

# THE UTILITY OF HETEROGENEOUS SWARMS OF SIMPLE UAVS WITH LIMITED SENSORY CAPACITY IN DETECTION AND TRACKING TASKS

Matthias Scheutz\*, Paul Schermerhorn\*, and Peter Bauer\*\*

(\*) Department of Computer Science and Engineering

(\*\*) Department of Electrical Engineering

University of Notre Dame

Notre Dame, IN 46556, USA

{mscheutz,pscherm1,pbauer}@nd.edu

## ABSTRACT

We present a physically realizable UAV model for locating and tracking chemical clouds. Simulation results are presented for implementations of this model with two configurations, one that is faster and requires more space to avoid collisions, and one that is slower and can cover an area more densely. Heterogeneous swarms of agents are shown to have better performance than homogeneous swarms of similar size because they take advantage of the strengths of each configuration.

## 1. INTRODUCTION

Biologically inspired swarms of autonomous agents have been used successfully in a variety of applications ranging from various kinds of ground-based robots, to unmanned aerial vehicles (UAVs). The agents employed in swarm systems are typically governed by simple rules and coordinate their behavior based on simple interaction principles that use, among other factors, the distance between all agents within a given local neighborhood (e.g., variants to Reynold's three rules "flock centering", "obstacle avoidance", and "velocity matching" [3]). Theoretical results have been proved about various properties of swarms governed by such rules (e.g., collision avoidance or stability [4, 5]).

From a practical perspective, however, some assumptions about the availability of information about the swarm are idealized and difficult to meet in implementation. In particular, it might not be possible to get exact distance readings for all agents within a given neighborhood (e.g., because of occlusion effects, or simply because sensors that could provide that information are too expensive or complex to be used in the agent implementation, lack of GPS, etc.). Moreover, it might not be possible to determine the overall goal direction (e.g., because the goal cannot be sensed at a distance). Finally, theoretical investigations typically limit swarms to homogeneous groups, while it might be advanta-

geous to use heterogeneous groups, e.g., because they give rise to a more robust system, improve overall task performance, or require fewer resources.

In this paper we address the above concerns and investigate the utility of heterogeneous swarms of extremely simple, physically realizable agents for a class of *local detection and tracking tasks*. Potential field-based goal detection and tracking methods are not directly applicable in this case because the objects to be tracked cannot be sensed at a distance (which is a prerequisite for standard potential-field based navigation). As a specific instance of this class of tasks, we consider a "chemical detection and tracking scenario", in which a cloud of chemicals of a particular kind with unknown density, extension, heading and dispersion over time has to be detected and enclosed by unmanned aerial vehicles without collisions among UAVs. Each UAV has an on-board sensor that can detect chemicals of this type whenever they come in touch with the sensor surface, but cannot detect chemicals otherwise. Consequently, it is not possible for any UAV to detect the chemical cloud at a distance using its chemical sensors.

Our solution to overcome this problem is to equip each UAV with an "attractive beacon" that it will turn on whenever it detects a chemical, thus subsequently attracting other UAVs to that area. The task is then accomplished in two phases: (1) a detection phase where UAVs swarm, find the cloud and once one UAV has detected it others will move toward it (top in Fig. 1), and (2) an enclosure and tracking phase, where the UAVs cover the cloud area as tightly as possible and move along with it (bottom in Fig. 1).

A related study employs biologically-inspired plume-tracking mechanisms to allow robots to track odor plumes [1, 2]. In some configurations, the robots work together in a manner similar to that described below, with a robot detecting an odor plume able to attract other robots to the plume. One main difference between this work and the present study is that the robotic agents use scent gradients,

whereas the UAVs described below sense only in direct contact with chemicals. The UAVs then impose an artificial gradient via the attraction beacons to make up for not being able to sense chemicals out of direct contact. Moreover, this study examines the benefit of heterogeneous swarms, whereas the robots used in the scent plume tracking studies were all configured identically.

The rest of the paper details our approach. First, we introduce the proposed control mechanisms for the UAVs. Then we briefly discuss the implications of these controllers for the whole UAV swarm and the resultant *emergent properties*, including the motivation and justification for considering heterogeneous swarms as well as the limitations of the approach. The experimental section then compares homogeneous and heterogeneous swarms to demonstrate that (1) the proposed heterogeneous swarm is capable of detecting and tracking the chemical cloud, and (2) that it performs significantly better in the tracking task than homogeneous swarms of identical size. The final discussion briefly analyzes the results and provides an outlook for future work.

## 2. THE UAV MODEL AND CONTROL

Since we are ultimately interested in swarm systems that can be directly employed to solve engineering problems, we will start with the restriction that sensors, effectors, and the employed control of each swarm agent should be realistic, simple, and easily realizable. Moreover, we are interested in swarm solutions that exploit emergent group behavior as much as possible and trade it for individual complexity (i.e., the simpler the individual agent while keeping the overall task performance of the group approximately the same, the better).

