Demonstration of Digital Pheromone Swarming Control of Multiple Unmanned Air Vehicles

John A. Sauter,[*] Robert Matthews,[†] H. Van Dyke Parunak,[‡] and Sven A. Brueckner[§]
Altarum Institute, Ann Arbor, MI 48105

The use of digital pheromones for controlling and coordinating swarms of unmanned vehicles has been studied under various conditions demonstrating their effectiveness in multiple military scenarios in simulation. An experiment was conducted to verify that these same algorithms could effectively coordinate unmanned vehicles in a simulated exercise. Two air vehicles (modified target drones) controlled by digital pheromones and four ground robots controlled by a related stigmergic algorithm successfully executed a two-hour multi-mission surveillance, patrol, target acquisition, and tracking scenario without any scripting. The vehicles were given only high-level instructions, such as "survey this area and identify and track any targets" or "patrol around this convoy". The air vehicles were able to dynamically adapt to new commands and coordinate their actions with each other and the ground robots to achieve the objectives. The algorithm's robustness was demonstrated when it dynamically adjusted to the unplanned failure of one of the ground robots without any operator intervention.

I. Introduction

The word "swarming" is used to describe two different types of systems. Biologists use it to describe decentralized self-organizing behavior in populations of (usually simple) animals.[2,3,8] Examples include path formation, nest sorting, food source selection, thermoregulation, task allocation, flocking, nest construction, and hunting behaviors in many species. Military historians use it to describe a battlefield tactic that involves decentralized, pulsed attacks.[1,4,6]

Insect self-organization is robust, adaptive, and persistent and military commanders understand the advantages those attributes can have in a military engagement. While examples of swarming behavior used by human commanders have been studied, little attention has been given to the application of these approaches to the control and coordination of unmanned vehicles. This paper describes an adaptation of insect behavior using digital pheromones to control and coordinate the behaviors of many heterogeneous unmanned air (UAV) and ground (UGV) vehicles.

First we briefly describe the basic mechanisms and how they are used to control UAVs. We summarize some key results from an extensive study of these techniques under various simulated military scenarios. Finally we describe the successful demonstration of these algorithms involving two UAVs (modified target drones) and four UGVs in an extended, multi-phase military scenario held at Aberdeen Proving Grounds in September and October of 2004.

II. Digital Pheromones

There are several methods available for controlling and coordinating swarms of unmanned vehicles. Parunak [7] reviews the major classes of algorithms that have been applied to the Command and Control (C2) of multiple robotic entities. Digital pheromones are a *stigmergic* mechanism for coordinating and controlling swarming vehicles. "Stigmergy" is a term coined in the 1950's by the French biologist Grassé[5] to describe a broad class of multi-agent coordination mechanisms that rely on information exchange through a shared environment. Examples from natural systems show that stigmergic systems can generate robust, complex, intelligent behavior at the system level even when the individual agents are simple and individually non-intelligent. In these systems, intelligence resides not in a single

---

[*] Group Leader, Emerging Markets Group.
[†] Chief Scientist, Emerging Markets Group.
[‡] Senior Systems Engineer, Emerging Markets Group.
[§] Senior Systems Engineer, Emerging Markets Group.

distinguished agent (as in centralized control) nor in each individual agent (the intelligent agent model), but in the interactions among the agents and the shared dynamical environment.

Stigmergic mechanisms have some attractive features.

Simplicity.—The logic for individual agents is much simpler than for an individually intelligent agent. They can easily run on the small platforms envisioned for swarming vehicles. These agents are easier to program and prove correct. They can be trained with genetic algorithms or other weak optimization methods without requiring any knowledge engineering.[9]

Scalabilty.—Stigmergic mechanisms scale well to large numbers of entities. In fact, stigmergy *requires* multiple entities, and performance typically improves as the number of entities increases.

Robustness.—Because stigmergic deployments favor large numbers of entities that are continuously organizing themselves, the system's performance is robust against the loss of a few individuals. The simplicity and low expense of each individual means that such losses can be tolerated economically.

A. Vehicle Control with Digital Pheromones

Digital pheromones are modeled on the pheromone fields that many social insects use to coordinate their behavior. Digital pheromones support three primary operations, inspired by the dynamics of chemical pheromones.

