

Effectiveness of Digital Pheromones Controlling Swarming Vehicles in Military Scenarios

John A. Sauter,* Robert Matthews,† H. Van Dyke Parunak,‡ and Sven A. Brueckner§
NewVectors LLC, Ann Arbor, MI 48105

The use of digital pheromones for controlling and coordinating swarms of unmanned vehicles has been studied under various conditions. This paper describes experiments that demonstrate their effectiveness in several militarily significant scenarios in simulation. The scenarios are described along with the results comparing the performance of swarming to traditional military tactics. A demonstration was conducted using these same algorithms to coordinate unmanned vehicles in a simulated exercise. Two air vehicles 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. 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.

Nomenclature

$P = \{p_i\}$	set of place agents
$N : P \rightarrow P$	neighbor relation between place agents. Thus the place agents form an asymmetric multigraph
$s(\Phi_f, p, t)$	strength of pheromone flavor f at place agent p and time t
$d(\Phi_f, p, t)$	sum of external deposits of pheromone flavor f within the interval $(t - 1, t]$ at place agent p
$g(\Phi_f, p, t)$	propagated input of pheromone flavor f at time t to place agent p
$E_f \in (0, 1)$	evaporation factor for flavor f
$G_f \in [0, 1)$	propagation factor for flavor f
T_f	threshold below which $s(\Phi_f, p, t)$ is set to zero
$V(p)$	Scaler value of a place agent’s “attractiveness”
α, β	weighting factors for evaluating response to pheromones
$P_{p'}$	probability that place p' is selected for a move

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.¹⁻³ Examples include path formation, nest

Presented as paper 2005-7046 at the Infotech@Aerospace 2005 Conference and Exhibit, Arlington, Virginia, 26 September 2005; received 7 August 2006; accepted for publication 1 March 2007. Copyright © 2007 by the American Institute of Aeronautics and Astronautics, Inc. All rights reserved. Copies of this paper may be made for personal or internal use, on condition that the copier pay the \$10.00 per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923; include the code 1542-9423/04 \$10.00 in correspondence with the CCC.

* Group Leader, Emerging Markets Group, john.sauter@newvectors.net

† Senior Systems Engineer, Emerging Markets Group, robert.matthews@newvectors.net

‡ Chief Scientist, Emerging Markets Group, van.parunak@newvectors.net

§ Senior Systems Engineer, Emerging Markets Group, sven.brueckner@newvectors.net

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.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, less attention has been given to the application of these approaches to the control and coordination of unmanned vehicles in military applications. 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 describe the requirements that military applications place on swarming control. We then review the common approaches that have been used and compare those to digital pheromones. This is followed by a description of the digital pheromone algorithm and how it is used to control UAVs. We then describe three experiments and 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. Requirements for Swarming UAV Control

Effective Command and Control (C2) for swarming UAV's has four requirements, which can be summarized as four D's: Diverse, Distributed, Decentralized, and Dynamic.

Diverse—A system to control swarming UAV's must be able to handle a wide range of functions such as surveillance, target recognition, tracking, and targeting. It must integrate diverse kinds of information from electro-optical, chemical, biological, radio frequency, and other sensors both onboard and offboard. It must coordinate the behavior of a diverse population of manned and unmanned air, land, and sea vehicles with varying capabilities. It must be capable of handling many kinds of fixed and mobile targets, in different environments (desert, mountain, forest, urban) and weather conditions. The complexity this introduces can overwhelm most traditional control methods.

Distributed—By its very nature, a swarm of UAVs is a distributed system. The entities in the swarm are distributed geographically. Information is generated and used in a distributive fashion throughout the swarm. The size of the swarm increases the complexity of the potential interactions and the size of the state space that must be managed. This can overwhelm some algorithms. The degree of connectedness of a distributed swarm will vary widely, but despite the promise of ubiquitous communications, there will never be unlimited bandwidth and range for swarming vehicles. Small UAVs and UGVs are often limited to line-of-sight communications and underwater vehicles have even greater restrictions on distance and bandwidth. This decreases the degree of connectedness among the units in the swarm that may not be in constant communications contact with either a central node or each other. Swarming mechanisms must be able to control large numbers of distributed entities with limited connections.

Decentralized—Currently each unmanned vehicle has at least one and usually two or three human controllers responsible for it. As we move to smaller and more numerous vehicles, the old model of control breaks down. Not only are the human costs too high, but the coordination complexity among the human controllers increases exponentially with the size of the swarm. The limitations of communications also demand the swarm be organized with a certain amount of autonomy. Each unit must be able to make decisions based on local information. Centralized control leads to bottlenecks for timely processing of information and execution of command and control. It results in a single point of failure or attack (the central communications pipe or C2 node) that could disable the entire swarm. Any method that requires the centralized processing of information in order to control the swarm will be ineffective. Methods that require extensive computational demands may be infeasible on small platforms with limited power and processing capability.

