Forthcoming in Proceedings of SPIE Defense & Security Conference, March 2008, Orlando, FL

Distributed Pheromone-Based Swarming Control of Unmanned Air and Ground Vehicles for RSTA

John A. Sauter, Robert S. Mathews, Andrew Yinger[*]

NewVectors
3520 Green Ct
Ann Arbor, MI 48105

Joshua S. Robinson, John Moody[†]

Augusta Systems, Inc.
3592 Collins Ferry Road
Suite 200
Morgantown, WV 26505

Stephanie Riddle

Naval Air Systems Command (NAVAIR)
48150 Shaw Road
Bldg. 2109 Suite S211-21
Patuxent River, MD 20670

Abstract

The use of unmanned vehicles in Reconnaissance, Surveillance, and Target Acquisition (RSTA) applications has received considerable attention recently. Cooperating land and air vehicles can support multiple sensor modalities providing pervasive and ubiquitous broad area sensor coverage. However coordination of multiple air and land vehicles serving different mission objectives in a dynamic and complex environment is a challenging problem. Swarm intelligence algorithms, inspired by the mechanisms used in natural systems to coordinate the activities of many entities provide a promising alternative to traditional command and control approaches. This paper describes recent advances in a fully distributed digital pheromone algorithm that has demonstrated its effectiveness in managing the complexity of swarming unmanned systems. The results of a recent demonstration at NASA's Wallops Island of multiple Aerosonde Unmanned Air Vehicles (UAVs) and Pioneer Unmanned Ground Vehicles (UGVs) cooperating in a coordinated RSTA application are discussed. The vehicles were autonomously controlled by the onboard digital pheromone responding to the needs of the automatic target recognition algorithms. UAVs and UGVs controlled by the same pheromone algorithm self-organized to perform total area surveillance, automatic target detection, sensor cueing, and automatic target recognition with no central processing or control and minimal operator input. Complete autonomy adds several safety and fault tolerance requirements which were integrated into the basic pheromone framework. The adaptive algorithms demonstrated the ability to handle some unplanned hardware failures during the demonstration without any human intervention. The paper describes lessons learned and the next steps for this promising technology.

Keywords: swarm intelligence, digital pheromone, autonomous vehicle, RSTA, stigmergy

1. INTRODUCTION

The emergence of decentralized, asymmetric threats to global security and stability is causing the United States military to rely, more than ever, on innovative technologies for Reconnaissance, Surveillance, and Target Acquisition (RSTA). Asymmetric threats easily get lost in a high clutter environment which makes them difficult to detect. To counter this threat the military is increasingly relying on remote sensing and monitoring systems deployed on Unmanned Vehicles (UxVs) for detecting and identifying potential threats while minimizing the harm to military personnel. "Reconnaissance and Surveillance" and "Target Identification and Designation" were named the number one and two priorities for unmanned systems in the recent OSD Unmanned Systems Roadmap[1]. Through deployment, UxVs have not only decreased risks confronted by military personnel, but have also become an indispensable weapon in the War on Terrorism, as well as an important tool for homeland security. In the near future, network-centric military operational environments will rely heavily on swarms of autonomous sensors deployed on UxVs and scattered throughout the battlefield to provide a superior intelligence gathering capability.

This sensor-enabled network-centric future will require advances in several technology areas. Current UxVs typically require multiple operators and a ground station for each vehicle. Operator to vehicle ratios greater than one are a scaling barrier that must be overcome to meet the needs of future networked sensor systems[2]. If we are to dramatically increase

[*] {john.sauter, robert.matthews, andrew.yinger}@newvectors.net; phone 734-302-5660
[†] {jrobinson, jmoody}@augustasystems.com; phone 304-599-3200

Copyright © 2008 Society of Photo-Optical Instrumentation Engineers. This paper will be published in Proceedings of SPIE Defense & Security Conference and is made available as an electronic preprint with permission of SPIE. One print or electronic copy may be made for personal use only. Systematic or multiple reproduction, distribution to multiple locations via electronic or other means, duplication of any material in this paper for a fee or for commercial purposes, or modification of the content of the paper are prohibited.

the number of sensors in the field we need systems that can operate semi-autonomously. This requires the development of robust coordination and control technologies that can manage large numbers of networked sensor entities cooperating on a common mission objective. Automated target detection and recognition (ATR) algorithms are needed so multiple sensors can function autonomously and reliably while being monitored by a single operator thus removing the operator-to-vehicle ratio barrier. We need adaptable network technologies that can function reliably over many mobile nodes moving in and out of the same space. Finally we need Operator System Interfaces (OSIs) that can intelligibly present information being collected from the sensor swarm to allow monitoring and management of the system with minimal human intervention.

The research described in this paper was sponsored by NAVAIR under two projects. A study project researched and prototyped different technologies to address the gaps described above in coordination, ATR, communications, and OSI. A companion demonstration project engineered selected technologies from the study into an integrated RSTA system. The goal of this project was to demonstrate the feasibility of using swarm intelligence to coordinate multiple heterogeneous vehicles in a realistic application. Table 1 lists specific technology objectives identified in the OSD Roadmap that were addressed in some fashion by these projects. In this paper we touch on each of the technologies demonstrated but focus on the collaborative coordination and control strategies.