Each UAV is equipped with two types of beacons, a collision avoidance beacon *col* (to repel UAVs from each other) and an attractive signal detection beacon *sen* that is only activated when the UAV detects a chemical. The collision avoidance beacons are always on and each UAV is equipped with an antenna pair to detect the collision avoidance beacons of other UAVs. In the same way, each UAV is also equipped with an antenna pair to detect the attraction beacons of other vehicles. In essence attraction beacons simply translate the detected signal into a on/off radio beacon and thus boost its strength. In contrast to the attraction beacon receivers, the collision avoidance receivers can extract the distance of the source using signal strength. This allows the collision avoidance algorithm to react only to UAVs within a certain circle of size  $\alpha$ . This  $\alpha$  is the UAV's collision avoidance range, and represents the distance a UAV must keep between it and its neighbors to leave enough space to turn. Therefore,  $\alpha$  is dependent on the UAV's turning radius, which, in turn is dependent on the UAV's speed; given a speed  $s$ , a UAV can turn only so fast.

We are now ready to provide some definitions for the control algorithm:

$$I_{y,i} = \sum_{j=1}^n A_y / (\| \underline{x}_j - \underline{x}_i \|_2^2)$$

where  $I_{y,i}$  is the received power of beacon type  $y$  at UAV  $i$  at location  $x_i$  with  $y \in \{col, sen\}$ , and  $A_y$  is the power of UAV beacon of type  $y$ . (There are  $n$  UAVs, and all beacons of the same type have the same power.) Including the directional sensitivity of two side looking directional antennas, we can find the following signal intensity for left and right looking antennas of each of the two modalities:

$$R_{y,i} = \sum_{j \in \Gamma_y} A_y f(\underline{x}_j - \underline{x}_i, \underline{\eta}_i) / (\| \underline{x}_j - \underline{x}_i \|_2^2)$$

$$L_{y,i} = \sum_{j \in \Gamma_y} A_y f(\underline{x}_j - \underline{x}_i, -\underline{\eta}_i) / (\| \underline{x}_j - \underline{x}_i \|_2^2)$$

with  $\Gamma_{col} = \{j \mid \| \underline{x}_j - \underline{x}_i \|_2 < \alpha\}$ ,  $\Gamma_{sen} = \{1, \dots, n\}$ ,  $\underline{\eta}_i$  being the right normal vector to the speed vector of UAV  $i$  in the flight plane, and  $f(x, \eta)$  being the directional sensitivity function of the antenna, where  $x$  is the vector from receiver to source and  $\underline{\eta}$  is the direction of highest sensitivity of the antenna.

In the case of  $y = col$ , the summation for the left and right antenna signal intensity  $L_i$  and  $R_i$  respectively are taken only over those UAVs  $j$  that satisfy  $\| \underline{x}_j - \underline{x}_i \|_2 < \alpha$ . This requires certain provisions in the modulation scheme that allow the UAV to distinguish each collision beacon. (In the attraction beacon receivers, distinguishing between different UAV beacons is not necessary and the above summation can actually be done by the antenna itself rather than in digital hardware.)

In order to decide in which direction to turn, one requires the two directional antennas on each side of the UAV facing in opposite directions ( $\underline{\eta}$ , and  $-\underline{\eta}$ ), but in direction perpendicular to the vehicle speed vector. Since the turn radius of the UAV is assumed to be constant in both directions (left and right), a simple intensity comparison between left and right directional antenna will allow to derive the new heading of the UAVs, which is either "turn left" or "turn right".

Now define the intensity sum and difference between the antenna pairs as:  $L_{y,i} + R_{y,i} = S_{y,i}$ ,  $L_{y,i} - R_{y,i} = D_{y,i}$ ,  $y \in \{col, sen\}$ ,  $i = 1, \dots, n$ , where  $S_{y,i}$  and  $D_{y,i}$  denote sum and difference of left and right antenna signal strength of modality  $y$  at UAV  $i$ .

The UAV guidance system implementation is shown in Figure 2. It consists of two power complementary antennas that are used for both, the sensor and the collision beacons. The left and right collision beacon receivers demodulate the unique collision beacon signals and provide at the output

