

Dispersion of a Team of Surveillance and Reconnaissance Robots Based on Repellent Pheromones

Janice L. Pearce, Bob Powers, Christopher Hess

Department of Mathematics and Computer Science, Berea College

Paul E. Rybski, Sascha A. Stoeter, Nikolaos Papanikolopoulos

Department of Computer Science and Engineering, University of Minnesota

Abstract

Search and rescue robots have become a viable alternative these days. Their mission success is highly correlated with their effective deployment. This paper looks at “repellent virtual pheromones” (inspired by insect colony coordination behaviors) to guide several miniature robots in order to be quickly deployed and dispersed. Based on simple vision-based control techniques, the objective is to cover an area of interest as fast as possible. We employ an overhead camera or a camera mounted on a command/control robot to provide each miniature robot with estimations of the positions of all of the other nearby robots in the robotic team. Each robot can then move away from the other nearby robots, resulting in the robot team swiftly dispersing through the local area. The approach been validated using the miniature Scout robots, developed by the Center for Distributed Robotics at the University of Minnesota. The Scout’s design is well-suited to surveillance and reconnaissance missions.

Index Terms: Visual Tracking, Mobile Robots, Robotic Teams, Robot Dispersion.

1 Introduction

With the capabilities of mobile robots improving dramatically and their cost dwindling, robot teams hold increasing promise for the successful execution of several challenging tasks associated with emergency response. For example, teams of flexible, reconfigurable, and inexpensive robots can be used to perform urban surveillance after a hurricane, execute remote operations in case of hazardous spills, participate in decontamination and decommissioning efforts immediately after a nuclear accident, accomplish search and rescue operations, and locate survivors in collapsed structures.

At the University of Minnesota, we have developed a prototype heterogeneous robotic team which can be used for a variety of missions. Larger “Ranger” robots which are capable of navigating long distances over rough terrain without needing to recharge their batteries are used

to traverse the environment and to deploy the miniature “Scout” robots (shown in Figure 1). Equipped with a magazine and a spring-based delivery mechanism, a Ranger can deploy up to ten Scouts into a target area. The Ranger’s powerful onboard computer can then be used to coordinate Scouts and relay status information [1].




Figure 1: Two Rangers with seven Scouts.

The need for a fast dispersal technique exists regardless of whether the coordinating agent is a human or a robot. When the coordinating agent is a human equipped with a wearable computer interface, commands can be sent to the robots from a hand-held controller and sensor information can be returned to the human through a head-mounted display. However, a human operator cannot be expected to simultaneously control many robots during the deployment, since the operator’s attention must be completely dedicated to the teleoperation of a single robot. Thus, the work described in this paper is designed to create a method for efficiently deploying and dispersing a team of robots regardless of the type of coordinating agent.

The paper first describes the repellent pheromones that our scheme employs. We then present an analysis of the results. Related work follows and the paper concludes with directions for future work.

2 Repellent Pheromones for Dispersion

Pheromones are chemicals employed in nature to provide indirect communication for the support of organized group activities. One may look at nature where ants leave a trail of pheromone to mark the path that they traverse between their nest and a food source. As more ants traverse this path, the pheromone trail is reinforced. The main objective this paper is to model “repellent pheromones” in order to bring about the dispersion of a robotic team. Many approaches to dispersion demand prior knowledge of the deployment area, but this approach has the advantage that it requires no map of the area and the robots need no self-knowledge of their location within the area. Other approaches to dispersion use a global geometric model, but our approach is based only upon decisions that are made locally. This has the distinct advantage of flexibility; suppose one robot loses its ability to travel, the robots in the local approach will automatically and continuously adjust to the given situation, while the system in the global approach would require a complete reworking.

We design the virtual pheromones to degrade as the distance from the virtual pheromone emitter increases. Since robots may be subjected to multiple virtual pheromone emitters from different directions, we model the cumulative repellent force of the virtual pheromones as a vector sum with the vector length given by a decreasing function. How quickly the virtual pheromones degrade as distance increases determines which decreasing function is employed. For our vector length, we choose the reciprocal of the distance the virtual pheromone travels. This choice mimics natural phenomena such as electric field strength, which is inversely proportional to the distance from a charged object. Simulations have been run by varying persistence of the pheromone over distance by varying the power on the reciprocal of the distance and runs have been carried out with unit power.

