

Heterogeneous Unmanned Vehicle Collaborative Control Demonstration

John A. Sauter, Robert S. Mathews,
Andrew Yinger

NewVectors
3520 Green Ct
Ann Arbor, MI 48105

{john.sauter, robert.mathews,
andrew.yinger}@newvectors.net

Joshua S. Robinson, John Moody

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

{jrobinson, jmoody}@augustasystems.com

Abstract

The use of swarming unmanned vehicles to support target detection and identification has been the subject of several research efforts. Cooperating land and air vehicles can support multiple sensor modalities providing pervasive and ubiquitous broad area sensor coverage. Coordination of air and land vehicles offers a wide range of stand-offs and viewing angles. However swarming systems are inherently diverse, decentralized, distributed, and dynamic making them nearly impossible to control through traditional command and control methods. Research in a class of algorithms based on digital pheromones has demonstrated their effectiveness in managing the complexity of swarming unmanned systems in simulated scenarios. This paper reports on the results of using digital pheromones in an onboard demonstration of unmanned land and air vehicles cooperating in a surveillance and automated target recognition scenario. Aerosonde air vehicles and Pioneer ground robots controlled by an onboard digital pheromone algorithm demonstrated the ability to self-organize and accomplish area surveillance, automatic target detection, sensor cueing, and automatic target recognition with no central processing or control and minimal operator input.

Introduction

Unmanned Vehicles (UxVs) have served to enhance and augment the Intelligence Surveillance and Reconnaissance (ISR) abilities of the military of the United States of America. 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. As the threats to global stability multiply and the United States relies, more than ever, on its intelligence gathering capabilities and on technologies that can provide superior tactical support to military personnel, the demand for sophisticated UxVs will only increase.

In developing the next generation of UxVs, the application of novel methodologies to UxV command and control will become more and more important as the number and mission criticality of deployed UxVs continues to increase. One of these novel approaches, long considered to possess some ideal characteristics for military application, is the use of swarming agents and their associated emergent behavior. This technique possesses the potential to provide minimal user intervention, a high level of robustness, and largely autonomous operation provided that suitable algorithms and techniques can be developed. The research described in this

paper provides the basic mathematical framework for swarming agent based Automatic Target Recognition (ATR). The technologies support the autonomous operation of multiple UxVs and satisfy avionic and sensor requirements for persistent tactical ISR missions.

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 [10], potential fields based on physics models [4, 11], and digital pheromones based on insect models [2, 5, 9, 12-14]. Digital pheromones are similar to potential fields, but they more naturally lend themselves to decentralized computation than potential fields.

Digital pheromones as described in [13, 14] 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 [8, 9] and coordination for unmanned vehicles[3, 15], positioning multi-sensor configurations[7], and maintaining line of sight communications in mobile ad hoc networks [6].

In vehicle control, the area of operation is tiled with a network of “place agents” which maintain the level of each flavor of pheromone present at that location and perform the basic pheromone operations. Digital pheromones support three primary operations, inspired by the dynamics of chemical pheromones.

1. *Deposit* – Deposits of a pheromone flavor, Φ_f are summed with the current amount of that flavor of pheromone located at that place. (Information fusion and aggregation).
2. *Evaporate* – Pheromones are evaporated over time. This serves to forget old information that is not refreshed. (Truth maintenance).
3. *Propagate* – Pheromones propagate from a place to its neighboring places. The act of propagation causes pheromone gradients to be formed. (Information diffusion).

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

$$s(\Phi_f, p, t) = E_f [(1 - G_f)(s(\Phi_f, p, t-1) + d(\Phi_f, p, t)) + g(\Phi_f, p, t)] \quad (1)$$

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 sum 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 Eq. (1) to each pheromone flavor once during every update cycle.

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

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

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

proportion depending on the parameter G_f and the total number of its neighbors.

Agents representing the unmanned platforms use the pheromone fields to plan vehicle paths. They are able to sense the level of pheromones in their current location and cells 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 simulated surveillance and target tracking application [12, 14]. 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 simple simulations and in several realistic battlefield scenarios. The algorithms demonstrated an ability to handle 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 online automatic target recognition using commercial unmanned platforms.