1. They can be deposited in an area. Deposits of a certain flavor are added to the current amount of that flavor of pheromone located at that place. (Information fusion and aggregation).
2. They are evaporated over time. This serves to forget old information that is not refreshed. (Truth maintenance).
3. They propagate from a place to its neighboring places. The act of propagation causes pheromone gradients to be formed. (Information diffusion and dissemination).

In vehicle control, the area of operation is tiled with a network of "place agents" which maintains the digital pheromone field. Place agents are responsible for maintaining the level of each flavor of pheromone present at that location, propagating those pheromones to neighboring place agents, and evaporating them over time.

Several options are available for implementing place agents. Agents can be embedded in the environment using unattended ground sensors (UGS) networked through wireless communications. Place agents can also be distributed on Command and Control (C2) nodes according to their area of responsibility. The swarming platforms only need to communicate with the local UGS or C2 node. Alternatively each swarming platform maintains a full or partial version of the pheromone map representing the immediate vicinity around the unit. Pheromone map updates (deposits and withdrawals) need only be communicated locally to maintain each map. Since the information content is low (8 bytes/pheromone) and frequency of map updates is low (on the order of once a second), low bandwidth communications are sufficient to maintain the information flow among place agents.

Walker agents (representing the swarming unmanned vehicles) can sense the level of pheromone present and deposit additional pheromones at these place agents. A digital pheromone represents information about the system. Different "flavors" of pheromones convey different kinds of information. Based on the pheromones they sense they make decisions about what they will do next and where they will go. Some flavors of pheromone are attractive to walkers depending on their state, some are repulsive and some are neutral. The walker agents can use various means to combine the different flavors in the neighboring place agents to determine what action to perform next. For this application, a simple equation was used to combine the levels of the different pheromone flavors in the current place agent and the 8 neighboring place agents on the square grid. The results were used to construct a 9-segment weighted roulette wheel that the walker used to determine its next move (Figure 1).

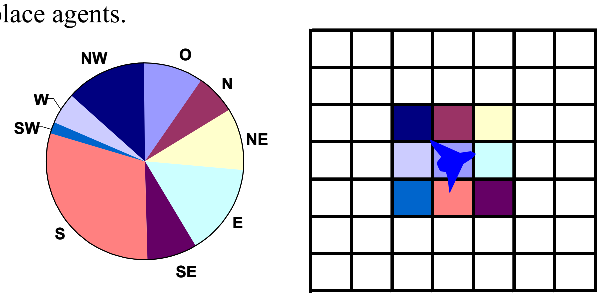

Figure 1. Weighted roulette wheel - Walker calculates difference between repulsive and attractive pheromone in each neighboring cell and its current location (O) to create a weighted roulette wheel used to determine the next location to move towards.

Avatars are used to represent the other entities in the system that are outside the scope of control of the digital pheromones. These could be friendly, enemy, or neutral entities, manned or unmanned, mobile or stationary. They can also sense and deposit pheromones primarily to inform the walker agents about their presence or to estimate their entity's movements in between sensory updates on their position.

A simple surveillance application will suffice to explain how digital pheromones can be used to control a UAV. Say the user has a set of irregular-shaped Areas of Interest (AOIs) that he wishes to have continuously monitored. Each area has differing priorities for surveillance that may change as new information arrives. The UAVs need to be able to configure themselves to survey the areas according to that priority (higher frequency of revisits for higher priority AOIs) regardless of the number of UAVs, the number, size or shape of the AOIs, and their varying priority. This particular application can be performed for any number of UAVs by two flavors of pheromone (one attractive and the other repulsive) and a combining equation that simply subtracts the repulsive pheromone from the attractive pheromone and zeros the result if negative.