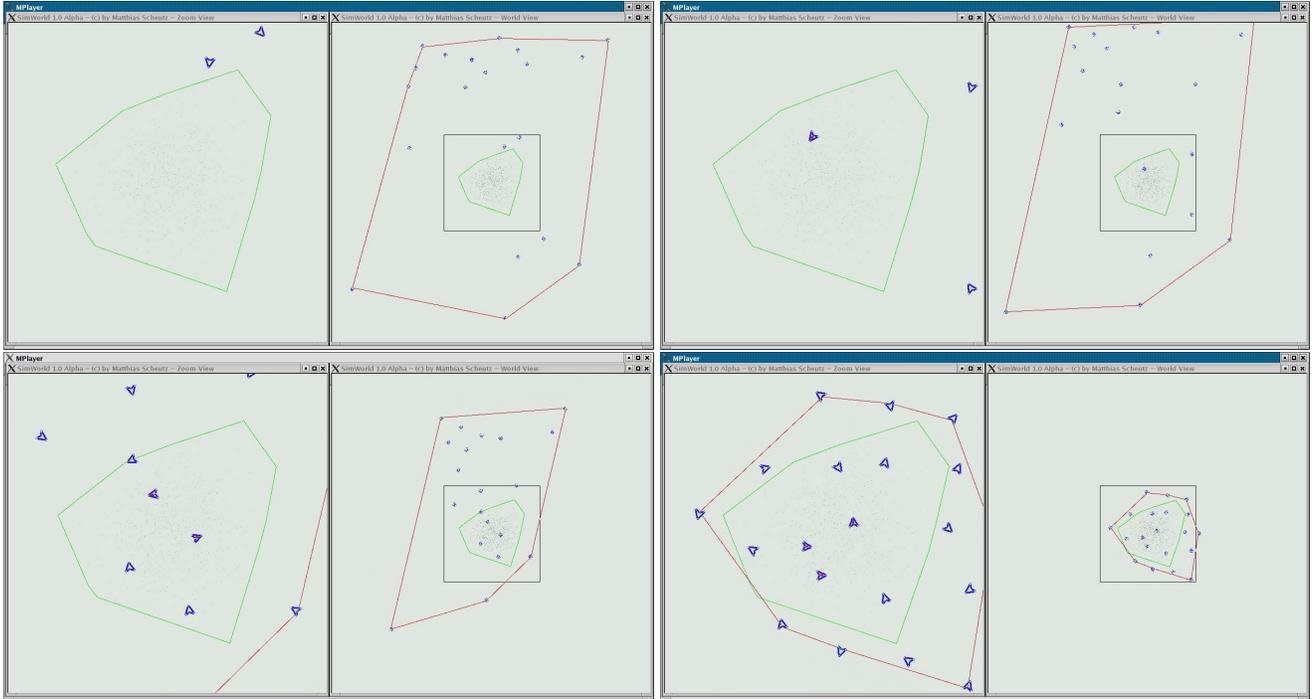


Figure 1. Different phases in the chemical detection and tracking task: searching for the cloud (top left), at least one UAV detects chemicals and turns on its beacon (top right), thus attracting others to the area (bottom left), which is subsequently tightly enclosed (bottom right). Each of the four pictures shows a magnified version on the left of the environment on the right (square on the right) as well as the bounding polygons of the chemicals (green) and the UAVs (red).

the detected signal strength for each of the  $n$  UAVs. Since the collision beacon strength of all UAVs are equal, and in free space the power drops with the square of the distance, one can easily determine the threshold for the received collision beacon signal that corresponds to the critical distance  $\alpha$ . This threshold is implemented in the two threshold filters, which pass the collision beacon signal of UAV  $i$  only if it satisfies the critical distance (received power) requirement. All the UAV collision beacon contributions that pass the threshold test are then added for the left and right side, and in a second step subtracted to find which side produces a stronger signal. This difference is then converted into a left or right turn signal using a (signum) hard limiter nonlinearity. The sensor beacon signals are compared in a similar way. However, in this case, all UAV contributions are used and are summed up by the antennas themselves. If the left plus right antenna signal strength passes the given threshold, the left and right turn signal is generated in the same way as in the collision beacon case. The channel selector is a priority selector, that selects channel 0 with highest priority (collision beacon generated command signals) if it is enabled, then channel 1 with second highest priority (sensor beacon generated command signals) and finally with lowest priority channel 2 which is the alternating left right sig-

nal. This signal results (after the compensator) in a constant heading. This figure only shows guidance mechanisms in the 2-D plane. Altitude control and the thresholding circuit for activating the sensor beacon are not shown.

## 2.1. Agent Model

Given the above mechanisms, we define an agent model that implements the control algorithm in Figure 3. During the sensing phase, agents detect the presence of chemicals, sum the attraction vectors of other agents' attraction beacons, and identify the collision vector of the nearest agent within collision avoidance range. These comprise the sensory information  $S$  used by the control system:

- Rule 1:** If a chemical is detected, activate attraction beacon
- Rule 2:** If no chemical is detected, deactivate attraction beacon
- Rule 3:** If some other agent is within collision avoidance range, turn away from nearest agent, reset  $rwc$  to 400
- Rule 4:** If no other agent is within collision avoidance range and beacons are detected, turn in the direction of the beacon vector, reset  $rwc$  to 400