Requirements for Realistic Swarming ATR

This section considers how the UxV platform constraints and Automatic Target Recognition (ATR) algorithms 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.

ATR Requirements

ATR 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 getting the right sensor, at the right altitude, and the right orientation to the target in order for the sensor to collect the data

necessary for the ATR algorithm to successfully complete its processing.

The target detection and recognition algorithms were developed by West Virginia University (WVU). An optical sensor captures images of varying resolution and fidelity. Images may have noise, focus blur, and motion blur and targets may be partially occluded. The detection algorithm utilizes 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 [1].

The initial guidance from WVU provided the following parameters for this work:

1. The swarming UAVs and sensors shall capture 64 x 64 pixels on the target. Lower resolutions are allowable for target detection with higher resolutions for target recognition.
2. The swarming UAVs shall respond to specific requests from the ATR to capture images of defined resolutions from defined orientation and declination angles.
3. The camera shall have zoom capability typical for commercial cameras.
4. The camera shall be gimballed 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.
5. All targets are on the surface, immobile, and between 1m x 1m and 6m x 6m in size. This requirement and #1 above require a pixel density between 114 and 4096 pixels/m².
6. The UAV shall fly at an altitude between 150m and 300m (nominal 230m) at 25m/s.

From these requirements the following additional specifications can be derived:

7. A 35 mm 2 megapixel (1600 x 1200) sensor requires a lens with a 420mm focal

length to image 4096 pixels/m² from 300m altitude. Alternatively higher pixel sensors or lower altitudes could be used to obtain the highest pixel density required.

8. A 1600 x 1200 sensor flying at 25 m/s capturing 114 pixels/m² 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.

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 certain level of certainty. If the uncertainty remains high after processing the image additional images of the area will be required to confirm the presence or absence of a target.

UxV Platform Requirements

Previous experience with adapting digital pheromones to controlling UxVs [12, 14] identified and resolved issues related to communications, constrained UAV turning radius, and no-go areas restricting UxV movement. For this work we also looked at the following factors:

- 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. Hardware and software failures need to be accommodated and backup systems put in place.
- Errors in communication and positioning, can lead to dropped messages, missed updates, and inaccuracies in computing paths and avoiding collisions.
- Turns and climbs consume more fuel decreasing the effective range and time on station.
- Ground speed is affected by wind speed and direction which in turn impacts ground coverage and image capture rates. Increasing the ground speed while flying downwind might increase the image capture rate above the processing capability of the ATR algorithm.

Description of the Modified Swarming Algorithm

This section describes the modified swarming algorithm to accommodate the ATR algorithm and unmanned platform requirements. Each vehicle maintains its own version of the pheromone map. There are four primary pheromones:

1. Φ_{\ominus} Search pheromone that attracts UxVs to areas that have not been searched yet. In more sophisticated versions this can be a measure of the ATR uncertainty.
2. Φ_r ATR Request pheromone deposited by the ATR algorithm that has detected a target and needs additional images for identification (UAV) or target confirmation (UGV).
3. Φ_x No-go pheromone deposited in areas that represent no-fly zones
4. Φ_v Vehicle path pheromone deposited along the planned path for each vehicle.

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 to areas where there is the most need for surveillance (highest uncertainty, or specific ATR request).
2. Prefer to move in straight lines to conserve fuel (UAVs) or time (UGVs).
3. Prefer to fly at constant altitude or move on level ground to conserve fuel.
4. Prefer to fly cross wind to maintain ground speed for ATR (UAVs).
5. Stay away from other vehicles to avoid collisions.
6. Stay away from no-go zones.

These high level objectives are implemented in the form of a benefit to cost ratio. The benefit for a path will be the sum of

the Search pheromone and ATR Request pheromone present in all the cells \aleph within the field of view of the sensor along the path:

$$B_p = \sum_{\aleph} (\theta\Phi_{\ominus} + \rho\Phi_r) \quad (3)$$

where θ and ρ are tuning constants that depend on the sensor capabilities of the UAV and its ability to reduce uncertainty or respond to specific ATR requests.

The cost for a path has three elements.