The user starts by outlining each AOI on a map and assigning a surveillance priority to that AOI. If at a later time the user wishes to change the priority or add a new area, they can simply make the entry on the map interface. The system interprets that input and places a pheromone pump at each place agent enclosed by the AOI. The pump will regularly deposit a fixed amount of the attractive pheromone at that place agent as long as it remains on. A Walker agent is created for each UAV. The Walker agent senses the level of attractive and repulsive pheromones in the current and neighboring place agents, subtracts the two, and uses the result to construct a weighted roulette wheel. Spinning the wheel, the walker agent makes a stochastic decision on where to move next. It then places a deposit of repulsive pheromone on the place agent representing its next move, removes all attractive pheromone from that place agent, and turns off any pheromone pump found there. The pump will automatically restart after a delay inversely proportional to the priority the user placed on surveying that AOI (see Figure 2).

This simple algorithm is all that is required to completely control and coordinate the surveillance activities of any number of UAVs assigned to the mission as described above. They will spread out because of the repulsive pheromones they emit on the paths they take, and will congregate in the AOIs because of the attractive pheromone. Higher priority AOIs will emit more pheromone overall, attracting more UAVs than the lower priority AOIs.

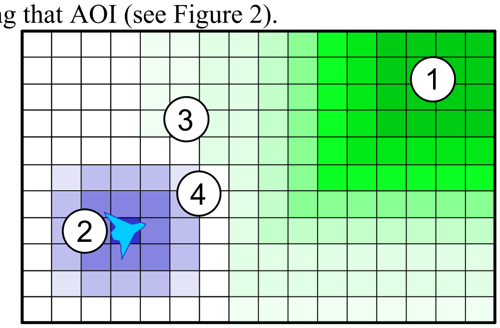

Figure 2. Attractive And Repulsive Pheromones For Surveillance, 1. Surveillance area deposits attractive pheromone, 2. Walker deposits its repulsive pheromone, 3. Pheromone infrastructure propagates both attractive and repulsive pheromone to form gradient, 4. UAV climbs net gradient, withdrawing attractive pheromone.

Coverage metrics for this system are good. In one experiment 30 UAVs moving at 90 kph were able to find and begin coverage of all 10 AOIs located in a 400 km^2 area after 35 minutes and reached a stable allocation of units across AOIs within 80 minutes on average. At any point in time a certain percentage of the population of UAVs will be outside any of the AOIs. Since the pheromones propagate outside the AOIs, some UAVs may decide to go to place agents immediately outside the AOI perimeter. These UAVs serve as a population that can be reassigned should there be a UAV lost to failure or refueling, a change in priority of an AOI, or the addition of a new AOI. The number of these wandering UAVs can be controlled by controlling the propagation factor on the attractive pheromone.

Other experiments showed that the pheromone algorithm is able to achieve impressive performance for its simplicity and ease of implementation. These experiments demonstrated that

1. Within 6 hours 30 UAVs were able to find all 10 mobile ground units that were hiding 33% of time in a 400 km^2 area divided into 40,000 cells (a UAV and a ground unit needed to be in the same cell for detection).
2. 30 UAVs were able to find and maintain tracks on 90% of the 20 ground units that hid 33% of the time in the same 40,000-cell grid with only a 0.5% track loss. The average revisit interval was controlled by a simple change in the amount of repulsive deposit used. When the speed of the ground units increased to half the speed of the UAVs, the UAVs were still able to track 80% of the units on average.
3. When the above experiment was repeated with the addition of a target identification task requiring confirmation from a UAVs equipped with a special ID sensor (only half of the UAVs had the sensor), the 30 UAVs were still able to track with the same efficiency.

These applications were accomplished with at most four pheromone flavors and a simple equation combining the various flavors. The equations governing the management of pheromones by the place agents, further details on the walker agents and their decision processes, as well as additional performance experiments on applications such as target acquisition, target tracking (intermittent and continuous), and sensor cueing are described in greater detail elsewhere.[10]