Dynamic—The battlespace is an uncertain and rapidly changing environment. Enemy forces will try to adapt to the swarm and attempt to deceive and disrupt it. The swarm will only have imperfect and inconsistent knowledge about the enemy's location and intent. The swarm itself may change dynamically as new units are added and others are disabled or destroyed. The nonlinear nature of warfare⁷ generates unexpected changes in the situation. Moltke's thesis that "no battle plan survives contact with the enemy" demands that swarming control systems must be ready to constantly change and adapt to their environment. Approaches that require complete knowledge or lengthy computations for replanning and optimization will not be viable.

III. Approaches to Swarming Control

There are several methods available for controlling and coordinating swarms of unmanned vehicles that attempt to address the four D's described above. Parunak⁸ reviews the major classes of algorithms that have been applied to the Command and Control of multiple robotic entities. There is abundant literature on centralized control schemes, but based on the observations above we will limit our discussion to algorithms that can support distributed, decentralized computation. Some work has been done on developing decentralized versions of centralized control strategies (such as distributed model predictive control⁹), but most of the work in distributed vehicle control involves various kinds of field-based mechanisms. In field-based systems a scalar field is generated by a combination of attracting and repelling elements, and the agents respond to those forces or follow gradients in this field. Within this class of algorithms are particle systems based on Reynold's model^{10,11}, potential fields based on physics models¹²⁻¹⁶, and digital pheromones based on insect models.¹⁷⁻²³ 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 Grasse²⁴ 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.

These methods have been used with varying success in addressing the four D's of swarming control. Table 1 summarizes the strengths and weaknesses of the three field based mechanisms against each of those requirements. A common question with stigmergic algorithms is how well they perform against traditional methods. Can they reach an optimal solution and can they do so in less time than other techniques? Unfortunately these questions are not easily answered. First, most of the benchmark problems do not themselves exhibit the four D's and hence are poor candidates for evaluating algorithms designed for those requirements. Standard problems typically exhibit limited diversity (so that simple comparisons are possible), lend themselves to centralized computation (since they are primarily used to evaluate centralized algorithms) and are almost always static. Even the meaning of optimality in a highly dynamic military environment is not well defined. In time-critical military scenarios where loss of life could result, having a timely solution that is good enough is always preferable to one that cannot be re-computed as quickly as the problem changes.

One would expect that stigmergic algorithms would not perform as well as traditional approaches on static benchmark problems. Still some researchers have begun to make comparisons between stigmergic and optimal solutions to dynamic problems with promising results.²⁵ However, these simple benchmark applications do little to convince military planners of the value of swarming technology. Rather than using simple benchmarks this paper compares the use of digital pheromones against models of traditional control in realistic military scenarios exhibiting the four D's.

A. Vehicle Control with Digital Pheromones

Digital pheromones as described in^{17,18} are modeled on the pheromone fields that many social insects use to coordinate their behavior. A digital pheromone represents information about the system. Different "flavors" of pheromones convey different kinds of information. Brueckner²⁶ proves certain convergence and stability theorems for digital pheromone fields following this model. They have been used to support a variety of traditional swarming functions including path planning^{21,27} and coordination for unpiloted vehicles,^{28,29} positioning multi-sensor configurations,³⁰ and maintaining line of sight communications in mobile ad hoc networks.³¹

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).

Table 1 Comparison of stigmergic swarming mechanisms against four D requirements.

Mechanism	Diverse	Distributed	Decentralized	Dynamic
Particle systems	Poor: Most research focused on one behavior (e.g. flocking) among similar entities	Excellent: each particle maintains local state. Information requirements are local. Interactions are local. Good scalability with size of swarm	Excellent: Each particle operates under its own rules independently of any central control or computation	Excellent: particles adapt dynamically as numbers and, environmental changes
Potential Fields	Good: Different kinds of entities can generate different fields. Fields can be constructed to accomplish many kinds of tasks	Poor to Good: Most systems rely on global knowledge and interactions, but some use only local information and interactions. Good scalability with size of swarm	Excellent: each entity makes decisions based on their own sensing of the fields and internal rules. Some field representations are maintained centrally, but decentralized, distributed versions are available	Poor to Good: some potential fields computed in their entirety before moves are planned
Digital Pheromones	Excellent: Any number of different kinds of entities and behaviors accommodated through different agent rules and pheromone flavors	Excellent: All information exchange and interactions among the agents are local. Field maintenance occurs through local interactions across a distributed architecture. Good scalability with size of swarm	Excellent: No centralized information processing or control is used. Each agent operates under its own rules	Excellent: Fields rapidly integrate new information (deposits and propagation) and purge old information (evaporation). Stochastic decision processes allow agents to deal with partial, uncertain information. Any-time nature of algorithm yields results that can adapt to rapidly changing environments

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. There are two primary equations governing the maintenance of the pheromone field. The first equation describes the evolution of the strength of a single pheromone flavor at a given place agent.