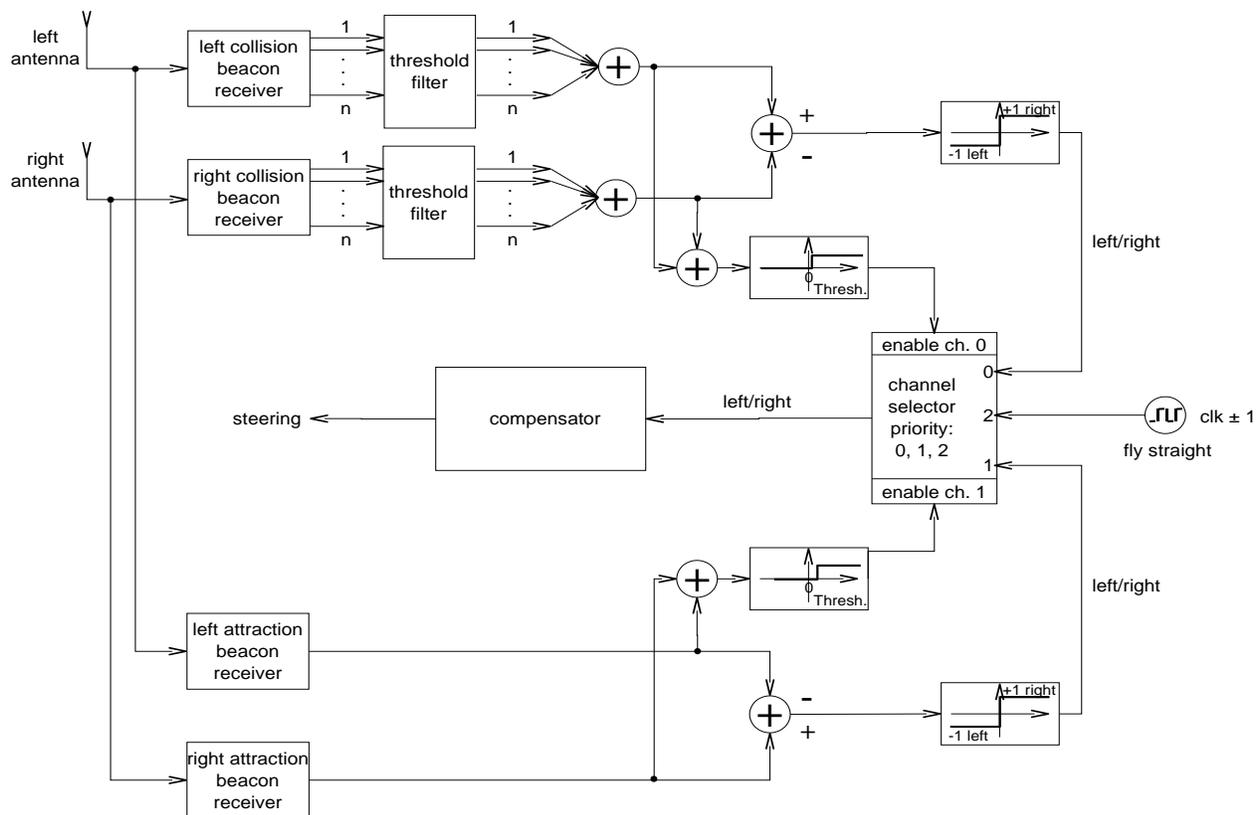


Figure 2. The UAV guidance system.

**Rule 5:** If no other agent is within collision avoidance range, no beacons are detected, and  $rwc > 0$ , decrement  $rwc$  and turn right if  $rwc$  is even, left if  $rwc$  is odd

**Rule 6:** If no other agent is within collision avoidance range, no beacons are detected, and  $rwc = 0$ , reset  $rwc$  to  $100 * \text{random}(4)$ , and turn 75–100 degrees (randomly selected) right or left (randomly selected)

These rules cause the agent to turn away from the nearest agent, if there is one within collision avoidance range, or turn toward the highest intensity of attraction beacons, if any are detected, otherwise patrol the environment randomly as follows: the agent is equipped with a counter  $rwc$  that determines how long it will proceed on a given heading (via alternating turns) before choosing a new heading. When  $rwc$  reaches 0, the agent randomly selects a left or right turn of 75–105 degrees and resets the counter. This can help to avoid simulation runs in which agents never detect a chemical and simply diverge outward, away from the chemical cloud; eventually, some agent may turn back toward the cloud, although the simulation may terminate before any agent makes it back to the cloud.

### 3. SIMULATION EXPERIMENTS

#### 3.1. Experimental Setup

All simulations were conducted using the artificial life simulator SWAGES under development in our lab.<sup>1</sup> The world in which the agents operate is an unbounded 2000 by 2000 square 2D surface. The number of agents ( $A$ ) was varied from 5 to 30 in steps of 5. In homogeneous experiments, all agents are configured with  $\alpha = 180$  and  $speed = 10$ . In heterogeneous experiments,  $m$  agents are configured this way, while the remaining  $n$  agents are configured with  $\alpha = 120$  and  $speed = 5$ . We will refer to these as  $Alpha_{180}$  and  $Alpha_{120}$  agents, respectively.