III. OASD Study Results

The OASD (NII) Decision Support Center sponsored a study entitled "Swarming Concept Development and Utility"**. The study investigated the performance of swarming assets in three joint capability areas: (1) Intelligence, Surveillance and Reconnaissance (ISR), (2) Communications and (3) Battle Damage Assessment (BDA). The Government furnished four operational situations (OPSITs) based on lessons learned from OIF and OEF (where swarming might have helped). These were elaborated by NAVAIR to compare swarming and non-swarming behavior. OPSIT 1 ("adversary command and control cell meeting in an urban area") has Blue reconnaissance detecting a meeting of high-level terrorist leaders, followed by a strike force to disrupt the meeting; this OPSIT was combined with OPSIT 4 ("friendly movement into a contested urban area") where Blue forces provide security to a military/civilian reconstruction team. Combining these two OPSITs makes it possible to simulate dynamic retasking and the interaction between multiple swarms since the two OPSITs are 30 miles apart and occur near-simultaneously. OPSIT 2 is based on a large-scale assault of a Red Weapons of Mass Effects (WME) site by Blue forces including air, land and sea assets. OPSIT 3 takes place in a hostile, high-desert mountainous area where the Blue forces must suppress terrorist/guerilla activities.

The Space Missile Defense Center Battle Lab simulated each of these scenarios using either the EADTB or SEAS simulation platforms. A base case was constructed using manned and unmanned platforms under traditional control approaches. This was compared to the same scenario using swarming entities that were controlled by Altarum's digital pheromone algorithms. The number of platforms and their capabilities were designed to provide roughly equivalent sensor coverage in each case. This was done to ensure that any improvement in performance from the swarming entities was not simply due to additional sensors in the field.

In OPSIT 1/4, a neighborhood suspected of being the location for a high-level terrorist C2 meeting is placed under surveillance. Thirty km away the area around a power plant reconstruction project is also under continuous surveillance. Six hours into the scenario Red forces attack the power plant. The Blue Commander requests additional swarming assets. The assets originally deployed to cover the C2 meeting site are closest, so some of them are tasked to provide target tracking and BDA for the power plant engagement until additional swarming assets arrive from the base to replace them. About the time the additional assets arrive, the C2 meeting begins. The SOF team is deployed to disrupt the meeting and the swarming assets are tasked to track any fleeing personnel or vehicles from the meeting site.

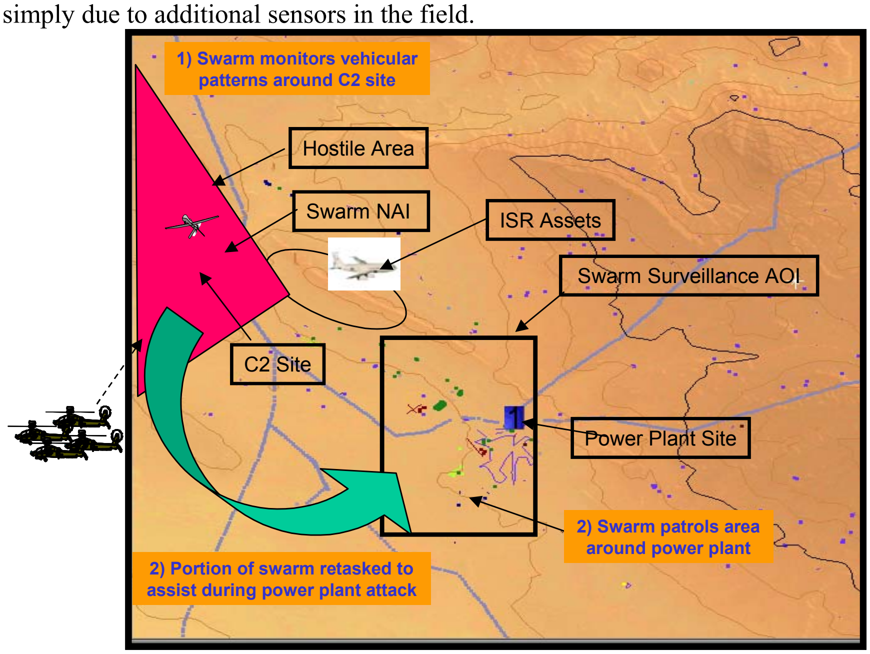

Figure 3. Opsit 1/4 – C2 meeting and power plant attack

In this scenario the swarming assets (UAVs) were able to detect the C2 meeting site 16 minutes earlier than the base case (using unattended ground sensors). The swarming UAVs also provided a 45-fold increase in tracking capability over the base case which used a single Global Hawk. The use of digital pheromones to control the UAVs resulted in an 18-fold increase in detections over the same UAVs flying fixed search patterns.

