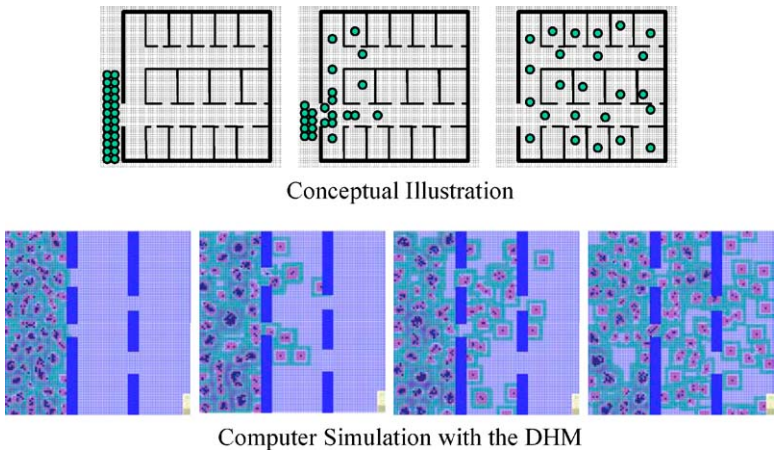

Figure 6. A swarm spreading into a building.

The second row in Fig. 6 shows the computer simulation of using the DHM to accomplish the task of spreading a swarm of robots into an unfamiliar building. The floor has three zones separated by walls and each wall has two "doors" to connect the zones. We assume the wireless signals can penetrate the walls. Initially, all robots are in the left zone. Driven by their hormone signals, some robots are pushed through the doors into the middle zone. They then gradually spread out into the third zone. Notice that the robots automatically and evenly distributed themselves in these zones without explicitly being commanded to do so. This is another demonstration that swarm-level self-organizing behaviors can achieved based on local interactions among robots. The degree of spread can be controlled by the strength of the repulsive hormones generated by the robots.

5.3. *Self-Repair Unexpected Damages*

The third experiment demonstrates that a swarm of robots controlled by the DHM can self-repair unexpected damages to their organization. Unlike classical network protocols that cannot adapt to dynamic network topology, hormone-controlled robots can use the presence/absence of hormone signals to self-adjust their topology connections (via changing locations in this example) to self-heal the damage.

The first row in Fig. 7 shows the task at a conceptual level. Assume that in a stabilized network of mobile

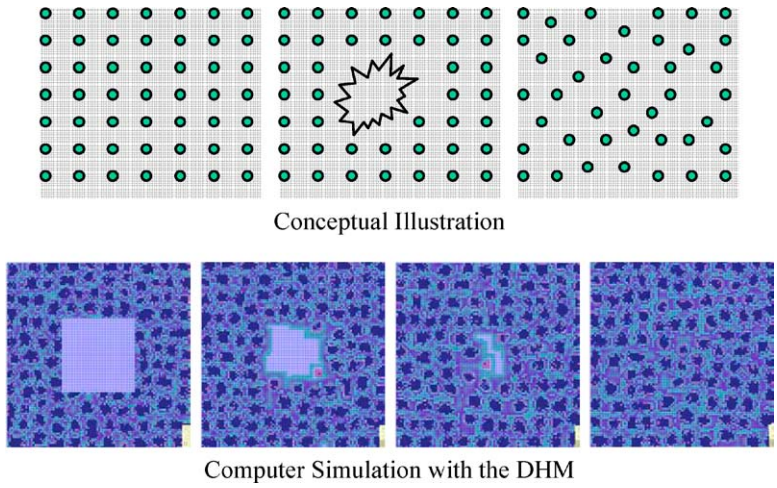

Figure 7. A swarm self-repairing behavior.

robots, a bomb explodes at the center of the network and damages many robots there. In this case, the remaining robots will move into the empty space and the result will be a new (thinner) network with fewer robots.

The second row in Fig. 7 illustrates the computer simulation with the DHM control. We first run the simulation and let the robots forming a stabilized network, and then we manually removed 15% of the robots from the middle of the global pattern (first column in Fig. 7). We then let the robots continue to run for 50, 100, 500, and 1000, and observed that the robots self-repaired the hole completely. This demonstrates that as long as the local hormone profiles are in effect, the global patterns can repair themselves even after severe damage. The speed of repair is quite fast in simulation (about 500 simulation steps).