Each experiment consists of 40 experimental runs, each of which begins with a distinct set of initial conditions (e.g., placement of agents, their headings, etc.). These initial conditions are consistent across all experiments, so that, for example, the initial placement of the chemicals is the same for experimental run 1 of the 15 homogeneous  $Alpha_{180}$  agent experiment as it is for experimental run 1 of the 15 hetero-

<sup>1</sup>SWAGES is an agent-based artificial life simulation experimentation environment freely available at <http://www.nd.edu/~airolab/software/>.

```

if  $S_{col,i} > 0$  then
  if  $D_{col,i} > 0$  then
    turn right
  else
    turn left
  end if
else if  $S_{sen,i} > threshold$  then
  if  $D_{sen,i} > 0$  then
    turn left
  else
    turn right
  end if
else
  fly straight (by alternating right and left)
end if

```

Figure 3. The UAV direction control algorithm.

geneous agent (10  $Alpha_{180}$  and 5  $Alpha_{120}$ ) experiment. This allows us to directly compare performance, since we know that the tasks were the same for all combinations of agents.

Over the course of the 10,000 cycles for each experiment, we tracked the area of the convex hull of the UAV agents, as well as the area of the convex hull of the chemicals. This allows us to see how tightly the UAVs surround the chemicals. At times, however, the UAVs may not completely contain the chemical cloud, so we also tracked the number of chemicals contained inside the convex hull of the UAV agents. The combination of these two measures gives us a complete view of how well the UAVs are surrounding the cloud. A good solution is defined as one in which the difference between the UAV area and the chemical cloud area is minimized and the number of chemicals enclosed by the UAV convex hull is maximized. It is easy to get complete enclosure of the cloud simply by expanding the UAV formation indefinitely. In fact, there are cases, especially when there are few agents, in which that happens. The UAVs miss the chemical cloud entirely and the area enclosed by the UAV convex hull increases. We characterize any experimental run in which the UAV area at the end of the simulation is greater than the 2000 by 2000 simulation area as a failure.

### 3.2. Results

The number of failures for each experiment is given in Table 1. When a simulation is classified as failed, its data is not included in the figures below. Practically speaking, when an experimental run fails, the area of the UAVs is so large that it dominates the areas of those that succeed, and the figure would depict little other than the failure if it were included. For this reason, we report the failure rates and ignore failed experimental runs in the remaining results.

The results indicate a tendency for failure that is more

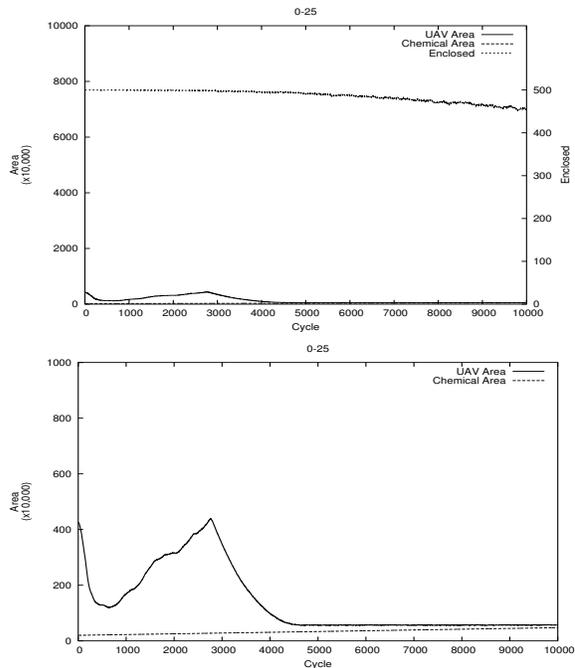


Figure 4. Homogeneous (25  $Alpha_{120}$ ) configuration. The top graph presents the average area enclosed by the UAVs in successful experimental runs, along with the number of chemicals enclosed by the UAVs. The bottom graph zooms in on the lower portion of the top graph, allowing a comparison with the area of the chemical cloud.

pronounced in homogeneous experiments, particularly homogeneous  $Alpha_{180}$ , than in experiments with both agent types. There were a small number of failures of heterogeneous groups with small total size, but other than these, heterogeneous experiments did not fail. The homogeneous  $Alpha_{180} = 5$  and  $Alpha_{180} = 10$  experiments failed to produce even one success.

The area results presented in Table 1 consist of two parts: the average final area covered by the UAVs in the experiment and the average maximum UAV area. The final UAV area is the area to which the convex hull of the UAVs converged. This can be compared to the final cloud area to determine how close the fit is (i.e., how well the UAVs achieved the first criterion, minimizing the difference between UAV area and cloud area).

Also included in Table 1 are measures of the minimum and final chemical enclosures. The cloud consists of 500 chemicals, so configurations that approach this (e.g., homogeneous  $Alpha_{180} = 30$  configurations) do well on this component of performance.