** The report is available at http://www.dsc.osd.mil/ under the link: "Swarming Concept Development and Utility, Final Report"

In OPSIT 2 a combined land and sea assault is planned on a suspected WME facility located on a wharf deep within enemy territory. The area around the wharf must first be cleared of mines using Unmanned Underwater Vehicles (UUVs). The SEALs assault the WME facility while the land forces and air surveillance assets are deployed to ensure that no material or personnel escape from the facility over land. Red forces in the hostile zone include BM-21s and chemical munitions. The campaign unfolds over multiple, coordinated phases between the land attack and the engagement from the SEA which must pass through the minefields and small attack boats.

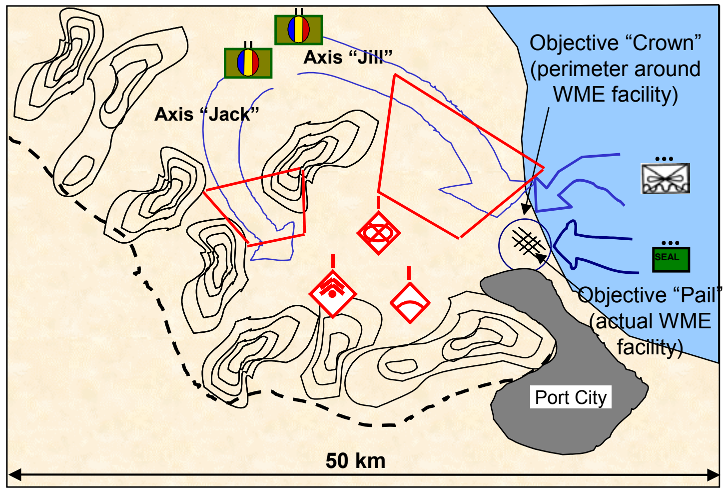

Figure 4. OPSIT2 – combined land and sea assault on WME facility

In this experiment, the swarming UAV assets were able to accurately map the extent of the chemical contamination caused by Red's chemical weapons giving the Blue commander a more direct route to engage the enemy. In the base case, the lack of that information delayed Blue's advance by three hours as they were forced to detour around the contaminated area thus engaging the enemy too late to capture the fleeing terrorists from the WME site. Overall the swarming entities compressed the mission timeline from twelve hours to six hours. Similarly the swarming UUVs were able to find a safe route through the minefield in 50% of the time it took UUVs searching with a traditional lawnmower pattern.

There were twice as many Blue casualties in the base case as the swarming case. The swarming entities detected about the same number of Red entities as the base case in the end, but they were able to detect them sooner. This allowed Blue to identify and neutralize more red forces (particularly the more lethal BM-21s) early in the operation before they were able to inflict casualties on Blue. When the swarming entities used fixed surveillance patterns, Blue losses increased 57% demonstrating the effectiveness of the digital pheromones over traditional control approaches.

Class III UAV survivability was increased three-fold since the swarming entities were able to quickly locate Red air defense artillery, which Blue destroyed. The increased number of Class III UAVs improved the situational awareness of Blue further reducing their causalities and increasing their OPTEMPO.

In OPSIT 3, the final scenario, 136 terrorists hiding in a large mountainous region are trying to escape on foot, mule, and trucks from Blue forces converging on their location by fleeing through the mountains in order to reach safe harbors in a city to the South (Figure 5). Blue deploys surveillance assets over a ten separate AOIs representing choke points through the 400 km^2 mountainous region. Their mission is to find the fleeing terrorists and maintain a track on them until attack forces can brought in to neutralize them.