In particular, if \mathbf{x} is the position of robot R , then $p_i = \frac{1}{\|\mathbf{x} - \mathbf{x}_i\|^l}$ will model the level of virtual pheromone emitted by robot R_i that is detected by robot R . Here the positive number l is called the localization factor and is used to vary the persistence level of the virtual pheromone. Clearly, increasing l will cause the pheromone to have a more localized effect on nearby robots because the strength of the pheromone will fall off more rapidly over distance. The direction of this virtual pheromone is $\mathbf{d}_i = \frac{\mathbf{x} - \mathbf{x}_i}{\|\mathbf{x} - \mathbf{x}_i\|}$. Thus, the total repellent direction of all of the detected virtual pheromones is $\sum_i p_i \mathbf{d}_i = \sum_i \frac{\mathbf{x} - \mathbf{x}_i}{\|\mathbf{x} - \mathbf{x}_i\|^{(l+1)}}$.

With n robots in randomly distributed starting positions, iterating this algorithm infinitely many times on an infinite plane, will, in nearly all cases, direct the robots to positions that approximate the vertices of a regular

n -gon, thus asymptotically approaching the perfect circular sweep coverage of the area. In situations with more pathological starting positions, such as with three or more robots whose centroids are exactly in a line, the noise inherent in the physical system should cause one or more of the robots to move off of the line and then asymptotically approach perfect circular sweep coverage. In simulations as the localization parameter l is increased, robots achieve a given approximation to perfect circular sweep coverage more quickly. This can be seen by considering the standard deviations of the robots’ nearest neighbor distances.

In our simulation work, we employ velocity vectors for the trajectories of our simulated robots. Since velocity is the change in position over the change in time, we consider velocity as a derivative and approximate the derivative via Euler’s method. In particular, the velocity equations yield a set of differential equations that is solved to obtain the new position vector for each robot.

As expected, increasing the localization parameter l that represents the the virtual pheromone’s persistence over distance causes the dispersion of the robots to be decreasingly effected by robots that are farther away than one those nearer.

3 Results and Discussion

3.1 Color Tracking of Scout Robots

A robot, such as a Scout, which is small enough to avoid detection and still able to access hard-to-reach areas, is extremely useful for surveillance and reconnaissance applications, but the small size brings certain challenges to the task of using color markers for tracking. The Scout robot is cylindrical, measuring only 11 cm long and 4 cm in diameter. The electronics of Scouts include microcontrollers, transmitters, magnetometers, tiltometers, and shaft encoders. The Scout has differentially-driven wheels and a leaf-spring tail jumping mechanism. Scouts also carry a sensor payload, usually a miniature video camera, used to gather information that is broadcast over an analog RF transmitter. For color tracking, two color bands of approximately 1 cm in length encircle the two opposite ends of the deployed Scouts.

To track the relative locations and orientations of each Scout robot at the point of deployment, an overhead camera and standard vision techniques are used to assist in the dispersion of the robots. Thus, each of the Scout’s color markers is tracked as a color “blob.”

In order to achieve accurate tracking, we adopt the ActivMedia Color Tracking System (ACTS) [2], a commercial color segmentation software package from ActivMedia Robotics, for performing color thresholding and computing blob statistics (interface shown in Figure 2). ACTS

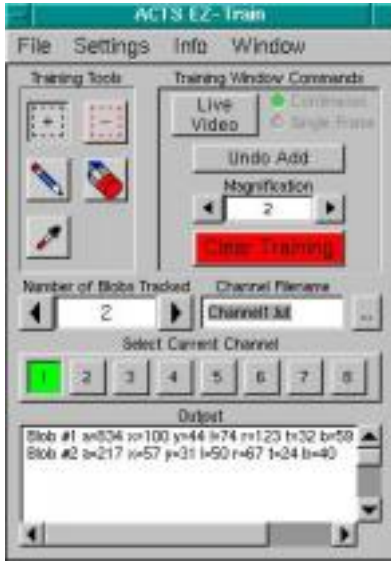


Figure 2: ACTS Interface.

allows users to select regions of color in an image not simply as ranges, but almost as finely as selected sets of points in RGB space. In fact, ACTS stores the projections of the RGB voxels onto the Red-Green, Blue-Green, and Red-Blue planes, giving far more flexibility than a simple sub-cube in RGB space. These color regions are then tracked as blobs, and statistics are computed for each blob. ACTS users can train color channels by selecting individual pixels of color from the image by clicking on a window to select the color markers directly. The ACTS-trained color channel files are stored as look-up tables. Using 160×120 pixel images, ACTS is optimized to track 32 individual color channels simultaneously at greater than 30 fps on a 160Mhz Pentium. Thus, its speed and number of channels are more than sufficient for our application since the limiting factor in our dispersion speed proved to be the RF command speed.