The fuel cost C_f includes the cost for a heading change, a change in altitude or elevation, and movement parallel or perpendicular to the wind. The path cost also includes two other factors 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 vehicle in the area is planning on going so that other vehicles can avoid those areas. The pheromone Φ_x is deposited in the no-go zones and propagates a short way into the no-go zones. This pheromone provides a "soft" boundary so that UAVs can sense when they are nearing a no-go boundary and begin maneuvers to avoid it. These two pheromones are summed along with the fuel cost for each segment of the move \aleph to arrive at a total cost for the path:

$$C_p = \sum_{\aleph} (\phi C_f + \nu\Phi_v + \chi\Phi_x) \quad (4)$$

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

$$T_p = \frac{\sum_{\aleph} (\theta\Phi_{\ominus} + \rho\Phi_r) + k}{\sum_{\aleph} (\phi C_f + \nu\Phi_v + \chi\Phi_x) + k} \quad (5)$$

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

simple approach repeatedly applies equation (5) to all cells reachable within a certain radius of the current point in the path and either picks the best point as the next waypoint or selects one stochastically using a weighted roulette wheel. Paths that end up leading to a collision or that enter no-go zones are eliminated from consideration. Multiple paths are evaluated in their entirety using equation (5) to select the final path.

Flight and Ground Tests

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

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



Figure 2: Pioneer 3-AT Ground Robot

Augusta Systems equipped both the UAV and the UGV with a Pentium-M payload computer running Windows XP Embedded on 1 GB Compact Flash and a MeshNetworks WMC6300 2.4 GHz subscriber card to support transmission of imagery and swarming data. A single laptop on the MeshNetwork is used as a “payload control station” for monitoring and manually controlling the UxVs.

Augusta Systems developed the interfaces between the swarming algorithms and the cameras, the MeshNetwork

communications network, the Piccolo autopilot, the robot microcontroller, the UGV GPS, and the payload control station. NewVectors developed the swarming algorithms operating on the payload computer and software for visualizing the 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.

The Test Scenario

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 three Pioneer UGVs. 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. When the ATR algorithm on the UAV detected a target an ATR Photo Request pheromone was deposited at that location. This would attract the same or another nearby UAV to make another pass over the target location to get another picture. Once the target was identified an ATR Confirmation Request pheromone was deposited which would attract a nearby UGV

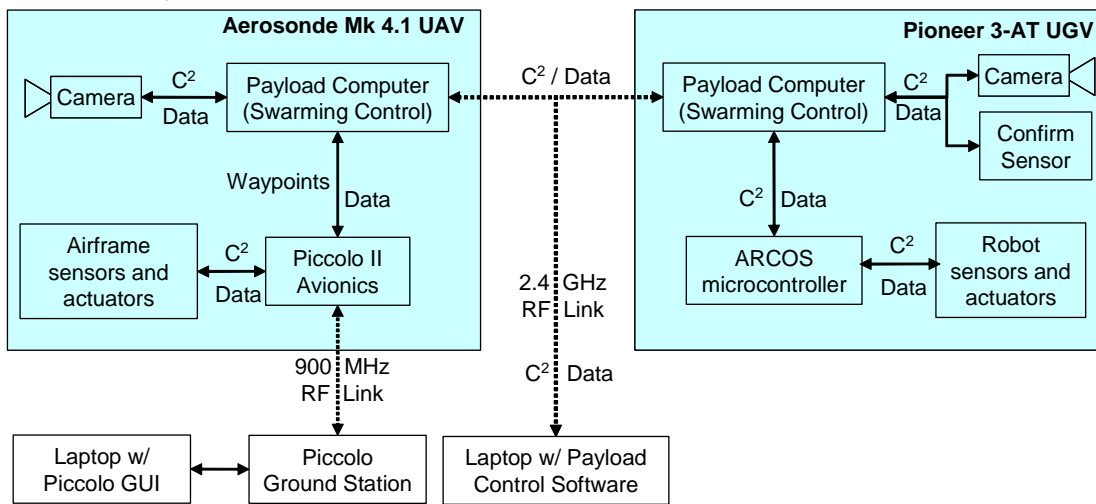


Figure 3: UAV and UGV System Configuration

to come over and confirm the target with its “confirmation” sensor.