Table 1. Aspects of OSD Unmanned Systems Roadmap Objectives Addressed by Project

Roadmap Area	Objective	Technology Applied	Extent Objective Addressed
Autonomy	Collaborative coordination	Pheromone-based swarming coordination plus ATR	High: able to coordinate multiple vehicles of different types driven by RSTA application demands
Cognitive Processes	Support human perception, reasoning and decision making	OSI: geospatial display of UxV status and combined sensor results	Medium: Further advances in combining sensory data for display necessary
	Capture cognitive capabilities in ATR	Haar and Bessel K ATR algorithms	Low: focus on mathematical rather than cognitive approaches to ATR.
Common Control	Control of multiple vehicle types from single station	Operator controls all vehicle types through common swarming algorithm	Medium: Same algorithm controls both UGS and UAS (extendable to maritime) enabling common control of multiple vehicle types.
Communications	Available bandwidth, range	Mesh network	Medium: demonstrated successes with commercial not with military comms.
	Multiple vehicles in same space	Mesh network	Medium: demonstrated multiple vehicles operating on an ad-hoc network.
Data Interfaces	Move from off-GIG to on-GIG	Network based communications	Medium: Each UxV a full node on the network, but not yet GIG compliant
	Serve as relay nodes	Mesh network	High: this is an inherent capability of mesh networks
Human Systems Integration	Provide high level objectives, ROE, constraints, weapons authorization, mission coordination, manage launch and recovery	Operator enters high-level guidance sufficient to direct pheromone-based swarming coordination	Medium: Command function specifies high level objectives, ROE, and constraints. Weapons authorization, multi-mission coordination and launch and recovery beyond current scope.
	Monitor, interpret, diagnose, manage resources, maintain components	OSI: graphical display elements and techniques to enhance operator understanding	Medium: operator can monitor locations, ATR results and limited status and diagnostics. New views into performance and additional diagnostic capabilities slated for next phase

In the following sections we summarize the approaches described in the literature to swarming control. We define the requirements that a realistic RSTA mission places on a swarm of networked UxV sensor platforms. We describe the basic mathematical framework for the swarming algorithms used for coordination and control and introduce the technologies integrated for ATR, networking, and the Operator System Interface. Finally we show how each of the technologies came together in the final demonstration and talk about lessons learned.

2. APPROACHES TO SWARMING UXV CONTROL

There are several approaches described in the literature for distributed control of swarms of unmanned vehicles. Most of the work in distributed vehicle control involves various kinds of field-based mechanisms where a scalar field is generated by a combination of attracting and repelling elements and the agents respond to those forces to follow gradients in this field. Within this class of algorithms are particle systems based on Reynold's model [3], potential fields based on physics models [4, 5], and digital pheromones based on insect models [6-11]. Digital pheromones are similar to potential fields, but they more naturally lend themselves to decentralized computation than potential fields. Field-based methods rely on *stigmergic* mechanisms for coordinating and controlling swarming vehicles. "Stigmergy" is a term coined in the 1950's by the French biologist Grassé[12] 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 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.

Digital pheromones as described in [6, 7] are modeled on the pheromone fields that many social insects use to coordinate their behavior. Different "flavors" of pheromones convey different kinds of information. They have been used to support a variety of traditional swarming functions including path planning [9, 13] and coordination for unmanned vehicles[14, 15], positioning multi-sensor configurations[16], and maintaining line of sight communications in mobile ad hoc networks [17].

In vehicle control the area of operations is tiled with a grid of *place agents*. Each place agent maintains the level of each flavor of pheromone present at that location and performs the basic pheromone operations. These pheromone maps are like layers in a Graphical Information System (GIS). Each layer represents the spatial concentration of a single pheromone flavor over the area of operations. Unlike a GIS layer, these maps are dynamic. The place agents perform three primary pheromone operations, inspired by the dynamics of chemical pheromones.

1. *Deposit (information aggregation)* – A deposit of a pheromone flavor, Φ_f is accumulated with the current amount of pheromone Φ_f located at that place. This represents the accumulation of information (such as the evidence for a target) at that location. Deposits are normally summed with the current amount of the pheromone flavor present, but other accumulation operations are possible depending on the rules for aggregating different kinds of information.
2. *Evaporate (truth maintenance)* – Pheromones are evaporated over time. This serves to forget old information that is not refreshed serving a kind of truth maintenance function over the information space.
3. *Propagate (information diffusion)* – Pheromones propagate from a place to its neighboring places. This serves to propagate local information to a broader neighborhood affecting decisions of more remote nodes. Propagation also forms gradients that can guide swarming agents in navigating through a complex information landscape.

There are two equations governing the maintenance of the pheromone field employed by the place agents. The evolution of the strength of a single pheromone flavor at a given place agent is commonly defined by:

$$s(\Phi_f, p, t) = E_f\left((1-G_f)\left(s(\Phi_f, p, t-1) + d(\Phi_f, p, t)\right) + g(\Phi_f, p, t)\right) \tag{1}$$

Where $s(\Phi_f, p, t)$ is the strength of pheromone flavor f at place agent p and time t, $d(\Phi_f, p, t)$ is the accumulation of external deposits of pheromone flavor f within the interval $(t-1, t]$ at place agent p, $g(\Phi_f, p, t)$ is the propagated input of pheromone flavor f at time t to place agent p, $E_f \in (0, 1]$ is the evaporation factor for flavor f, and $G_f \in [0, 1)$ is the propagation factor for flavor f. Each place agent applies equation (1) to each pheromone flavor once during every update cycle.

The second equation describes the propagation received from the neighboring place agents:

$$g(\Phi_f, p, t) = \sum_{p' \in N(p)} \frac{G_f}{|N(p')|} \left(s(\Phi_f, p', t-1) + d(\Phi_f, p', t)\right) \tag{2}$$

Where $N: p \rightarrow p$ defines the neighbor relation between place agents. Each neighbor place agent p' propagates a portion of its pheromone to p each update cycle. The fraction propagated depends on the parameter G_f and the total number of its neighbors.

Deposits increase the amount of pheromone at a place agent, while evaporation and propagation reduces the amount present. Brueckner[18] has proven certain stability theorems regarding pheromone fields. In particular the equations ensure that with regular deposits of a pheromone the amount of pheromone in a cell eventually asymptotes to a constant level.

Place agents maintain information about the geospatial area they represent. They can deposit pheromones to model information or requirements (such as a surveillance need) for that area. *Avatars* are agents that control the unmanned platforms. They use the pheromone fields to plan the future actions of the vehicle they are managing. They are able to sense the level of pheromones at their current location and in the immediate vicinity and they can deposit and remove pheromones of different flavors.