Using ACTS with a window size of only 320×240 pixels and the above-described technique for color-tracking gave us average standard deviations of Scout positions of less than 2 pixels and of orientations of less than 3 degrees for a set of up to eight selected colors. We elected to track only blobs larger than 4 pixels wide, since smaller blobs are likely to be simply noise due to such issues as the combined effects of fluorescent and natural lighting condition changes.

3.2 Issues in Radio-Controlling the Scout Robots

After implementing color-tracking and employing the virtual pheromone model to determine a Scout's trajectory, the next step is to direct the Scout's movement in the appropriate direction. Due to the small size and power constraints of the current version of the Scout robots, very limited on-board computational power is available, since they require their two CPUs for network communications

and actuator control. Thus, intelligence for control decisions must be provided externally. This implementation of dispersion of the Scout robots involves external visual servoing [3] and requires the auxiliary hardware of a computer equipped with a framegrabber to run the image-processing algorithms. In the field, this computer will be either a Ranger or another machine within the reception range of the Scout's analog video transmission.

To overcome the Scouts' size-imposed limitations and to connect multiple computers for complex missions, a distributed software architecture has been developed that supports the transparent integration of remote resources [4]. A functional view of missions is taken, so all hardware resources, including the robots, are partitioned into finely grained resources that can be requested by functional components.

This distributed software control architecture dynamically coordinates hardware resources and shares them between the various clients, allowing for simultaneous control of multiple robots. Each Scout has a unique network ID which allows commands to be routed to specific robots while being ignored by others, allowing multiple robots to be controlled from a single RF modem. Motion commands are then transmitted from a remote source and are received and executed by the Scout robot.

However, due to noise in the system such as radio interference, Scouts do not always receive the commands that are sent to them. Even when they do receive the intended commands, they may not receive the commands for the duration that the command was intended. These issues cause control difficulty, particularly in finely adjusting the pose of the Scout, since turning by varying degrees is accomplished by directing the Scout to rotate one or both wheels for a certain duration of time.

Orienting the body of a Scout so that it faces a target head on is an important task apart from the dispersion problem. In a variety of situations, control is needed over Scout alignment. Such scenarios include docking with a larger robot such as the Ranger for pickup, and using landmarks for tasks such as localization. Previous experiments in tracking Scouts have utilized pattern matching [5] and active contour models [6], but neither approach proved completely adequate.

Scouts are able to receive radio signals on two different radio frequencies. However, in this work, all of the Scout commands are sent on a single one of these two frequencies, because the single serial port of the processing computer is used to direct the commands to the radio. The speed of the radio commands proved to be the limiting factor in the speed of dispersion, because commands sent too quickly in succession interfered with each other, causing the Scout to fail to recognize the command sent. Thus, a pause of between 150 and 300 ms was added to

slow the command transmission rate. In addition, because commands to different Scouts are sent on the same frequency, only a single Scout is actually able to receive a command at a given time instant.

3.3 Dispersing the Scout Robots

Dispersion runs were completed using the virtual pheromone model described above with two, three, and four Scouts. Given the unreliability of the radio signal being correctly received, the dispersion results were satisfactory. In Table 1, results are shown as deviations from optimal beginning with the three Scout poses shown in Figure 3 ($(0^\circ, 0^\circ)$, $(0^\circ, 90^\circ)$, and $(0^\circ, 180^\circ)$). Error runs with more Scouts show similar errors.

Starting Poses	Average Turn Error	Communication Failures
$(0^\circ, 0^\circ)$	4.7°	1
$(0^\circ, 90^\circ)$	5.7°	0
$(0^\circ, 180^\circ)$	6.7°	0

Table 1: Dispersion errors with two Scouts from a given pose. Five runs were done for each starting pose. Communication failures indicate the number of times a Scout failed to receive commands and move during an experiment.

Figures 4(a), 4(b), and 4(c) show initial positions and dispersion paths taken by the Scouts during various runs. Once started, the Scouts followed the indicated paths until they left the field of view of the camera. One notes that while some of the Scouts choose their direction quickly and drive away, it is possible to see the adjustments of other Scouts as they correct their turns. In the run with four Scouts, it also is possible to note an error in the color tracking. The Scout in the upper right-hand corner of Figure 4(c) did not actually head the wrong direction and then retrace its path exactly. An error in the blob detection, occurring for a single frame, caused the Scout only to appear to traverse an incorrect dispersion path. Though such errors are infrequent, they are obvious when viewing the data visually. Even given these challenges, the trajectories of the Scouts create a reasonable dispersion pattern in each case.

4 Related Work


Team behaviors have been studied in a variety of disciplines from biological studies of herds and swarms to sociological studies of societies of humans. Physicists and chemists have studied the behaviors of a variety of interacting bodies: from gravitational planetary forces to the movements of various particles. Many of these studies of interactions in the natural world have become models for the behaviors of teams of robots, particularly as the robotic teams engage in tasks such as dispersing and attaining area coverage.