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

E_f models evaporation of pheromone, $1 - G_f$ calculates the amount remaining after propagation to its neighbors, $s(\Phi_f, p, t - 1)$ represents the amount of pheromone from the previous cycle, $d(\Phi_f, p, t)$ represents the total deposits made since the last update cycle and $g(\Phi_f, p, t)$ represents the total pheromone propagated in from all the neighbors of p . Each place agent applies Eq. (1) to each pheromone flavor once during every update cycle.

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

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

This equation states that 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.

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 local C2 nodes according to their area of responsibility. The swarming platforms only need to communicate with the local UGS's 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 and gossip²³ or synchronization³² mechanisms used to propagate updates through the swarm. 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. Based on the pheromones they sense they make decisions about what they will do next and where they will go. Depending on the state of the Walker, some flavors of pheromone are attractive, 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. Normally a simple algebraic equation is used to evaluate the levels of the different pheromone flavors in the current place agent and the neighboring place agents to establish a score, $V(p)$ for each possible destination. Infeasible locations due to vehicle maneuverability constraints are eliminated from consideration. The probability that a given place is chosen as the next move is given by the following equation.

$$P_{p'} = \frac{V(p')}{\sum_{p' \in N(p)} V(p')} \quad (3)$$

This is commonly known as a weighted roulette wheel. UAVs typically plan multiple steps into the future by repeating this process.

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.

B. Example Surveillance Application

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 need to be continuously monitored. Each area has differing priorities for surveillance that may change as new information arrives. The UAVs must configure themselves to survey the areas according to that priority (higher frequency of revisits for higher priority AOIs) regardless of the




Fig. 1 Attractive and repulsive pheromones for surveillance, 1. Surveillance area deposits attractive pheromone, 2. Walker deposits repulsive pheromone, 3. Pheromone infrastructure propagates both attractive and repulsive pheromone to form gradient, 4. UAV climbs net gradient, withdrawing attractive pheromone.

number of UAVs, the number, size or shape of the AOIs, and their varying priority. This particular application can be performed with only two flavors of pheromone. An attractive pheromone, Φ_l is deposited on each place agent that requires surveillance. The amount of the deposit depends on the place agent's surveillance priority or frequency requirements. When it was last surveyed determines when new deposits are made. By controlling the amount and frequency of deposits, the frequency of surveillance can be varied. A repulsive pheromone, Φ_v is deposited by vehicles along their planned paths to deconflict the airspace and avoid duplication of efforts (see Fig. 1). The equation that calculates the score used in the weighted roulette wheel movement decision is simply:

$$V(p) = \alpha s(\Phi_l, p) - \beta s(\Phi_v, p) \quad (4)$$

where α and β , typically set to 1, can be used to tune the system's response.

The user starts by outlining each AOI on a map and assigning a surveillance priority or frequency 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 Φ_l pheromone at that place agent ($d(\Phi_l, p)$) as long as it remains on. No Φ_l pheromone is deposited in the places agents outside any AOI, but they will propagate and evaporate any pheromone that propagates to them from the AOIs according to Eqs. (1) and (2).

A Walker agent is created for each UAV. The Walker agent senses the level of Φ_l and Φ_v pheromones ($s(\Phi_l, p, t)$ and $s(\Phi_v, p, t)$) in the current and neighboring place agents scores each using Eq. (4), and uses the scores to select the next place to move by spinning the weighted roulette wheel. It then places a deposit of Φ_v pheromone ($d(\Phi_v, p)$) on the place agent representing its next move, removes all Φ_l pheromone from that place agent, and turns off the Φ_l pheromone pump found there. The pump will automatically restart after a delay inversely proportional to the priority the user placed on surveying that AOI.

The best parameter settings determined through simulation trade studies for this application are as follows:

$$\begin{aligned} d(\Phi_l, p) &= 10 \\ E_l &= 0.005 \\ G_l &= 0.4 \\ d(\Phi_v, p) &= 10 \\ E_v &= 0.1 \\ G_v &= 0.4 \end{aligned}$$

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 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.