Previous work has described how digital pheromones were used to control multiple UAVs in a simulated surveillance and target tracking application [7, 8]. A simplified target recognition algorithm was also described (requiring one sensor to cue another for image recognition). The performance of these algorithms was studied in both abstract simulations and in several realistic battlefield scenarios demonstrating superior mission effectiveness over traditional approaches. The algorithms have proven adept at handling a diverse set of requirements in dynamic environments using very simple mechanisms. This prompted the follow-on work reported in this paper to adapt these algorithms to the specific requirements for a swarming unmanned system in a realistic RSTA application using commercial platforms.

3. REQUIREMENTS FOR SWARMING RSTA

This section considers how the Automatic Target Recognition (ATR) algorithms and the UxV platform and sensor constraints impact the swarming algorithms. We first describe the ATR algorithms employed and the requirements they place on the swarming algorithm. We then consider the issues that arise when deploying a swarming algorithm on air and ground vehicles.

3.1 ATR Algorithm

An optical sensor on a UAV captures images of varying resolution and fidelity. In the study West Virginia University (WVU) used a realistic computer simulation to create synthetic images. These images were distorted to simulate the effects of noise, focus blur, motion blur, low contrast, low light, and occlusion. The detection algorithm utilized Haar-based features trained on standard images of the targets at different angles. Regions within images where potential targets have been detected are further subjected to recognition algorithms based on Bessel K forms [19]. The performance of the detection and recognition algorithms was evaluated on the images to determine the effect of optical and environmental distortions. Since swarming UAVs can produce multiple images of the same target, the improvements possible with super-resolution techniques were also studied.

3.2 Swarming Algorithm Requirements from ATR

The ATR algorithm involves processing one or more images, from one or more sensors, from one or more angles, and with possibly varying resolutions. For the swarming algorithm this means directing the right sensor, to the right location, with the right attitude to the target so the sensor can collect the data necessary for the ATR algorithm to successfully complete its processing. Thus the swarming algorithm is responsible for allocating the sensor tasks among all the platforms and planning and deconflicting the paths for all the units to ensure that these requirements are met.

For any given image, the ATR algorithm generates a best estimate for what target(s) may or may not be located within the frame with a level of certainty. If the uncertainty remains high after processing the image additional images of the area may be required to confirm the presence or absence of a target. Alternatively other sensor modalities may need to be brought to bear on the area (high resolution images, infrared, Synthetic Aperture Radar, etc.). The swarming algorithm is responsible for ensuring that a sufficient number of samples are taken with the right sensors for each area to enable the ATR algorithm to perform its function.

3.3 RSTA Platform and Sensor Constraints

For the purposes of the study project a RSTA application scenario was defined. The initial guidance from NAVAIR and WVU provided the following parameters for this study:

1. The UAV shall fly at an altitude between 150m and 300m (nominal 230m) at 25m/s.
2. The camera shall have zoom capability typical for commercial cameras.
3. The camera shall be gimbaled to allow it to maintain a constant orientation with respect to the ground while the aircraft makes maneuvers, or to take side glancing photos at different declination angles for 3-D target recognition.

Lightweight, inexpensive gimbal systems suitable for small UAS are becoming increasingly available. The use of a gimbal improves ATR and target location performance, but is not a hard requirement for the algorithms. A gimbal was not used for the demonstration project.

4. The swarming UAVs and sensors shall capture 64 x 64 pixels on the target. Lower resolutions are allowable for target detection with higher resolutions enabling better target recognition.
5. The swarming UAVs shall respond to specific requests from the ATR to capture images of defined resolutions from defined orientation and declination angles.
6. All targets are on the surface, immobile, and between 1m x 1m and 6m x 6m in size. The size requirement and #4 above require a pixel density on the ground between 114 and 4096 pixels/m^2.
7. A 35mm, 2 megapixel (1600 x 1200) sensor requires a lens with a 420mm focal length to image 4096 pixels/m^2 from 300m altitude. This is a fairly large lens for a small UAV payload. Alternatively sensors with higher pixel counts or lower altitudes could be used to obtain the highest pixel density required. This suggests that still images rather than video should be used. For example a 10.1 megapixel APS-C sized sensor would only require a lens with a 110mm focal length to image 4096 pixels/m^2 from 300m altitude.
8. A 1600 x 1200 sensor flying at 25 m/s ground speed capturing 114 pixels/m^2 on the ground must sample the sensor every 0.5 sec to capture at least one full image of a 6m target in the image frame. This sampling rate is well within the range of most commercial cameras once initial auto-focusing and auto-exposure compensation has occurred.

3.4 Other Requirements

While simulation studies and hardware prototypes provide valuable experience in developing new technologies, there are always numerous issues that arise when you move to an integrated hardware demonstration of the entire system in a realistic scenario. Previous experience with adapting digital pheromones to controlling real UxVs [7, 8] identified and resolved issues related to communications, constrained UAV turning radii, and no-go areas restricting UxV movement. For this work we also identified and addressed the following issues:

- Safety issues are critical in unmanned systems. Collision with other UxVs, and collision with non-swarm entities in the air or on the ground must be avoided. Sometimes additional safety factors must be incorporated in the design of the algorithms when hardware methods alone are insufficient.
- Hardware and software failures need to be accommodated and backup and recovery procedures put in place. The algorithms must be designed to be robust in the face of different kinds of failures.
- Errors in communication and positioning, can lead to dropped messages, missed updates, and inaccuracies in computing vehicle and target locations. This complicates the navigation, collision avoidance, and target acquisition functions requiring strategies to accommodate these errors.
- Turns and climbs consume more energy decreasing the effective range and time on station for the UxV. The swarming algorithm needs to consider the cost for such maneuvers in making its decisions.
- UAV ground speed is affected by wind speed and heading which in turn impacts ground coverage and image capture rates. Increasing the ground speed while flying downwind might induce motion blur or increase the image capture rate above the processing capability of the ATR algorithm. The swarming algorithm must be able to accommodate varying wind speeds and headings in order to provide the right sensor collection conditions for effective ATR.