In 1992 Gage categorized the concept of “area coverage” by a robotic team into three basic types of coverage: “blanket coverage,” in which the main objective is to maximize the total detection area; “barrier coverage,” where the objective is to minimize the possibility of undetected penetration of a defined barrier; and “sweep coverage,” where the objective is to cover an area with a sweeping or moving barrier [7]. Using this taxonomy, the objective of the work described in this paper is to quickly deploy robots and achieve either a blanket or circular sweep coverage of an area.


In 1992 Gage also designed some robot coordination simulations, such as “condensation,” based loosely upon biosystem analogies such as pheromones [7]. In the early work by Arkin and Ali, the dispersion of a robotic team was carried out by a random-wandering behavior coupled with moderate robot repulsion as well as more significant obstacle repulsion [8]. In 1995, directly inspired by animal navigation routines, M. Mataric and her research group designed a dispersion algorithm that moves an agent away from the centroid of the local density distribution of the other agents that are visible to that agent’s sensors [9].

In 1999, Spears and Gordon provided distributed control of large collections of agents by having agents react to artificial forces motivated by natural laws of physics, observing that in the real physical world surprisingly complex behaviors arise from simple interactions between entities. However, their applications were self-assembly and self-repair rather than dispersion for the purpose of surveillance [10]. In another virtual physics approach, Howard *et al.* used a “potential-field-based approach” to the deployment of a mobile sensor network by treating their robots as virtual particles subjected to virtual forces [11]. These forces cause each given robot to be repelled from the other robots as well as from other obstacles in the environment with a potential that is proportional to the sum of the reciprocals of the distances from the first given robot. Though this portion of the algorithm is somewhat similar to the work presented in this paper, Howard *et al.* continued to run their algorithm until the whole network reached a static equilibrium, while in this paper after the initial dispersion, other robot behaviors such as locating a specific goal are allowed to operate.


The work that perhaps shares the most motivational similarities with the techniques described in this paper is the research of Payton *et al.* which employs techniques for coordinating the actions of large numbers of small-scale robots used in surveillance, reconnaissance, hazard detection, and path finding [12]. As in our project, they exploit the biologically inspired notion of a “virtual pheromone”, but with an implementation using transceivers mounted atop each robot rather than with global information from an overhead camera.



(a)




(b)




(c)


Figure 3: Scout poses of (a) $(0^\circ, 0^\circ)$, (b) $(0^\circ, 90^\circ)$, and (c) $(0^\circ, 180^\circ)$.



(a) Three robots



(b) Three robots



(c) Four robots

Figure 4: Example dispersion runs with Scout robots (axes in pixels). Initial starting positions of the robots are shown in the picture. The lines extending from the robots show the path the robot took until it left the field of view of the camera.

Batalin and Sukhatme also address the problem of multi-robot area coverage from the premise that local dispersion of robots will ultimately achieve good global area coverage [13]. As in this paper, their algorithms result in their robots being “mutually repelled” from one another, however like Payton *et al.*, they depend upon their robots to be able to sense or recognize one another rather than on global information such as from an overhead camera.

Stoeter *et al.* effectively use an overhead camera to track and direct a miniature Scout robot marked with color-markers as it travels, orients on a target, and climbs stairs. Though the extension to the problem of multiple robots is discussed in the article, the experiments were carried out with a single robot [14].

5 Conclusions and Future Work

This paper describes the use of repellent virtual pheromones in the dispersion of a team of miniature robots. It needs no prior map of the area and requires no localization, yet it leads to a reasonable dispersion for broadcast coverage, even when noise is present in the system. Though improvements can certainly be made in

future implementations, this technique can be used as is for the dispersion of a robot team.

In this paper, the Scout’s location and orientation are calculated from vision-analysis of the position of the colored markers, then the pheromones are modeled virtually. While the color-tracking analysis has proved quite accurate using the ACTS software, improved results might be achieved by applying a Kalman filter. The Scout’s color markers could be supplemented and/or replaced by colored wheels, possibly yielding an increase in tracking accuracy because the markers would be larger and farther apart. It would also allow observers on the ground to better track Scouts from the side. However, in place of the color-tracking implementation, the implementation of virtual pheromones by using a short-range transceiver should be considered even though it would require additional on-board power since it would offer the following additional benefits:

- Transceiver-implemented virtual pheromones would operate anywhere the Scouts were operating, so no overhead camera or Ranger-mounted camera would need to be present.