This, however, does not tell the entire story. In some cases, the UAVs may converge on the cloud only after diverging broadly. Figures 4, 5, and 6 compare the “histories”

UAV Ratio	Failures	Max. UAV Area	Final UAV Area	Final Cloud Area	Min. Cloud Enclosure	Final Cloud Enclosure
0:5	21	3853130.0	58320.9	477473.0	103.895	128.579
0:10	3	7101330.0	172445.0	474930.0	290.459	292.811
0:15	0	7044920.0	299943.0	476834.0	378.25	387.45
0:20	0	7639880.0	443059.0	477258.0	425.3	432.6
0:25	0	4397880.0	580109.0	475186.0	451.85	456.625
0:30	0	4554040.0	697956.0	477175.0	471.025	471.85
5:0	40	-	-	-	-	-
5:5	9	6030390.0	322516.0	472572.0	346.806	353.065
5:10	2	17785300.0	503823.0	475191.0	414.342	427.737
5:15	0	6318570.0	645109.0	476675.0	449.325	449.325
5:20	0	4287510.0	797641.0	475490.0	470.475	476.75
5:25	0	4580980.0	953815.0	477175.0	479.1	479.75
10:0	40	-	-	-	-	-
10:5	2	31505800.0	654625.0	477278.0	445.158	449.395
10:10	1	21042600.0	796044.0	476286.0	468.205	473.897
10:15	0	10100300.0	962614.0	475517.0	481.925	483.2
10:20	0	4629140.0	1129680.0	477175.0	487.1	487.725
15:0	15	94065100.0	818284.0	483257.0	451.48	466.52
15:5	0	16029600.0	999767.0	477501.0	476.4	483.45
15:10	0	16637000.0	1128350.0	475497.0	486.375	487.675
15:15	0	4646460.0	1326000.0	477175.0	491.475	491.925
20:0	12	43254000.0	1225820.0	480183.0	478.143	491.75
20:5	1	4396200.0	1319350.0	473228.0	488.795	493.41
20:10	0	4694200.0	1462930.0	477175.0	494.25	495.375
25:0	2	16174600.0	1526080.0	472502.0	493	494.237
25:5	0	4734490.0	1653220.0	477175.0	495.15	496.325
30:0	2	11399300.0	1899280.0	476831.0	496.105	497.842

Table 1. Overall Results

of three configurations with similar final UAV Areas and Cloud Enclosures (averaged over successful experimental runs). Figure 4 depicts the homogeneous  $Alpha_{120} = 25$  configuration. This configuration has good absolute performance, according to the criteria mentioned before: minimizing the difference between UAV area and cloud area and maximizing the number of chemicals enclosed by the UAVs. There are configurations that do better on each count, but this has a nice balance of both. The progression of the task is also very good, with the UAV area never growing much beyond its initial size. It converges quickly and remains tight.

Turning to Figure 5, which depicts heterogeneous ( $Alpha_{180} = 10$ ,  $Alpha_{120} = 5$ ) experiments throughout the 10,000 cycles, we find the UAV area expanding substantially before converging on the chemical cloud. The final values for UAV area and chemical enclosure are similar to the homogeneous  $Alpha_{120} = 25$  configuration. While the behavior of the heterogeneous configuration is more unruly, it requires only 15 total agents, compared to the 25 in Figure 4.

Finally, for the sake of comparison, Figure 6 presents the same history for homogeneous  $Alpha_{180} = 15$  simulations. In this case, UAV area diverges significantly before converging near the end of the simulation duration. The final area is somewhat worse, but the enclosure is correspondingly better. However, as shown in Table 1, 15 of these experimental runs failed to complete.

In all cases, the cloud enclosure value decreases throughout the simulation. The chemical cloud tends to become elongated as it travels, whereas the UAVs tend to maintain a roughly circular formation, making it difficult to enclose all chemicals without exaggerating the UAV area. Including a drifting influence in the UAV control system similar to that of the cloud may be one way to help the UAVs match more closely the area of the cloud.

### 3.3. Analysis

Overall, the homogeneous  $Alpha_{180}$  configurations are unstable compared with heterogeneous configurations. They tend to fail more and their average maximum UAV area is