4. DESCRIPTION OF THE MODIFIED SWARMING ALGORITHM

This section describes the modified swarming algorithm to accommodate the ATR algorithm and unmanned platform requirements. The flavors of pheromones used in the control of the UxVs are defined. The evaluation function for the path planning algorithm is introduced. Finally a description of how the distributed pheromone map is maintained is included.

4.1 Pheromone Flavors

Each vehicle maintains its own version of the pheromone map. There are four primary pheromones involved in the control of the UxVs:

1. Φ_Θ Search pheromone that attracts UxVs to areas that have not been searched yet. This is a measure of the uncertainty about an area. High uncertainty (Φ_Θ) attracts sensors that can reduce the level of uncertainty about the presence or absence of targets in that area.

2. Φ_r Target Recognized pheromone deposited by the ATR algorithm that has detected a possible target. In our demonstration an ATR algorithm onboard the UAVs identified targets from images taken by the UAV's camera. Additional ATR hits from the same or other UAVs increased the evidence for a target present at that location by depositing additional Target Recognized pheromone.
3. Φ_x No-go pheromone deposited in areas that represent no-fly zones for UAVs or no-go zones for UGVs.
4. Φ_v Vehicle path pheromone deposited along the planned path for each vehicle.

4.2 Path Planning

The UxVs plan their paths independently and broadcast their planned path to other vehicles in their vicinity. Each path consists of a set of waypoints. The length of the path is long enough so that potential collisions can be detected and corrective measures taken. The heart of every UxV planning algorithm is the evaluation function. The purpose of the function is to evaluate different paths based on the following high-level objectives:

1. Move quickly to areas where there is the most need for my sensor (highest uncertainty, or possible target detected that needs additional confirmation my sensor can provide).
2. Prefer to move in straight lines to conserve fuel (UAVs) or time (UGVs).
3. Prefer to move at optimal airspeeds (UAVs) or more slowly (UGVs) to conserve energy.
4. Prefer to fly at constant altitude or move on level ground to conserve energy.
5. Prefer to fly at optimal ground speed for ATR (UAVs). In high winds this and objective #3 leads to a preference to fly cross wind to maintain constant ground speed near the optimal airspeed.
6. Stay away from other vehicles and planned paths to avoid collisions and duplication of effort.
7. Stay away from no-go zones.

These high level objectives are implemented in the form of a benefit to cost ratio to translate them into a more precise mathematical formulation that can drive swarming decisions. A simple way to calculate the benefit for a path (objective 1) is the sum of the expected change in Search pheromone and Target Recognized pheromone in all the cells $\aleph$ within the field of view of the sensor along the path:

$$B_p = \sum_{\aleph} \left(\theta \Delta \Phi_{\Theta} + \rho \Delta \Phi_r \right) \tag{3}$$

where θ and ρ are tuning constants. The expected change in Search pheromone and Target Recognized pheromone depends on the sensor's capabilities and its ability to reduce uncertainty or to improve the confidence in the identification of a target. For example, if the uncertainty remaining in an area is distinguishing whether the target recognized is friendly or hostile and the sensor onboard cannot make that distinction, then the expected change in Search pheromone would be zero and it would not be attracted to that area. For our demonstration the tuning constants θ and ρ were set to 1.

While attempting to maximize the benefits, the swarming algorithm must also attempt to minimize the costs (objectives 2-7). The cost for a path has three elements: energy used, potential for collision and duplication, and proximity to no-go zones. The energy cost C_f includes the energy cost for speed, heading change, change in altitude or elevation, and movement parallel or perpendicular to the wind. This is dependent on the platform. The other two costs are designed to support (but not enforce) the rules that there be no collisions and no violation of the no-go zones. The path pheromone Φ_v describes where each sensor in the area is planning on surveying in the future so that other vehicles can avoid searching those areas. It is used to calculate the cost for a potential collision and duplication of search effort and helps support the no collision rule. The pheromone Φ_x is deposited in the no-go zones and propagates a short way into the no-go zones. It is used to calculate the cost of proximity to a no-go zone and helps support the rule forbidding violation of that space. This pheromone provides a "soft" boundary so that UxVs can sense when they are nearing a no-go boundary and can begin planning maneuvers to avoid it. These two pheromones are summed along with the energy cost for each segment of the move $\Im$ to arrive at a total cost for the path:

$$C_p = \sum_{\Im} \left(\phi C_f + \upsilon \Phi_v + \chi \Phi_x \right) \tag{4}$$

where ϕ, υ, and χ are tuning constants. The evaluation function for a path is then:

$$\mathrm{T}_p = \frac{\sum_{\aleph}\left(\theta\Delta\Phi_\Theta + \rho\Delta\Phi_r\right) + k}{\sum_{\Im}\left(\phi C_f + \upsilon\Phi_v + \chi\Phi_x\right) + k} \tag{5}$$

where k is a constant to avoid irregularities when the benefit or cost evaluate to zero.

When a new path is required, a search is performed to find the best path. Different exploration strategies can be employed. One approach repeatedly applies equation (5) to all cells reachable within a certain radius of the current waypoint in the path and either picks the best point as the next waypoint or selects one stochastically using a weighted roulette wheel. This is repeated to identify each waypoint in the path. Multiple paths are evaluated in their entirety using equation (5) to select the best path. Paths that end up leading to a collision with other UxV planned paths or that enter no-go zones are eliminated from consideration. The application of this rule is one example of how the emergent properties of these algorithms can be managed to enforce hard constraints imposed on the system.

The tuning constants in the path evaluation function can be varied to improve the mission effectiveness of the system based on key performance parameters. Tuning can be performed manually using a Design of Experiments approach and simulation studies. Preferably the constants are tuned using an optimization algorithm such as the genetic algorithm described in previous work[6]. In this study the constants were manually modified from the tuned results of previous work. Experience has shown that the swarm performs well over a wide range of values for the tuning constants. Once the constants of equation (5) are tuned they are suitable for a number of applications.