- When a Scout power supply becomes exhausted, the color-marked Scout still appears on camera, but a transceiver would stop broadcasting. Thus, a Scout that is “down” would disappear and other Scouts would move in to cover the area.
- Obstacles that block a Scout’s view would likely also block the transceiver signal, so coverage of areas with short obstacles would likely be improved.
- In addition to dispersion (and grouping), virtual pheromones implemented with a transceiver could be employed in additional applications such as finding a shortest path through a maze-like site.
- With more processing power available on-board and transceiver-implemented virtual pheromones, the decision-making of the next generation of Scout robots can be much more distributed, including dispersal without use of the communications channel.

Future work will expand on the Scout’s autonomous capabilities, which will include more advanced sensor interpretation and spatial reasoning techniques. The software control architecture is also being expanded to allow more types of hardware resources, such as larger robots, to be controlled.

6 Acknowledgements

This work has been in part supported by the National Science Foundation through grant #EIA-0224363, the Microsoft Corporation, and the Defense Advanced Research Projects Agency, Microsystems Technology Office (Distributed Robotics), ARPA Order No. G155, Program Code No. 8H20, issued by DARPA under Contract #MDA972-98-C-0008.

References

- [1] P. E. Rybski, S. A. Stoeter, M. D. Erickson, M. Gini, D. F. Hougen, and N. Papanikolopoulos. A team of robotic agents for surveillance. In *Proc. of the Int’l Conf. on Autonomous Agents*, pages 9–16, Barcelona, Spain, June 2000.
- [2] P. E. Rybski. *ACTS: Activmedia Color Tracking System Manual*. ActivMedia Robotics, LCC, 44 Concord Street, Peterborough, NH 03458, USA, 2002.
- [3] S. Hutchinson, G. D. Hager, and P. I. Corke. A tutorial on visual servo control. *IEEE Transactions on Robotics and Automation*, 12(5), 1996.
- [4] S. A. Stoeter, P. E. Rybski, M. D. Erickson, M. Gini, D. Hougen, D. G. Krantz, N. Papanikolopoulos, and M. Wyman. A robot team for exploration and surveillance: Design and architecture. In *Proc. of the Int’l Conf. on Intelligent Autonomous Systems*, pages 767–774, Venice, Italy, July 2000.
- [5] P. Rößler, S. A. Stoeter, P. E. Rybski, M. Gini, and N. Papanikolopoulos. Visual servoing of a miniature robot toward a marked target. In *Proc. of IEEE International Conference on Digital Signal Processing*, July 2002.
- [6] D. P. Perrin, E. Kadioglu, S. A. Stoeter, and N. Papanikolopoulos. Localization of miniature mobile robots using constant curvature dynamic contours. *Proc. of the IEEE Int’l Conference on Robotics and Automation*, May 2002.
- [7] D.W. Gage. Command control for many-robot systems. In *AUVS-92, the Nineteenth Annual AUVS Technical Symposium. Reprinted in Unmanned Systems Magazine, Fall 1992, Volume 10, Number 4, pp 28-34.*, Huntsville AL, 22-24 June 1992.
- [8] R.C. Arkin and K. Ali. Integration of reactive and teleroctic control in multi-agent robotic systems. In *Proc. Third International Conference on Simulation of Adaptive Behavior, (SAB94) [From Animals to Animats]*, pages 473–478, Brighton, England, August 1994.
- [9] M. J. Matarić. Designing and understanding adaptive group behavior. *Adaptive Behavior*, 4(1):51–80, 1995.
- [10] W. Spears and D. Gordon. Using artificial physics to control agents. In *IEEE International Conference on Information, Intelligence, and Systems*, 1999.
- [11] A. Howard, M.J. Matarić, and G.S. Sukhatme. Network deployment using potential fields: A distributed, scalable solution to the area coverage problem. In *The 6th International Symposium on Distributed Autonomous Robotic Systems*, pages 299–308, Fukuoka, Japan, June 25-27, 2002.
- [12] D. Payton, M. Daily, R. Estkowski, M. Howard, and C. Lee. Pheromone robotics. *Autonomous Robots*, 11(3):319–324, Nov 2001.
- [13] M. Batalin and G.S. Sukhatme. Spreading out: A local approach to multi-robot coverage. In *The 6th International Symposium on Distributed Autonomous Robotic Systems*, pages 373–382, Fukuoka, Japan, 2002.
- [14] S. A. Stoeter, P. E. Rybski, M. Gini, and N. Papanikolopoulos. Autonomous stair-hopping with scout robots. In *Proc. of the IEEE/RSJ Int’l Conf. on Intelligent Robots and Systems*, 2002.