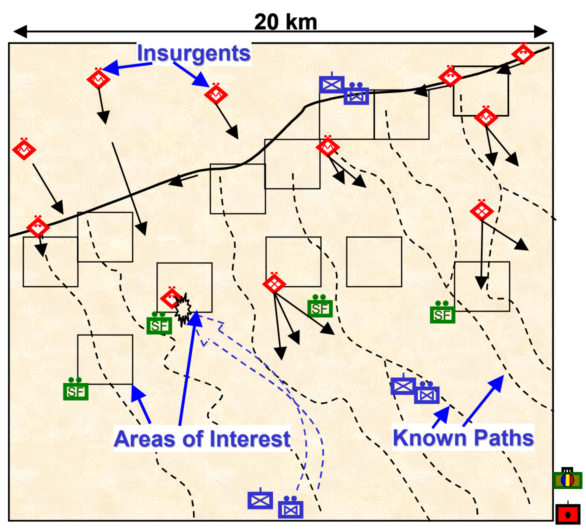

Figure 5. OPSIT 3 – detecting and tracking terrorist forces fleeing through mountainous region.

In the swarming case, 85% of the insurgents were captured or killed while only 59% were found in the base case. The swarming performance improved to 96% when the swarming assets provided target information to bring effects directly to the target.

The swarming system had several attributes that contributed to this success:

- The tendency of the swarming entities to occasionally wander outside the AOI led to additional detections. The AOIs in this scenario are the best estimates of the experts as to where surveillance assets should be concentrated. By not interpreting those areas as hard boundaries, the UAVs were able to find Red forces outside those areas.

- The swarming entities' AOI revisit rates were greater than legacy assets. Legacy assets tend to have wider field of view, but when there are several AOIs spread over a large area, a single Global Hawk must move between each one sequentially. By spreading out over the entire area, the swarming entities can more frequently cover the different AOIs then a single Global Hawk.
- The swarm was able to respond more quickly to the initial ground sensor detections than the legacy assets (for reasons similar to above). That timeliness meant that the target was more likely still in the area when the UAV came on the scene.
- Once a swarming entity detected a Red unit, it trailed the unit until the Blue forces were able to take it out. With the legacy assets, the surveillance could not be interrupted to trail a single target. Thus some detections did not result in a successful engagement when the insurgents evaded the Blue forces advancing on their last known position.

Finally the study looked at the effect of reducing both the number of UGSs and UAVs. Despite reducing the number of UAVs 60% (from 20 to 8) and the number of UGSs by 100% (from 24 to 0) the percentage of Red losses only changed by 9% (from 85% to 76%).

In summary the study came to the following conclusions:

- In all OPSITS, the swarming entities demonstrated an improvement in Force Effectiveness over the base case.
- The swarming entities improved inter-service synchronization, which helped improve mission success.
- Blue had better, (not just more) situational awareness with the swarming entities than without even after controlling for total sensor coverage.
- Mission success and OPTEMPO was increased by the improved BDA and situational awareness from the presence of swarming entities.
- The performance of the swarming entities was degraded when they were controlled by pre-planned behavior (such as fixed surveillance routes) rather than digital pheromones.
- The performance of the swarming entities degraded smoothly as the number of entities decreased.

IV. Demonstration

In October 2004 Altarum, Johns Hopkins University APL, and the Army Research Laboratory demonstrated the use of these swarming algorithms to control a heterogeneous population of air and ground unmanned vehicles in an urban combat scenario at Aberdeen Proving Grounds. The demonstration used four ground robots controlled by APL's co-fields algorithm (a stigmergic algorithm similar to digital pheromones), a mock urban area, and two Unmanned Air Vehicles (UAVs) controlled by Altarum's digital pheromone technology. The demonstration showed how these stigmergic swarming algorithms can control and coordinate the behaviors of a heterogeneous mix of vehicles.

The unmanned ground vehicles were research quality robots made by iRobot, Inc. All four robots used short range fixed acoustic sensors, laser range finders for obstacle detection and avoidance, and commercial GPS receivers for localization. One of the ground vehicles was equipped with a simulated target identification system (based on discriminating an acoustic signal of a specific frequency emitted by the target).

Figure 6. Mig 117 Bravo modified target drones were equipped with modern engine and Micro-Pilot autopilot

The air vehicles were modified Mig 117 Bravo target drones with a 6 ft wingspan (Figure 6). These target drones exist in large numbers in Army warehouses. The basic airframe was fitted with a modern engine, an autopilot by Micro-Pilot, and low light or infrared video camera. The autopilot was taught to take-off, hand launch, fly, and land completely autonomously.