4.3 Maintenance of Distributed Pheromone Maps

There are several methods that can be used for maintaining the pheromone maps. A centralized approach is the easiest, but requires that each node in the swarm have access to the centralized map which may not always be possible. In the demonstration we chose to implement a distributed pheromone map. Each UxV maintains its own grid of place agents defining the pheromone map over the entire area of interest. UxVs communicate events to their peers that affect the state of the pheromone map. While the events are communicated locally, they are propagated through the swarm. Distant UxVs may not receive some updates or may receive them late. This means that UxVs that are far apart may have different pheromone views of each others' areas. However, since UxVs plan locally, the errors in the pheromone maps for events occurring far away have little impact on the performance of the swarm. For this demonstration a mesh network ensured that all nodes in the swarm received the update messages, so the impact of errors in the pheromone maps was not studied.

UxVs broadcast the following information to their peers:

- *Current location, heading, and speed* – used to locate all the other UxVs in 3D space. Receiving UxVs deposit Φ_v vehicle path pheromone at that location.
- *Planned path* – is the current path plan as described above. Receiving UxVs deposit Φ_v vehicle path pheromone along the entire path.
- *Uncertainty reduced* – each sensor collection event causes a reduction in the uncertainty for the cells covered by its field of view. This message updates the uncertainty for all those cells. Receiving UxVs remove Φ_Θ search pheromone in proportion to the reduction of uncertainty reported.
- *Target detected or recognized* – the result of the onboard ATR algorithm when it detects or recognizes a target. It includes target location and level of confidence. Receiving UxVs deposit Φ_r target recognized pheromone at the location specified in proportion to the confidence reported.

The pheromone fields of all UxVs are initialized with $\Phi_\Theta = 1$ (except in no-go zones), $\Phi_v = 0$, and $\Phi_r = 0$. The no-go pheromone, Φ_x is initially deposited in the no-go areas and propagated a short way into the adjacent cells. After that Φ_x remains static unless the user makes a change in the no-go areas. The evaporation and propagation factors for Φ_x are set to 1 and 0 respectively so it does not evaporate or propagate. Search pheromone is regularly deposited by the place agents. The deposit amount is adjusted so Φ_Θ asymptotes within the range [0,1] depending on the current uncertainty for that place agent's location. This level is only changed when a sensor covers that location. The evaporation factor for target recognized pheromone Φ_r is set near 1 so it does not evaporate before confirming sensors can be attracted to the area.

5. FLIGHT AND GROUND TESTS

The next step was to demonstrate the feasibility of swarming multiple vehicle types in a RSTA application. This section describes the vehicles, the Operator System Interface, and the changes to the swarming algorithms required by the demands of the hardware and safety issues encountered. The test scenario is described along with the results from the demonstration and the lessons learned.

5.1 Description of the Vehicles

The AAI Aerosonde Mk 4.1 UAV was chosen as the air platform (Figure 1). This UAV cruises at 25 m/s, carries a maximum payload of 5 kg, can operate over 30 hours and has a minimum turning radius of roughly 140m. The nominal operating altitude for the aircraft in this test was 230m. The UAVs were equipped with a Canon PowerShot S80 or a Lumenera Lw235 color camera, but they were not gimbaled.

Figure 1. Aerosonde Mk 4.1 UAV

Modified Pioneer 3-AT robots were used for the ground vehicles (Figure 2). They can move at 3 kph, climb 45^o grades, carry 30 kg of payload, operate 3-6 hours, and turn in a 40 cm radius. The UGV is equipped with 8 fore acoustic proximity sensors, GPS, digital compass, video camera, and a simulated target confirmation sensor (an RF receiver).

Figure 2. Pioneer 3-AT Ground Robot

Both the UAV and the UGV were equipped with an Augusta Systems SensorPort payload computer utilizing a 1.4 GHz, low voltage, Pentium-M processor module running Windows XP Embedded on a 1 GB Compact Flash. A MeshNetworks WMC6300 2.4 GHz subscriber card providing a 1.5 Mbps (6 Mbps burst rate) ad-hoc mesh network supported communications of command and control and imagery data with the ground stations. A single laptop on the MeshNetwork is used as a “payload control station” for monitoring the vehicles and providing manual control in emergencies. A second laptop was used for the prototype Operator System Interface to demonstrate techniques for human systems integration.

Augusta Systems developed the software to interface the swarming algorithms with the other system components including the cameras, the MeshNetwork communications network, the autopilot, the robot microcontroller, the GPS, and the payload control station. NewVectors developed the swarming algorithms operating on the payload computer and software for visualizing the pheromones and status of the swarming algorithms on the payload control station.

Figure 3 shows the architecture of the systems and the communications links among the components.