As described in,¹⁸ coverage metrics for the algorithm described above are good. In one experiment 30 UAVs moving at 90 kph were able to find and begin coverage of 10 AOIs 1 km² located randomly in a 400 km² 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 Φ_l pheromone propagates outside the AOIs where it is deposited, 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.

One of the challenges with all stigmergic algorithms is how to define the rules and behaviors of the individual agents to achieve the overall performance or objective of the system. In the example above, one must determine the deposit amount and frequency and E_f and G_f of the two pheromones to achieve the desired surveillance rates over the different AOIs. This was done through simulation-based trade studies that establish the relationship between individual parameters (such as deposit rate) and a global metric or behavior (such as revisit frequency). Our experience has shown that these relationships can be well defined. If they are difficult to determine ahead of time, then on-line tuning algorithms such as genetic algorithms¹⁷ have proven successful in automatically tuning the system to achieve the level of performance desired.

IV. OASD Study Results

The Office of the Assistant Secretary of Defense sponsored a study entitled “Swarming Concept Development and Utility Study.” 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 Operation Iraqi Freedom (OIF) and Operation Enduring Freedom (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 friendly (Blue) reconnaissance detecting a meeting of high-level terrorist leaders (Red), 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 simultaneous military operations and the ability to dynamically retask multiple swarms to support those operations. OPSIT 2 is based on a large-scale assault of a Red Weapons of Mass Effect (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 Extended Air Defense Test Bed (EADTB) or System Effectiveness Analysis Simulation (SEAS) simulation platforms. A Base case was constructed using manned and unmanned platforms under traditional control approaches. This was compared to a Swarming case using the same scenario but with swarming entities controlled by a digital pheromone algorithm. In some scenarios 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.

The following describes each of the OPSITs in more detail. For each OPSIT, the scenario is described, along with the experimental setup, followed by a summary of the key results.

A. OPSIT 1/4 Terrorist Meeting Surveillance and Power Plant Protection

1. Scenario Description

In OPSIT 1/4, a neighborhood suspected of being the location for a high-level terrorist meeting is placed under surveillance. Thirty kilometers away the area around a power plant reconstruction project is also under continuous surveillance to protect the civilian reconstruction crew. 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 terrorist

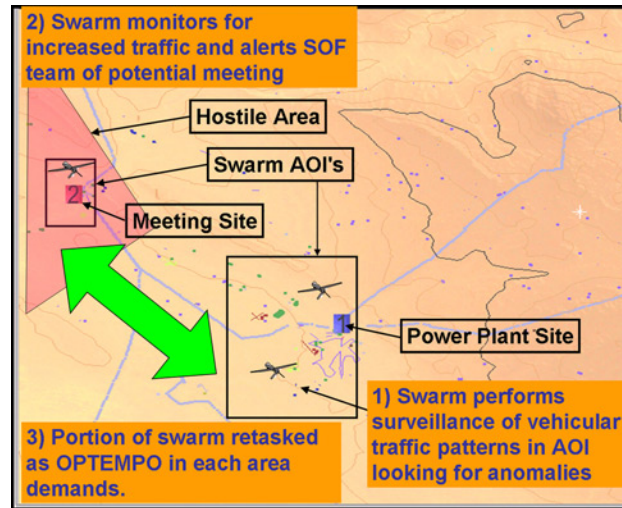


Fig. 2 Opsit 1/4—Terrorist meeting and power plant attack.

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 terrorist meeting begins. A Special Operations Force (SOF) team is deployed to disrupt the meeting and the swarming assets arriving from the power plant AOI are tasked to track any fleeing personnel or vehicles from the meeting site.

2. *Experimental Setup*

The force structures used for OPSIT 1/4 is shown in Table 2. Excursions were tested varying the number of UAVs to determine the impact of the swarm size on performance. Experiments were also performed using a fixed pattern search algorithm to control the UAVs rather than digital pheromones.

3. *Key Results*

In this scenario the swarming assets (UAVs) were able to detect the terrorist meeting site 16 minutes earlier than the Base case (using unattended ground sensors and Global Hawk alone). The swarm provided more regular coverage of the area than was possible with Global Hawk. 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 swarming UAVs were able to track the targets whereas assets flying fixed patterns could not deviate from their flight plan.

Table 2 Opsit 1/4 experimental setup.