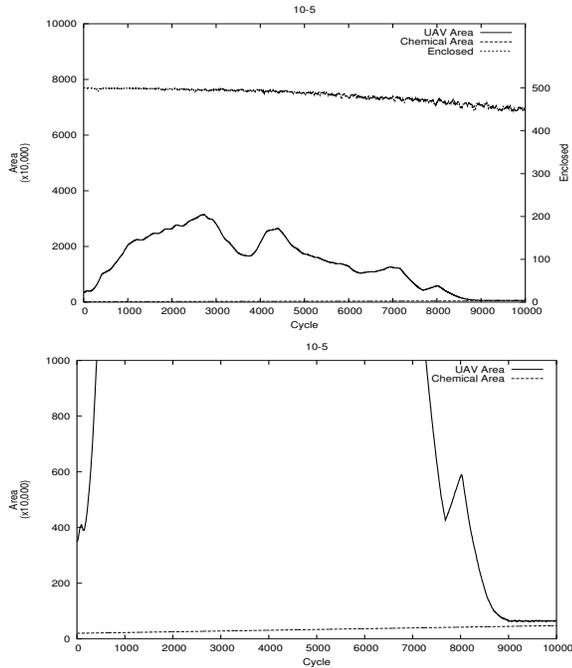


Figure 5. Heterogeneous (10  $Alpha_{180}$  and 5  $Alpha_{120}$ ) configuration. The top graph presents the average area enclosed by the UAVs in successful experimental runs, along with the number of chemicals enclosed by the UAVs. The bottom graph zooms in on the lower portion of the top graph, allowing a comparison with the area of the chemical cloud.

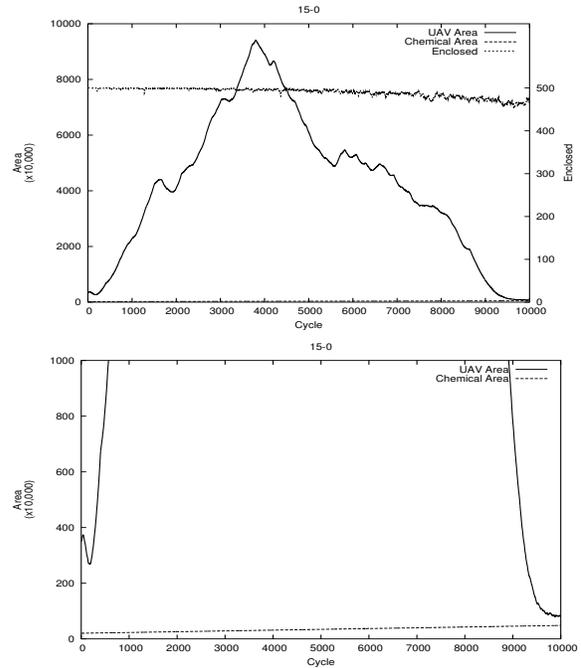


Figure 6. Homogeneous (15  $Alpha_{180}$ ) configuration. The top graph presents the average area enclosed by the UAVs in successful experimental runs, along with the number of chemicals enclosed by the UAVs. The bottom graph zooms in on the lower portion of the top graph, allowing a comparison with the area of the chemical cloud.

higher. Their enclosure is better than heterogeneous configurations with the same number of UAVs, but that is more a reflection of the greater area covered by the  $Alpha_{180}$  agents due to their higher  $\alpha$  than to a greater ability to track the cloud. Homogeneous  $Alpha_{120}$  configurations tend to fail less often, when given sufficient agents. The reason  $Alpha_{120}$  agents are less susceptible to losing the cloud is partly their lower speed. The random patrol mechanism described above induces agents to change heading after a set number of cycles. The faster  $Alpha_{180}$  agents will be able to move farther away from the chemical cloud in the course of those cycles, making it more difficult for them to return.  $Alpha_{120}$  agents tend to stay closer to home, which helps them locate the cloud, even when they have temporarily lost it. The reduced speed of  $Alpha_{120}$  agents is needed to avoid collisions, given their smaller collision range  $\alpha$ , but it leads to an unintentional benefit. Slowing down the  $Alpha_{180}$  agents may produce similar effects, however, there is a minimum speed required to remain airborne, so this cannot be decreased arbitrarily.

Homogeneous  $Alpha_{180}$  configurations tend to enclose the chemical cloud better than homogeneous  $Alpha_{120}$  con-

figurations of the same size, by virtue of their higher  $\alpha$ ; they cannot pack the area as tightly as the  $Alpha_{120}$  agents, and therefore require fewer agents to enclose the same area.

Combining the two in heterogeneous configurations takes advantage of the strengths of both: the  $Alpha_{120}$  agents remain close to home, and are more likely to locate the chemical cloud. Once the cloud is detected, all agents converge on it, and the more spread-out  $Alpha_{180}$  agents contribute greater average area enclosed. At this point, an interesting emergent behavior is seen (Figure 7): the  $Alpha_{120}$  agents cluster in the middle of the UAV group, and the  $Alpha_{180}$  agents surround them. The smaller  $\alpha$  allows  $Alpha_{120}$  agents to “squeeze” in between  $Alpha_{180}$  agents, making it easier for them to get to the center of the cloud, whereas  $Alpha_{180}$  agents very seldom get into the center because they turn away from other agents sooner.