5.4. Surmount and Detour in Mission Execution

In the process of achieving their goals, a swarm of robots will inevitably face barriers, obstacles, and pitfalls. For example, as shown in the conceptual illustration at the first row in Fig. 8, when more and more robots are trapped at a barrier, the repulsive hormone will be so strong there that some robots will be pushed away to a "detour" route. These "free" robots will in turn attract those robots that are trapped at the barrier. Because there are more robots outside the pitfall than inside, robots that are at the critical forking point between the correct path and the trapped path will be attracted more to the outside group. As more and more robots choose the correct path, the correct signals will become stronger and stronger, and eventually overcome the signals from the trapped robots. As a final result, the majority of robots will bypass the barrier.

The second and third rows in Fig. 8 show the computer simulation of this behavior. Initially, all robots are located at the up-left corner and we placed an L-shaped barrier around these robots. In the first experiment we examined whether the robots can detour the barrier using only hormone-induced diffusion. As shown in the figure, many robots first are trapped in the barrier. However, the robots are able to detour the obstacle, without any introduced bias. In the second experiment, we deliberately introduce a bias so that all robots are moving towards a target area that is located in the lower right corner. As it can be seen, after sufficient time the majority of the robots escapes the obstacle and founds the target. In general, the situations of barriers and traps can be arbitrarily complex and impossible to predict. But with the ability to self-organize and self-repair, a swarm of robots can find a way to bypass the barrier and traps under the control of the DHM.

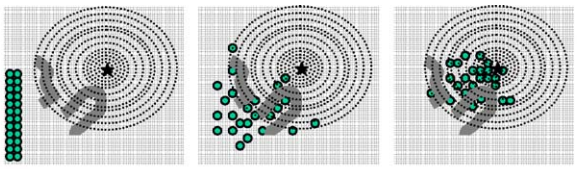

Conceptual Illustration

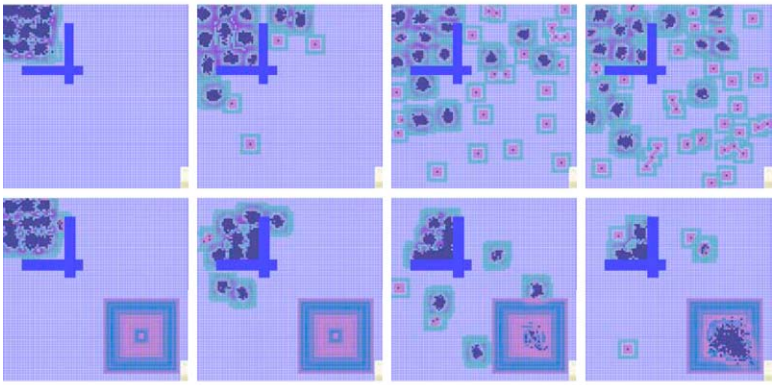

Computer Simulation with the DHM.

Figure 8. Bypassing a dead-end barrier in operation.

6. Conclusion

This paper presents the Digital Hormone Model (DHM) as a distributed control method for robot swarming behaviors and self-organization. The model combines advantages from Turing's reaction-diffusion model, stochastic reasoning and action, dynamic network reconfiguration, distributed control, self-organization, and adaptive and learning techniques. Such a model allows many simple robots in large-scale systems to communicate and react to each other to self-organize into global patterns that are suitable for the given task and environment, and it requires no unique identifiers for individual robots. The paper first demonstrates the utility of the Digital Hormone Model by simulating self-organization in the forming of feathers in biological systems. It then proposes an physical implementation of hormone diffusion and reaction among mobile robots in swarms, and demonstrates the implementation in simulation for swarming behaviors such as searching and seizing targets, distributing and forming sensor networks, self-repairing unexpected damages, and avoiding pitfalls by detouring. The advantages of the DHM include its locality, simplicity, robustness, and self-organization.

Acknowledgment

We are grateful that this research is in part supported by AFOSR under contract numbers F49620-01-1-0020 and F49620-01-0441, and by grants from NIH and NSF.