Red forces	Blue base case	Blue swarming case
Terrorist units (Power plant)	ISR Assets	ISR Assets
8—terrorist trucks/vehicles	1—JSTARS	1—JSTARS
1—truck bomb vehicle	1—Global hawk (Power plant site)	10–50—UAVs
	2—SOF teams (Meeting site)	2—SOF teams (Meeting site)
Terrorist units (Meeting site)	2—UGS (Meeting site)	2—UGS (Meeting site)
8—terrorist trucks/vehicles		
24—terrorist dismounts	Attack assets (Meeting site)	Attack assets (Meeting site)
	Assault teams w/ Helicopters	Assault teams w/ Helicopters
Other		
112—Civilian vehicles	Other	Other
	JAOC	JAOC

This demonstrated that the improved performance did not come from simply having more sensors in the field but resulted from the ability of the swarm to adapt to the target movement and simultaneously track detected targets while continuing to survey for new ones.

B. OPSIT 2 Theater WME Engagement

1. Scenario Description

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 under the protection of Unmanned Surface Vehicles (USVs) 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 tactical ballistic missiles (BM-21s), Surface to Air Missiles (SAMs) and chemical munitions. They will attempt to deny Blue access to Axis “Jack” and force them on the longer “Jill” route. 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 launched from shore.

2. Experimental Setup

The force structure for this scenario is more involved. The force structure for the Base case is shown in Fig. 4. The Swarming case added 25, 50, or 100 UAVs. Eight of the UAVs were always dedicated to surveying Area C (see Fig. 3) to find the BM-21’s and SAM launchers. The remaining units were split between Areas A and B (see Fig. 3) to look for chemical munitions and enemy force activity along the attack routes. The Class I and Class II UAVs remained with the forces to provide force protection. The Class III UAVs surveyed Areas A, B, and C cooperating with the Swarming UAVs in the Swarming case. Individual UUVs and USVs were not modeled. Only their effects were modeled as determined by separate studies performed by Johns Hopkins University. In one of the experiments, the eight swarming UAVs in Area C were set to perform a fixed pattern search rather than employing the swarming algorithm. The remaining swarming UAVs in Areas A and B continued to use the swarming algorithm.

3. Key Results

In this experiment, the swarming UAV assets were able to accurately map the extent of the chemical contamination caused by Red’s chemical weapons used to block Axis “Jack” allowing the Blue commander to maneuver around the plume 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.

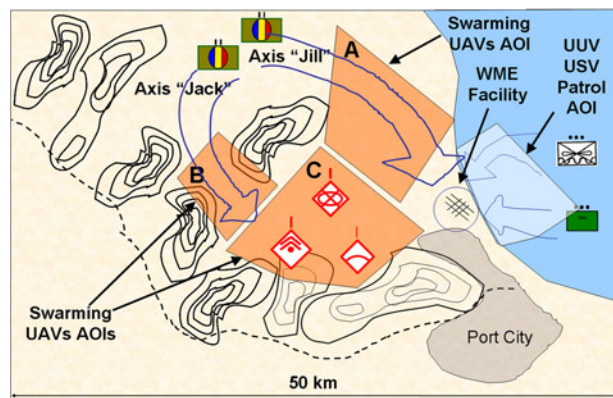


Fig. 3 OPSIT 2—combined land and sea assault on WME facility. 8 swarming UAVs survey Area C and the rest are split between Areas A and B.

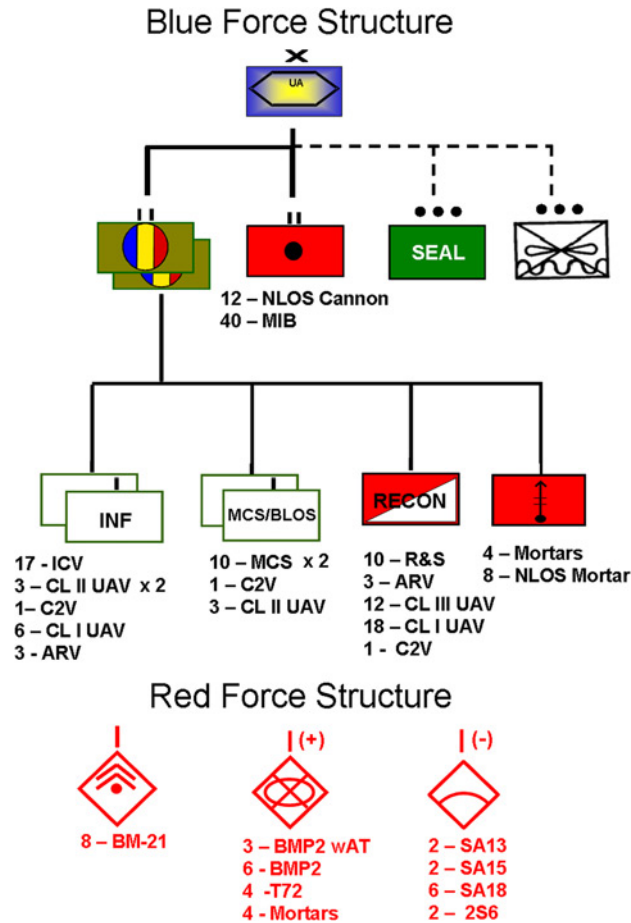


Fig. 4 Opsit 2 force structure.