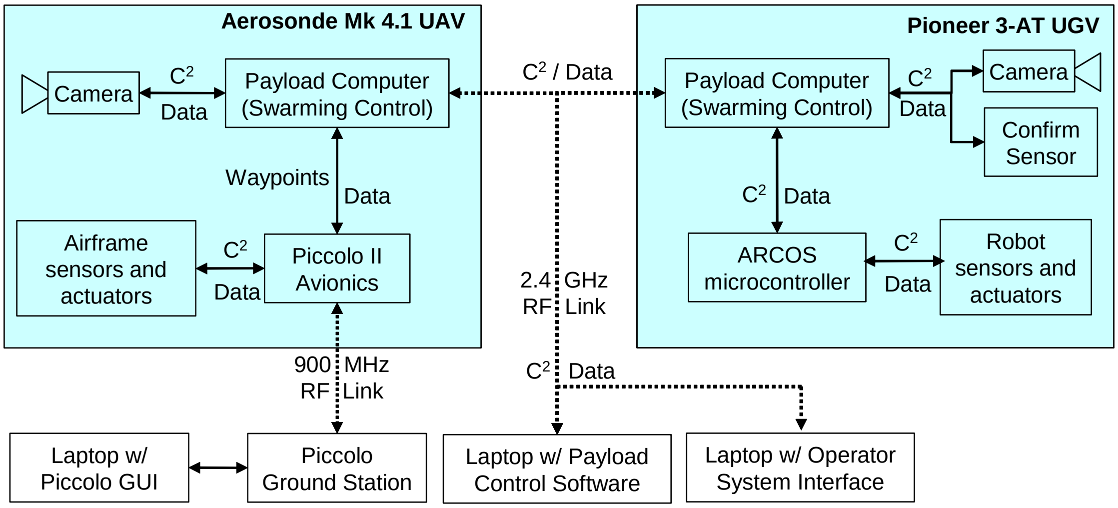

Figure 3. UAV and UGV System Configuration

5.2 Operator System Interface

The OSI was developed to evaluate techniques for enabling a single operator to monitor and manage multiple unmanned vehicles of different types in a RSTA application. The OSI displays advisories, cautions, and warnings; system status; time-stamped events; imagery from UAVs and UGVs, and a scalable bird's-eye view of the area of interest. This bird's-eye view includes the real-time position of all UxVs, as well as targets as they are located and identified (see Figure 4). This view also shows no-fly zones and high-priority zones for both UAVs and UGVs.

In addition to a traditional zoom-in/zoom-out feature, the OSI has a "zoom-to" feature. For example, if an operator is interested in UGV activity, he/she can use a simple pull-down menu to instantly zoom to a level that includes all UGVs. Other zoom-to options include "UAVs", "Targets", and "All Entities". Another unique feature of the OSI is the ability to instantly switch between symbology sets. Current options include MIL-STD 2525, MIL-STD 1787, MIL-STD 1477, and a custom set. Other symbol libraries can be easily added and once a standard symbol set is agreed upon, this interface will be able to incorporate it.

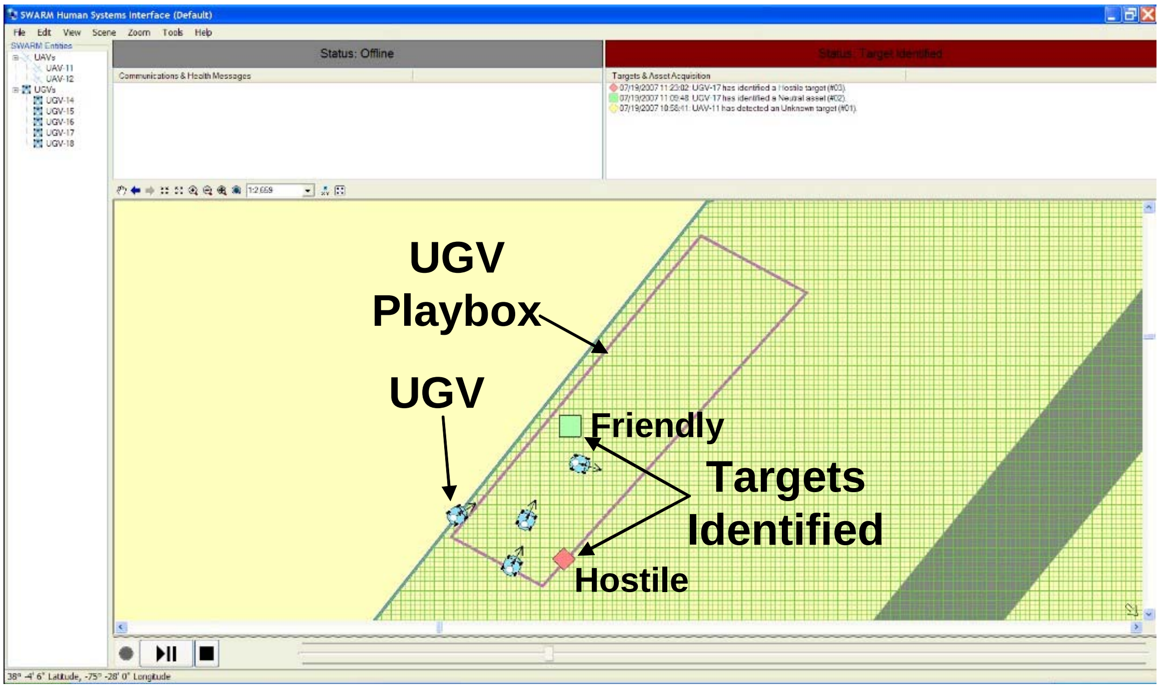

Figure 4. Graphical Operator System Interface Showing the UGV Playbox, UGVs, and Targets Identified

5.3 Swarming Algorithms Meet Reality

As expected the swarming algorithms required some adjustments to meet the demands of the demonstration. The safety precautions for the range imposed some new constraints on the algorithms. The flight plans developed by the swarming algorithm were downloaded to the autopilot where a human pilot had the option of rejecting the plan if it was deemed unacceptable for any reason. If a plan was rejected the aircraft would begin a 200m orbit around the last waypoint of the previous path while the swarming algorithm planned a revised route. Since the orbit could not violate the no-go zone, the last waypoint could not be closer than 200m to a no-go zone boundary. This ended up posing a particular problem for the demonstration since the ground targets were placed next to the observation area which was part of the no-go area. For the demo the swarming algorithm kept the UAVs 200m away from all no-go zones and the camera lens was zoomed out so it could still capture targets near the edge of the no-go zone. For future tests the algorithms will generate a special waypoint at the end of the path that is guaranteed to be more than 200m away from any no-go zone.

A highway crossed the area we planned to use for the flight demonstration (see Figure 5). NASA imposed a restriction that the UAV should never fly along or loiter over the road. This was solved by making a new kind of zone we called a *no loiter* zone along the road which was similar to the no-go zone. Like a no-go zone it had no Search pheromone and enforced the rule of no waypoint within an orbit radius of the zone, but it did not have no-go pheromone and did allow paths to cross over the zone. This simple change resulted in paths that would jump over the road and kept the UAVs from spending too much time near the road. This modification proved to be an elegant way to add new functionality to the system that can be used for road crossing and other similar requirements in the future. It also demonstrated the ease with which the algorithms can be adapted to new requirements.