The demonstration occurred on an airfield at Aberdeen Proving Grounds. An urban area with a power plant to be protected was simulated with some mock buildings (Figure 7). The area was initially surveyed by the unmanned ground vehicles (UGVs) to verify the area was clear. The UGVs then set up a perimeter patrol to protect power plant from intrusion.

Altarum's pheromone algorithms controlled and coordinated the flight of the two UAVs as they performed continuous surveillance over the urban area and adjacent territory looking for potential adversaries (see Figure 7). The two air units worked together to ensure even, thorough, and continuous coverage of all areas in the surveillance region while avoiding any collisions.

As the scenario unfolds, a convoy (simulated by a single van) enters the scenario and requests a UAV aerial patrol. The two UAVs patrol around the convoy as it moves up the airfield towards the mock power plant. Before the convoy reaches the power plant, the UAV's detect a potential threat in a crowd of people milling at an intersection ahead of the convoy. The convoy is alerted to the potential threat and halts awaiting verification before proceeding. Lacking the necessary sensors to make positive identification the UAVs deposit a pheromone on the potential threat that attracts the UGV with the target identification sensor. It makes a positive identification of the threat. The UGV and UAVs then track the threat until it can be neutralized by nearby Blue forces. Once the threat has been removed, the UAVs and UGVs perform a search of the area to ensure no further threats exist before the "all-clear" signal is given to the convoy so it can continue on its mission.

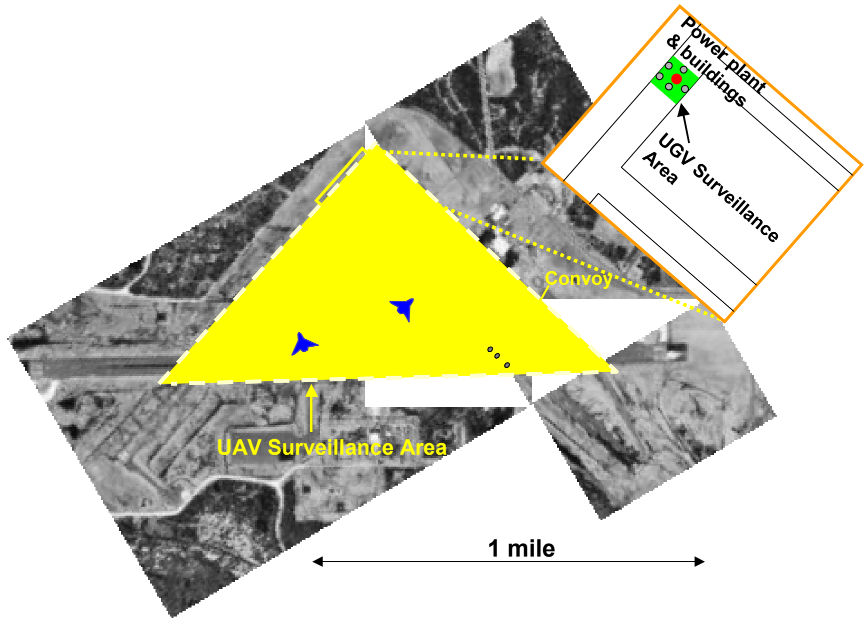

Figure 7. UAVs and UGVs cooperated in a successful demonstration of a multi-mission surveillance, patrol, and tracking mission at Aberdeen Proving Grounds in September and October of 2004.

The whole scenario took about two hours to complete. The only operator intervention involved high-level commands such as identifying the area to be surveyed, or the area or object to be patrolled. During the demonstration, one of the ground robots suffered an unplanned malfunction. The other ground robots were able to dynamically readjust their patrol patterns to accommodate the missing unit without any intervention by the operator. This unplanned event helped to demonstrate the robustness of these algorithms to unexpected events and failures.

The demonstration showed cooperative behavior between the air and ground units and two related, but different stigmergic algorithms. The UAV detected a potential threat and then had to coordinate with the ground vehicle with the necessary target identification sensor to verify the threat. The two algorithms controlling the UAVs and UGVs accomplished this cooperation through the deposit of a single pheromone.