There were twice as many Blue casualties in the Base case as the Swarming case (Table 4). 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 two hours 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.

Table 3 OPSIT 2 results with varying swarm size

	Base	Swarm		
Number swarm UAVs	0	25	50	100
Increase in total red losses over Base	–	31%	55%	55%
Increase in BM-21 and SAM launcher losses over base	–	2.5x	2.5x	2.5x
High priority target (HPT) Engagements	65	187	216	216
Hours to 65 HPT engagements	2.2	0.35	0.3	0.3
WME trucks intercepted	0	100%	100%	100%

Table 4 OPSIT 2 results swarming vs. non-swarming

	Blue	Swarm	Fixed pattern
Number of swarm UAVs	0	8	8
Total blue losses	39	16	25
Loss exchange ratio	1.2	3.1	2

Class III UAV survivability was increased three-fold since the swarming entities were able to quickly locate Red air defense artillery, which Blue destroyed. In the Swarming case the additional surviving Class III UAVs improved the situational awareness of Blue, reduced their casualties, and increased their OPTEMPO.

Other performance measures are listed in Tables 3 and 4. Table 3 shows some of the results from the experiments that varied the number of units in the swarm in Areas A and B. The experiments showed that even with only 25 swarming UAVs there were significant performance improvements over the Base case. The change in performance from 50 to 100 UAVs was negligible.

Table 4 compares the effect of using a swarming algorithm versus a fixed pattern search in Area C. Using a Fixed Pattern search Blue losses increased 57% over the swarming algorithm demonstrating the effectiveness of the digital pheromones over non-adaptive control approaches. This was primarily due to the ability of the swarming algorithms to perform both tracking and surveillance activities. The Loss Exchange Ratio (the ratio of Total Red Losses to Total Blue Losses) improved from 1.2 to 3.1 using the swarming algorithms.

C. OPSIT 3 Terrorist Suppression in Mountainous Terrain

1. Scenario Description

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 (Fig. 5). Blue deploys surveillance assets over a ten separate AOIs representing choke points through the 400 km² mountainous region. Their mission is to find the fleeing terrorists and maintain a track on them until attack forces can be brought in to neutralize them. In the Base case, the Blue forces rely on Special Operations Force (SOF) teams, Chinook (CH 47) helicopters, Global Hawks, Tactical UAVs (TUAVs), and Unattended Ground Sensors (UGSs) to detect terrorist movement through the mountains.




Fig. 5 OPSIT 3—Detecting and tracking terrorist forces fleeing through mountainous region.

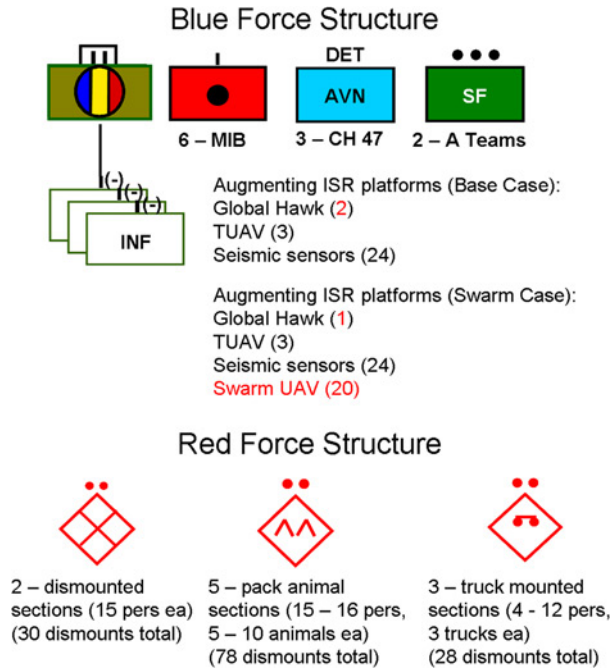


Fig. 6 OPSIT 3 force structure.