5.4 The Demonstration

The flight tests were held at NASA's Wallops Island test range. Two Aerosonde UAVs were launched and placed under the control of the swarming algorithm along with four Pioneer UGVs. The UAVs were responsible for a 2.5km by 1km playbox, while the UGVs were responsible for a smaller 250m by 75m portion of that playbox (see Figure 5). Four targets were placed within the UGV playbox and two targets just outside that playbox but still within the UAV playbox.

Both the UAVs and UGVs executed the swarming algorithm described above except that the altitude and cross wind cost factors were not included since they were not required for the demonstration. A simplified ATR algorithm was implemented on the UAVs. When a UAV or UGV sensor viewed an area the Search pheromone was removed. When the ATR on the UAV identified a friendly target (a white circle, see Figure 6) the target location was designated with a box and the image sent to the OSI. When the UAV's ATR detected an unknown target (a white cross) they deposited Target Recognized pheromone at the detected target's location. This attracted the UGVs which possessed the necessary target identification sensor: an RF receiver detecting an RF transmitter embedded in the targets. UGVs needed to be within 6-8 feet of the target to pick up the RF signal to identify the target. Once a UGV identified a target it was reported to the OSI and the rest of the swarm so that further sensor hits on that target would be ignored. The UAV's ground projection of target location was within 50 meters of the actual location, a function of GPS error and UAV avionics error. This error wasn't a problem since the Target Recognized pheromone would propagate and the UGVs would survey around the location estimate until it found the actual target.

On the day of the demonstration the swarming algorithm had to deal with a few new surprises. All six UxVs were successfully launched, but payload communications with one of the UAVs failed. Without any human intervention the second UAV automatically adapted to the missing UAV and surveyed the area by itself. The simplified ATR on the UAV performed as expected and the images of detected friendly targets were sent to the OSI for display. The UAV was supposed to find the unknown targets within the UGV playbox as well. However, due to the size of the UAV playbox that had to be covered by the one remaining UAV and the need to stay 200m away from the targets within the UGV playbox the UAV did not have time to survey that area and missed all those targets. Still, without the expected help of the UAV, the UGVs' normal swarming activity brought them within the requisite range of 6-8 feet to find and identify three out of the four targets placed within the 200,000 square foot playbox. Finally, prior to the demonstration, the acoustic collision detection sensors on the UGVs started generating spurious contacts. The sensors were turned off for the demonstration so the collision prevention function of the swarming algorithm was entirely responsible for guaranteeing that no two robots collided during the demonstration.

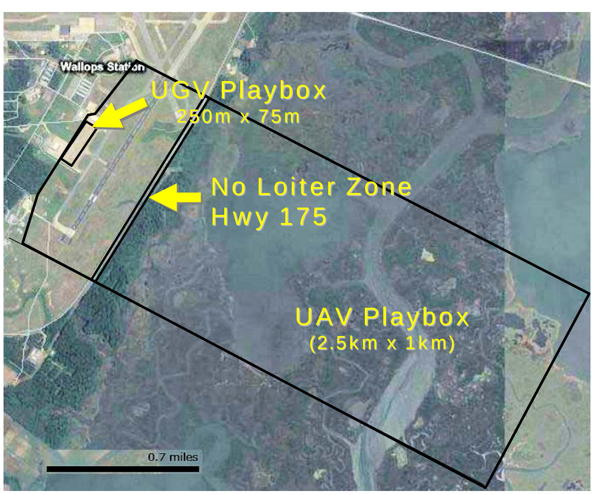

Figure 5. UAV and UGV playboxes

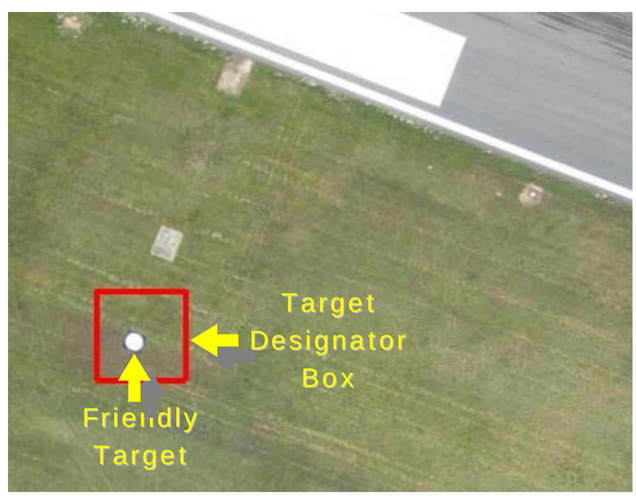

Figure 6. Image from UAV detecting friendly target.

5.5 Lessons Learned and Future Research

The swarming system was able to meet the demands of a realistic RSTA application. We demonstrated how autonomous swarming nodes with onboard swarm intelligence and ATR coupled over a mesh network successfully cooperated and self organized to achieve the overall mission objective. While the onboard ATR and ad hoc mesh networks performed as expected, further work is still required to demonstrate reliable onboard target recognition and mobile ad hoc networking to support swarm-based RSTA applications.

As is common with prototype systems, several failures were experienced before and during the demonstration. However, despite encountering unexpected problems, the swarming software performed remarkably well and was able to adapt and accommodate a number of the failures. When a failure happens in a single large platform, it can disable the whole vehicle possibly jeopardizing the entire mission. With robust swarming algorithms, when a failure happens in the swarm it may disable a whole vehicle but the swarm itself continues to function though in a slightly degraded mode.

Future work will focus on expanding into other application areas, adding new types of sensors into the swarming system, and adding humans as additional sensory nodes in a collaborative mixed initiative scenario all guided by the swarming algorithms.