References

Abelson, H., Allen, D. Coore, D., Hanson, C., Homsy, G., Knight, T.F., Nagpal, R., Rauch, E., Sussman, G.J., Weiss, R. 1999. *Amorphous Computing*. MIT: Boston.

Arkin, R.C. 1992. Homeostatic control for a mobile robot: Dynamic replanning in hazardous environments. *Journal of Robotic Systems*, 9(2):197–214.

Bonabeau, E., Dorigo, M., Theraulaz, G. 1999. *Swarm Intelligence: From Natural to Artificial Systems*. Oxford University Press.

Brooks, R. 1991. Integrated systems based on behaviors. Special Issue of SIGART on *Integrated Intelligent Systems*.

Chuong, C.-M., Chodankar, R., Widelitz, R.B., Jiang, T.X. 2000. Evo-devo of feathers and scales: Building complex epithelial appendages. *Current Opinion in Development and Genetics*, 10:449–456.

Gutowitz, H. (Ed.). 1991. *Cellular Automata—Theory and Experiment*. The MIT Press: Cambridge, MA.

Howard, A. and Mataric, M. 2002. Development and localization for mobile robot teams. In *Multi-Robot Systems: From Swarms to Intelligent Autonomous*, A.C. Schultz and L.E. Parker (Eds.). Kluwer Academic Publishers, pp. 41–51.

Ihara, H. and Mori, K. 1984. Autonomous decentralized computer control systems. *IEEE Computer*, 17(8):57–66.

Jiang, T.-X., Jung, H.S., Widelitz, R.B., and Chuong, C.-M. 1999. Self organization of periodic patterns by dissociated feather mesenchymal cells and the regulation of size, number and spacing of primordia. *Development*, 126:4997–5009.

Kravitz, E.A. 1988. Hormonal control of behavior: Amines and the biasing of behavioral output in lobsters. *Science*, 241(4874):1775–1781.

Lee, Y.C., Qian, S., Jones, R.D., Barnes, C.W., Flake, G.W., O'Rouke, M.K., Lee, K., Chen, H.H., Sun, G.Z., Zhang, Y.G., Chen, D., Giles, G.L. 1991. Adaptive stochastic cellular automata: Theory and experiment. In *Cellular Automata*, H. Gutowitz (Ed.). The MIT Press: Cambridge, MA, pp. 159–188.

Meinhardt, H. 1982. *Models of Biological Pattern Formation*. Academic Press: London.

Mori, K. 1999. Autonomous decentralized system technologies and their application to train transport operation system. In *High Integrity Software Conference*. Albuquerque, New Mexico.

Mori, K., Miyamoto, S., and Ihara, H. 1985. Autonomous decentralized computer system and software structure. *Computer Systems Science and Engineering*, 1(1):17–22.

Murray, J.D. 1989. *Mathematical Biology*. New York: Springer-Verlag.

Nagpal, R., 1999. Organizing a global coordinate system from local information on an amorphous computer. MIT: Boston.

Ouyang, Q. and Swinney, H.L. 1991. Transition from a uniform state to hexagonal and striped Turing patterns. *Nature*, 352:610–612.

Parunak, H.V.D. 2003. Making swarming happen. In *Conference on Swarming and C4ISR*. Tysons Corner, VA.

Parunak, H.V.D. and Brueckner, S. 2001. Entropy and self-organization in multi-agent systems. In *International Conference on Autonomous Agents*. Montreal, Canada.

Parunak, H.V.D., Purcell, M., and Connell, R.O. 2002. Digital pheromones for autonomous coordination of swarming UAVs. In *The First AIAA Unmanned Aerospace Vehivales, Systems, Technologies, and Operations*. American Institute of Aeronautics and Astronautics.

Payton, D., Estkowski, R., and Howard, M. 2002. Progress in pheromone robotics. In *The 7th International Conference on Intelligent Autonomous Systems*. Marina del Rey, CA: IOS Press.

Shen, W.-M., Chuong, C.-M., and Will. P. 2002. Digital hormone models for self-organization. In *The 8th International Conference on Simulation and Synthesis of Living Systems (ALIFE VIII)*. Sydney, Australia.