2. *Experimental Setup*

The force structure is shown in Fig. 6. Note that the Base case includes two Global Hawks while the Swarming replaces one of those Global Hawks with 20 swarming UAVs to keep the total sensor footprint roughly the same. The swarming UAVs were equipped with EO/IR sensors and were cued by the Global Hawk and the TUAVs. Additional experiments were run with varying numbers of swarming UAVs and UGSs to determine the impact of those factors on the performance of the system. Finally, in one experiment, the Rules of Engagement (ROEs) were altered to allow the swarming UAVs to provide targeting information to bring precision Non-Line of Sight (NLOS) effects directly to the target. In all other cases, targeting information was delayed until a CH-47 or SOF team could maneuver into position to confirm target identity and location.

3. *Key Results*

In the Swarming case, 85% of the insurgents were captured or killed while only 59% were found in the Base case (Table 5). 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%). 96% of the insurgents were captured or killed when the swarming assets provided target information to bring effects directly to the target.

Table 5 OPSIT 3 results

	Base	Swarm	Targeting
Number swarming UAVs	0	20	8
Number UGSs	24	8	0
Total red loss	59%	85%	76%
Increase in red losses over base		44%	28%

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 requiring those areas to be hard boundaries, the UAVs were able to find Red forces outside the AOIs.
- 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 than 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.

D. OASD Study Conclusions

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.

V. Demonstration

In October 2004 we demonstrated the use of these swarming algorithms with Johns Hopkins University APL, and the Army Research Laboratory to control a heterogeneous population of air and ground unmanned vehicles in an urban combat scenario at Aberdeen Proving Grounds. The demonstration included two UAVs controlled by digital pheromones and four ground robots controlled by a decentralized potential field algorithm¹⁵ similar to digital pheromones. The demonstration showed how these decentralized, stigmergic swarming algorithms can control and coordinate the behaviors of a heterogeneous mix of vehicles in a dynamic environment executing a diverse set of military operations.

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).

The air vehicles were modified Mig 117 Bravo target drones with a 6 ft wingspan (Fig. 7). 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 (Fig. 8). The area was initially marked by the operator for surveillance by the unmanned ground vehicles (UGVs) to verify the area in and around the buildings was clear. Once the area had been swept the UGVs autonomously organized a perimeter patrol pattern to protect the power plant from intrusion.



Fig. 7 Mig 117 Bravo modified target drones were equipped with modern engine and Micro-Pilot autopilot.

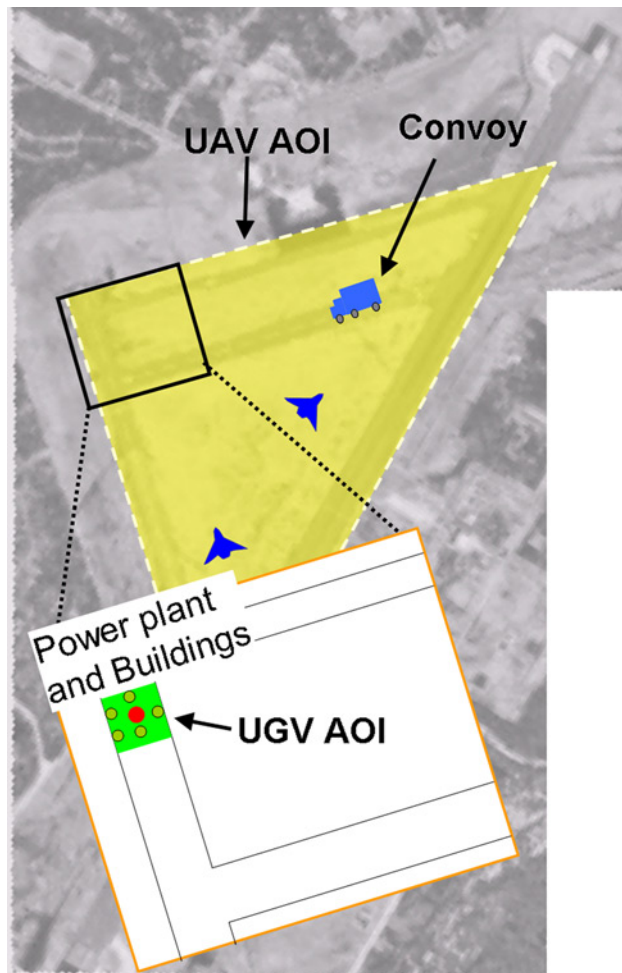


Fig. 8 Layout of the UAV and UGV surveillance areas at Aberdeen Proving Grounds.

Digital pheromones 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 Fig. 8). The operator selected the area for surveillance (the triangular area in Fig. 8). Using the surveillance algorithm described above, the two air units worked together to ensure even, thorough, and continuous coverage of all areas in the surveillance region.

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 towards the mock power plant. Before the convoy reaches the power plant, the UAVs 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 place where the potential threat is located that attracts the UGV with the target identification sensor. The UGV 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 operator requests the UAVs and UGVs to 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.