6. CONCLUSION

During previous experiments the researchers noted the ease with which the swarming algorithms were adapted to new applications and requirements. Increasing sophistication was added with very little change in the basic algorithm. During this round of experiments many new factors had to be considered. Despite some of the unusual requirements placed on the algorithms by the hardware, the application, and safety precautions, the swarming algorithms were easily modified to meet all the demands of the application. Adding special rules to encode hard constraints was easily incorporated into the planning routines of the swarm. Special behaviors (such as restricted flight over highways) were accomplished using the standard pheromone mechanisms. The OSI demonstrated how one could monitor and visualize the results of multiple diverse swarming sensors building a common operating picture over a large area. While the focus of the demonstration was on the swarming algorithms, we were also able to demonstrate the integration of an onboard ATR algorithm coupled with an ad hoc mesh network for communications. In summary the onboard digital pheromone swarming algorithms successfully coordinated the behaviors of multiple air and ground vehicles in a realistic ATR application.

7. ACKNOWLEDGEMENTS

This paper is based on work supported by NAVAIR with Augusta Systems as the prime contractor. NAVAIR Public Release 08-067. Distribution: Statement A – “Approved for public release; distribution is unlimited.” 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.

REFERENCES

[1] OSD, Unmanned Systems Roadmap 2007-2032. 2007, Office of the Secretary of Defense. URL: http://www.acq.osd.mil/usd/Unmanned%20Systems%20Roadmap.2007-2032.pdf

[2] Franke, J. L., Zaychik, V., Spura, T. M., and Alves, E. E. Inverting the Operator/Vehicle Ratio: Approaches to Next Generation UAV Command and Control. in *Proceedings of AUVSI Unmanned Systems North America*. 2005. Baltimore, MD. URL: http://www.atl.lmco.com/papers/1261.pdf

[3] Reynolds, C. W. Flocks, Herds, and Schools: A Distributed Behavioral Model. in *Proceedings of Computer Graphics (ACM SIGGRAPH '87 Confererence Proceedings)*. 1987. Anaheim, CA: ACM. URL: http://www.cs.umu.se/kurser/TDBD12/HT02/papers/ReynoldsBoids1987.pdf

[4] Rimon, E. and Kodischek, D. E., Exact Robot Navigation Using Artificial Potential Functions. *IEEE Transactions on Robotics and Automation*, 1992. **8**(5): p. 501-518.

[5] Eun, Y. and Bang, H., Cooperative Control of Multiple Unmanned Aerial Vehicles Using the Potential Field Theory. *Journal of Aircraft*, 2006. **43**(6): p. 1805-1814.

[6] Sauter, J. A., Matthews, R., Parunak, H. V. D., and Brueckner, S. A. Evolving Adaptive Pheromone Path Planning Mechanisms. in *Proceedings of First International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2002)*. 2002. Bologna, Italy. URL: www.altarum.net/~vparunak/AAMAS02Evolution.pdf

[7] Sauter, J. A., Matthews, R., Parunak, H. V. D., and Brueckner, S. A. Performance of Digital Pheromones for Swarming Vehicle Control. in *Proceedings of Fourth International Joint Conference on Autonomous Agents and Multi-Agent Systems*. 2005. Utrecht, Netherlands: ACM. URL: http://www.newvectors.net/staff/parunakv/AAMAS05SwarmingDemo.pdf

[8] Sauter, J. A., Matthews, R., Parunak, H. V. D., and Brueckner, S. A., Effectiveness of Digital Pheromones Controlling Swarming Vehicles in Military Scenarios. *Journal of Aerospace Computing, Information, and Communication*, 2007. **4**(5): p. 753-769.

[9] Payton, D., Daily, M., Estowski, R., Howard, M., and Lee, C., Pheromone Robotics. *Journal Autonomous Robots*, 2001. **11**(3): p. 319-324.

[10] Panait, L. and Luke, S. A Pheromone-Based Utility Model for Collaborative Foraging. in *Proceedings of Third International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS'04)*. 2004. New York, NY.

[11] Dasgupta, P. Distributed Automatic Target Recognition Using Multiagent UAV Swarms. in *Proceedings of Fifth International Joint Conference on Autonomous Agents and Multiagent Systems*. 2006. Hakodate, Japan.

[12] 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*, 1959. **6**: p. 41-84.

[13] Parunak, H. V. D., Purcell, M., and O'Connell, R. Digital Pheromones for Autonomous Coordination of Swarming UAV's. in *Proceedings of First AIAA Unmanned Aerospace Vehicles, Systems,Technologies, and Operations Conference*. 2002. Norfolk, VA.

[14] Dubik, J., Richards, R., and Trinkle, G. Joint Concept Development and Experimentation. in *Proceedings of Swarming: Network Enabled C4ISR*. 2003. Tysons Corner, VA: ASD C3I.

[15] SMDC-BL-AS, Swarming Unmanned Aerial Vehicle (UAV) Limited Objective Experiment (LOE). 2001, U.S. Army Space and Missile Defense Battlelab, Studies and Analysis Divisio: Huntsville, AL.

[16] Parunak, H. V. D., Brueckner, S. A., and Odell, J. J. Swarming Coordination Of Multiple UAV's for Collaborative Sensing. in *Proceedings of Second AIAA Unmanned Unlimited Systems Technologies and Operations Aerospace Land and Sea Conference and Workshop & Exhibit*. 2003. San Diego, CA: AIAA.

[17] Parunak, H. V. D. and Brueckner, S. A. Stigmergic Learning for Self-Organizing Mobile Ad-Hoc Networks (MANET's). in *Proceedings of Third International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS'04)*. 2004. Columbia University, NY: ACM. URL: http://www.newvectors.net/staff/parunakv/AAMAS04MANET.pdf.

[18] Brueckner, S., Return from the Ant: Synthetic Ecosystems for Manufacturing Control. 2000, Humboldt University: Berlin. URL: http://dochost.rz.hu-berlin.de/dissertationen/brueckner-sven-2000-06-21/PDF/Brueckner.pdf

[19] Chen, X., Gong, S., Schmid, N. A., and Valenti, M. C. UAV Based Distributed ATR under Realistic Simulated Environmental Effects. in *Proceedings of 2007 SPIE Symposium on Defense and Security, Conf. on Signal Processing, Sensor Fusion, and Target Recognition XVI*. 2007: SPIE. URL: http://link.aip.org/link/?PSI/6567/65670E/1