The selection of a best overall configuration is subjective; the area fit and the enclosure must be good, and they must find and converge on the cloud in a reasonable amount of time with few failures. Moreover, given similar absolute performance, a configuration that deploys fewer UAVs is preferable, for reasons of cost. A more objective selection would require an assignment of relative importance to

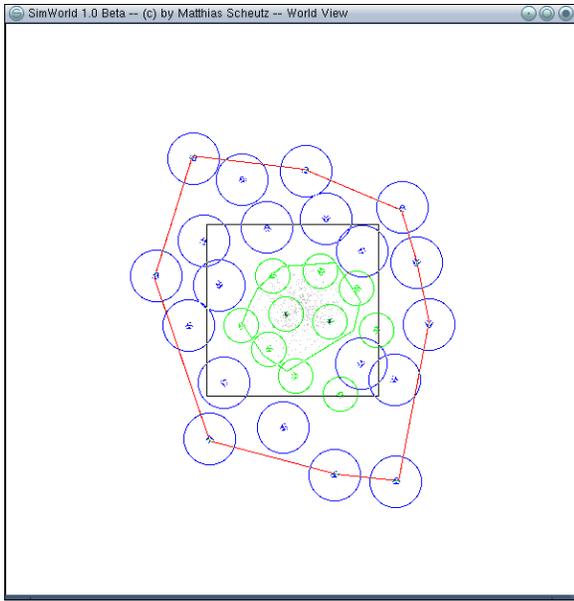


Figure 7. Heterogeneous behavior: Slower agents with smaller collision avoidance ranges tend to cluster in the middle of the UAV formation.

the two objectives. If, for example, the most important objective were to enclose the entire cloud, larger UAV areas would be (more) acceptable. The relative importance of the objectives will depend on the task (e.g., the nature of the chemicals in the cloud), so no assignment has been attempted here.

#### 4. CONCLUSION

This paper presents a simple UAV agent model that is capable of locating and tracking a target chemical cloud. We have taken care to ensure that the model is simple and physically realizable. While agents are limited to sensing only chemicals with which they are in direct contact, they collectively impose upon the chemical cloud an artificial gradient of attraction beacons (activated when a chemical is sensed) that can be detected remotely.

Simulation results show that swarms of these agents are effective at locating and tracking a slowly-moving chemical cloud. Heterogeneous swarms demonstrate certain advantages over homogeneous swarms. First, heterogeneous systems are more robust than homogeneous systems of agents with higher speed and larger collision avoidance range, because the latter are susceptible to losing the chemical cloud and spreading out indefinitely through the environment with little chance of reacquiring the location of the cloud. Second, heterogeneous systems demonstrate better enclosure behaviors than homogeneous systems with lower speed and

smaller collision range, because the latter tend to cluster closely together in the center of the cloud. Heterogeneous systems utilize the advantage of the slower agents to locate the chemical cloud, and the advantage of the faster agents to enclose it. This allows heterogeneous systems to achieve good performance with fewer agents than either homogeneous configuration.

For future work, we plan to look into the effect of reducing the speed of the  $Alpha_{180}$  agents; they may be able to maintain their good coverage properties while improving their ability to locate the cloud quickly. However, the larger  $\alpha$  will lead to fewer UAVs inside the cloud, in turn leading to fewer chemical detections. With fewer attraction beacons, the likelihood of losing the cloud will increase; these tradeoffs should be examined. Similarly, adaptive control systems that adjust their speed and  $\alpha$  could address the issue of the cloud growing as it disperses. While one option would be to deploy more UAVs as the cloud gets larger, it may also be feasible to simply increase  $\alpha$ , thereby increasing the area enclosed.

#### 5. REFERENCES

- [1] A. T. Hayes, A. Martinoli, and R. M. Goodman, "Distributed Odor Source Localization," IEEE Sensors, Special Issue on Artificial Olfaction, Vol. 2, No. 3, pp. 260-271, 2002. IEEE Press.
- [2] Kazadi S., Goodman R., Tsikata D., Lin H., "An Autonomous Water Vapor Plume Tracking Robot Using Passive Resistive Polymer Sensors," Autonomous Robots, Vol. 9, No. 2, pp. 175-188, 2000. Kluwer Academic Publishers
- [3] C. W. Reynolds. Flocks, herds, and schools: a distributed behavioral model. Computer Graphics (ACM SIGGRAPH 87 Conference Proceedings), 21(4):25-34, July 1987
- [4] Herbert G. Tanner, "Flocking with Obstacle Avoidance in Switching Networks of Interconnected Vehicles," 2004 IEEE International Conference on Robotics and Automation, New Orleans LA, April 26-May 1, pp. 3006-3011.
- [5] Herbert G. Tanner, Ali Jadbabaie and George J. Pappas, "Stable Flocking of Mobile Agents, Part I: Fixed Topology," 42nd IEEE Conference on Decision and Control, Maui Hawaii, December 2003, pp. 2010-2015.
- [6] Herbert G. Tanner, Ali Jadbabaie and George J. Pappas, "Stable Flocking of Mobile Agents, Part II: Dynamic Topology," 42th IEEE Conference on Decision and Control, Maui Hawaii, December 2003, pp. 2016-2021