Shen, W.-M., Chuong, C.-M., and Will, P. 2002. Simulating self-organization for multi-robot systems. In *International Conference on Intelligent and Robotic Systems*, Switzerland.

Shen, W.-M., Salemi, B., and Will, P. 2002. Hormone-inspired adaptive communication and distributed control for CONRO self-reconfigurable robots. *IEEE Transactions on Robotics and Automation*, 18(5).

Takagi, H. and Kaneko, K. 2002. Dynamic relationship between diversity and plasticity of cell types in multi-cellular state. In *The 8th International Conference on the Simulation and Synthesis of Living Systems (ALife VIII)*. University of New South Wales, Sydney, NSW, Australia.

Toffoli, 2000. Cellular automata. In *The Handbook of Brain Science*, M. Arbib (Ed.). MIT Press.

Turing, A.M. 1952. The chemical basis of morphogenesis. *Philos. Trans. R. Soc. London B*, 237.

Witkin, A., Kass, M. 1991. Reaction-diffusion textures. *Computer Graphics*, 25(3).

Wolpert, L., 1969. Positional information and the spatial pattern of cellular differentiation. *Journal of Theoretical Biology*, 25:1–47.

Yu, M., Wu, P., Widelitz, R.B., and Chuong, C.-M. 2002. The Morphogenesis of feathers. *Nature*, 420(21).

Wei-Min Shen is the Director of Polymorphic Robotics Laboratory (http://www.isi.edu/robots) at Information Sciences Institute (ISI), an Associate Director at the Center for Robotics and Embedded Systems, and a Research Assistant Professor in Computer Science at University of Southern California (USC). He received his Ph.D. under Nobel Laureate Professor Herbert A. Simon from Carnegie Mellon University in 1989. His current research interests include Self-Reconfigurable and Metamorphic Robotics, Artificial Intelligence, Machine Learning, and Life Science. He is the recipient of a Silver-Medal Award in 1996 AAAI Robotics Competition, a World Championship Award in 1997 Middle-sized RoboCup Competition, a Meritorious Service Award at ISI in 1997, and a Phi Kappa Phi Faculty Recognition Award at USC in 2003. He is the author of the book "Autonomous Learning from the Environment" published by W.H. Freeman in 1994. His research activities have been reported by media such as SCIENCE Magazine, CNN, PBS, Discovery channel, and other newspapers and magazines in the world. His research is sponsored by NSF, AFOSR, DARPA, and NASA

Peter M. Will is the Director of the Distributed Scalable Systems Division at USC/Information Sciences Institute and is a Research Professor in Industrial and Systems Engineering Department and Material Science Department at USC, and has over 35 years research experience in industry. He spent 16 years at IBM's Yorktown Research Lab, 7 years with Schlumberger, and 5 years at HP Labs. He has over 50 publications and 10 patents. He has served as chair of three NSF advisory committees and as Chair of the National Academy Study on Information Technology in Manufacturing. For six years he was a member of the ISAT group working with DARPA. In 1990, he was awarded the International Engelberger Prize in robotics. He received a B.Sc. degree in Electrical Engineering and a Ph.D. in non-linear Control Systems from the University of Aberdeen.

Aram Galstyan received his Ph.D. in physics from the University of Utah in 2000. He then joined USC Information Sciences Institute as a research associate. The main focus of his current research is emergent phenomenon in large scale multi-agent systems.

Cheng-Ming Chuong is a professor of pathology and Chair of Graduate Program for Experimental and Molecular Pathology in USC. He received his M.D. from Taiwan University in 1978, his Ph.D. from The Rockefeller University in 1983, working with Dr. Gerald Edelman on establishing the roles of cell-adhesion molecules in topobiology. Dr. Chuong directs Laboratory of Organ Development and Engineering (http://www-hsc.usc.edu/ cmchuong/index.html) that studies how stem cells are guided to form special tissues and organs of specific size and shape. Dr. Chuong has published more than 100 papers on the biology of integuments, and is the author of the book, "Molecular basis of epithelial appendage morphogenesis". He was honored by the invitation to give many lectures inside and outside US including John Ebling lecture in Europe and Don Orwin Lecture in Australia. He is an associate editor of J. Investigative Dermatology.