The actions of the vehicles were not scripted as evidenced by their adapting to the unplanned failure of one of the ground robots. Rather than specify each vehicle's task, the operator simply gave a high level command to the whole swarm. The vehicles autonomously configured themselves to determine which vehicle would perform what task in order to accomplish the overall objective. The operator was free to monitor their behavior, receive their reports, and provide additional guidance as needed when priorities or mission objectives changed. The swarm did not need any special configuration to meet a wide variety of mission requirements, irrespective of the operating environment or the number and type of vehicles involved.

V. Conclusion

At the start of this study there was concern about the whether the wide range of scenarios and the requirements they placed on the swarm would require a sophisticated and complex algorithm in order to meet the wide range of mission objectives. This study was able to demonstrate that a simple pheromone mechanism can be used to perform all the functions required by these scenarios. The surprising versatility arising from such simple mechanisms is one of the more promising aspects of this new class of algorithm.

The mechanism proved to be surprisingly robust to large variations in the parameter settings. Certain parameters had a greater influence than others, but the mechanism performed well even when those were varied by a factor of 10 or 100. Adding a new function typically involved at most

- Adding a new pheromone
- Adding a new term to the equation combining the pheromones
- Conducting some experiments to get the right settings

The study demonstrated that swarming provides a number of advantages over legacy systems. Swarms controlled by Altarum's digital pheromones improved situational awareness, OPTEMPO, and force effectiveness. The stigmergic pheromone algorithms have demonstrated great promising as a means to control swarms of unmanned vehicles. They are adaptable to a wide variety of scenarios, robust against change and failures, easy to program and tune, and effective in controlling both large and small swarms distributed over large areas. The ability to achieve complex coordination and control of large swarms of heterogeneous vehicles without relying on heavy computation or centralized control makes this class of algorithms ideal for the smaller autonomous platforms of the future.

Acknowledgments

This study was supported by the Decision Support Center of ASD NII. The views and conclusions in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Department of Defense, or the US Government. This study was supported and monitored by the Air Force Research Laboratory (AFRL) under contract number F30602-02-C-0196.

References

[1]Arquilla, J. and Ronfeldt, D., Swarming and the Future of Conflict, RAND, Report No. DB-311, Santa Monica, CA, 2000.

[2]Bonabeau, E., Dorigo, M., and Theraulaz, G., *Swarm Intelligence: From Natural to Artificial Systems*. Oxford University Press, New York, 1999.

[3]Camazine, S. et al., *Self-Organization in Biological Systems*. Princeton University Press, Princeton, NJ, 2001.

[4]Edwards, S. J. A., Swarming on the Battlefield: past, Present, and Future, RAND, Report No. MR-1100-OSD, Santa Monica, CA, 2000.

[5]Grassé, P.-P., "La Reconstruction du nid et les Coordinations Inter-Individuelles chez *Bellicositermes Natalensis et Cubitermes sp.* La théorie de la Stigmergie: Essai d'interprétation du Comportement des Termites Constructeurs," *Insectes Sociaux*, Vol 6, 1959, pp. 41-84.

[6]Inbody, D., "Swarming: Historical Observations and Conclusions," *Swarming: Network Enabled C4ISR*, Tysons Corner, VA, 2003.

[7]Parunak, H. V., "Making Swarming Happen," *Swarming: Network Enabled C4ISR*, Tysons Corner, VA, 2003.

[8]Parunak, H. V. D., "'Go to the Ant': Engineering Principles from Natural Agent Systems," *Annals of Operations Research*, Vol 75, 1997, pp. 69-101.

[9]Sauter, J. A., Matthews, R., Parunak, H. V. D., and Brueckner, S., "Evolving Adaptive Pheromone Path Planning Mechanisms," *Autonomous Agents and Multi-Agent Systems (AAMAS02)*, Bologna, Italy, 2002, pp. 434-440.

[10]Sauter, J. A., Matthews, R., Parunak, H. V. D., and Brueckner, S. A., "Performance of Digital Pheromones for Swarming Vehicle Control," *Fourth International Joint Conference on Autonomous Agents and Multi-Agent Systems*, Utrecht, Netherlands, 2005, pp. 903-910.