The swarming algorithm generates a path plan as a series of five waypoints representing on average 90 sec of flight time for the UAV or 100 seconds of movement for the UGV. Flight plans for the UAV were first stored in the Piccolo autopilot and a human pilot was given at least 20 seconds to review and override it otherwise it was automatically activated. If the plan was overridden, the swarming algorithm continued to propose new plans until one was accepted by the pilot while the UAV orbited the last waypoint.

Lessons Learned

The Piccolo autopilot design caused two issues with the swarming algorithm. Initially we assumed that the autopilot would fly to the waypoint and then make a heading change to fly straight to the next waypoint. The swarming algorithm could then pick waypoints that included specific areas to be photographed. In practice the Piccolo autopilot attempts to maintain the flight path rather than flying over the waypoint. This means that the smoothest way to control the aircraft is to begin the turn before it reaches the waypoint so that it can approach the next segment without having to execute an “S” turn. This required a minor adjustment in the path planning algorithms to plan the turn prior to the waypoint rather than after.

The Piccolo also has a communications timeout timer, set by default to 60 seconds. If it receives no communications on any of its serial ports it assumes it has lost communications with its ground station and activates its Lost Comms flight plan. Unfortunately the payload computer communicates with the Piccolo on one of those serial ports, so each command the swarming algorithm sends to the Piccolo autopilot causes a reset of the Lost Comms timer possibly masking a true communications blackout with the ground station. To address this problem, the Lost

Comms timer was set to 30 seconds and the swarm to Piccolo interface ensured that there was 35 seconds of communications silence from the payload computer before sending the next command to the Piccolo. This resulted in some unavoidable delays between when the flight plan was planned and it could be activated on the Piccolo.

The safety precautions for the range imposed some unusual constraints on the algorithms as well. Since the last waypoint must also be able to serve as a safe orbit point in addition to a normal waypoint it must not be in a location where the orbit would cross a no-go zone. Since the orbit radius is 200m that meant that the swarming algorithm would have to avoid flying closer than 200m to the no-go boundary. This ended up posing a particular problem for the demonstration since the ground targets were supposed to be placed next to the observation area which was part of the no-go area. For the demo the lens was zoomed out to capture the targets at the edge of the FOV. For subsequent work the algorithms will generate a special 6th waypoint that is always 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. 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 had no Search pheromone and enforced the rule of no waypoint within an orbit radius of the zone, but did not include no-go pheromone and did allow paths to cross over the zone. This resulted in paths that would jump over the road and kept the UAVs from spending too much time near the road.

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 basic swarming algorithms continue to operate as before. Adding special rules to encode hard constraints was easily incorporated into the planning routines of the swarm. Creating special behaviors (such as restricted flight over highways) was also easy to develop using the standard pheromone mechanisms. Digital pheromone swarming algorithms have successfully coordinated the behaviors of multiple air and ground vehicles in a realistic ATR application.

Acknowledgements

This paper is based on work supported by NAVAIR. 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] Chen, X., Gong, S., Schmid, N. A., and Valenti, M. C. UAV Based Distributed ATR under Realistic Simulated Environmental Effects. in *Proceedings of Signal Processing, Sensor Fusion, and Target Recognition XVI*. 2007: SPIE. URL: <http://link.aip.org/link/?PSI/6567/65670E/1>
- [2] 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.
- [3] 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.
- [4] 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.
- [5] 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.
- [6] 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>.
- [7] 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.
- [8] 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.
- [9] Payton, D., Daily, M., Estowski, R., Howard, M., and Lee, C., Pheromone Robotics. *Journal Autonomous Robots*, 2001. **11**(3): p. 319-324.
- [10] Reynolds, C. W. Flocks, Herds, and Schools: A Distributed Behavioral Model. in *Proceedings of Computer Graphics (ACM SIGGRAPH '87 Conference Proceedings)*. 1987. Anaheim, CA: ACM. URL:

- <http://www.cs.umu.se/kurser/TDBD12/H T02/papers/ReynoldsBoids1987.pdf>
- [11] 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.
- [12] 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.
- [13] 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/AAMAS02 Evolution.pdf
- [14] 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/parunak v/AAMAS05SwarmingDemo.pdf>
- [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 Division: Huntsville, AL.