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 equipped with the 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 in a shared environment.

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 identified the function to be performed by the swarm. The vehicles autonomously configured themselves, determined which tasks to perform, and then cooperated in accomplishing 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.

VI. 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 all the 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 very 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 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 OASD 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

- ¹Parunak, H. V. D., "Go to the Ant": Engineering Principles from Natural Agent Systems," *Annals of Operations Research*, Vol. 75, 1997, pp. 69-101, <http://www.newvectors.net/staff/parunakv/gotoant.pdf>
- ²Bonabeau, E., Dorigo, M., and Theraulaz, G., *Swarm Intelligence: From Natural to Artificial Systems*. Oxford University Press, New York, 1999, <http://www.santafe.edu/sfi/publications/Bookinforev/icnew.html>.
- ³Camazine, S., Deneubourg, J.-L., Franks, N. R., Sneyd, J., Theraulaz, G., and Bonabeau, E., *Self-Organization in Biological Systems*. Princeton University Press, Princeton, NJ, 2001.
- ⁴Edwards, S. J. A., "Swarming on the Battlefield: Past, Present, and Future," RAND, MR-1100-OSD, Santa Monica, CA.
- ⁵Arquilla, J. and Ronfeldt, D., "Swarming and the Future of Conflict," RAND, DB-311, Santa Monica, CA, <http://www.rand.org/publications/DB/DB311>.
- ⁶Inbody, D., "Swarming: Historical Observations and Conclusions," *Proceedings of Swarming: Network Enabled C4ISR*, ASD C3I, Tysons Corner, VA, 2003.
- ⁷Lauren, M. K., "Describing Rates of Interaction between Multiple Autonomous Entities: An Example Using Combat Modelling," URL <http://aps.arxiv.org/pdf/nlin/0109024> [retrieved 26 Sept, 2001].
- ⁸Parunak, H. V. D., "Making Swarming Happen," *Proceedings of Swarming and Network-Enabled C4ISR*, ASD C3I, Tysons Corner, VA, 2003, <http://www.newvectors.net/staff/parunakv/MSH03.pdf>
- ⁹Richards, A. G. and How, J., "A Decentralized Algorithm for Robust Constrained Model Predictive Control," *Proceedings of American Control Conference*, Boston, MA, 2004, pp. 4261-4266.
- ¹⁰Reynolds, C. W., "Flocks, Herds, and Schools: A Distributed Behavioral Model," *Computer Graphics*, Vol. 21, No. 4 (July), 1987, pp. 25-34.
- ¹¹Vicsek, T., Cziroók, A., Ben-Jacob, E., Cohen, O., and Shochet, I., "Novel Type of Phase Transition in a System of Self-Driven Particles," *Physical Review Letters*, Vol. 75, No. 6, 1995, pp. 1226-1229.
- ¹²Rimon, E. and Kodischek, D. E., "Exact Robot Navigation Using Artificial Potential Functions," *IEEE Transactions on Robotics and Automation*, Vol. 8, No. 5 (October), 1992, pp. 501-518.
- ¹³Khatib, O., "Real-Time Obstacle Avoidance for Manipulators and Mobile Robots," *International Journal of Robotics Research*, Vol. 5, No. 1, 1986, pp. 90-98.
- ¹⁴Masoud, S. A. and Masoud, A.A., "Constrained Motion Control Using Vector Potential Fields," *IEEE Trans. on Systems, Man, and Cybernetics*, Vol. 30, No. 3, 2000, pp. 251-272.
- ¹⁵Chalmers, R., Scheidt, D., Neighoff, T., Witwicki, S., and Bamberger, R., "Cooperating Unmanned Vehicles," *Proceedings of AIAA 1st Intelligent Systems Technical Conference*, Chicago, IL, 2004.
- ¹⁶Eun, Y. and Bang, H., "Cooperative Control of Multiple Unmanned Aerial Vehicles Using the Potential Field Theory," *Journal of Aircraft*, Vol. 43, No. 6, 2006, pp. 1805-1814.
- ¹⁷Sauter, J. A., Matthews, R., Parunak, H. V. D., and Brueckner, S., "Evolving Adaptive Pheromone Path Planning Mechanisms," *Proceedings of Autonomous Agents and Multi-Agent Systems (AAMAS02)*, ACM, Bologna, Italy, 2002, pp. 434-440, www.newvectors.net/staff/parunakv/AAMAS02Evolution.pdf
- ¹⁸Sauter, J. A., Matthews, R., Parunak, H. V. D., and Brueckner, S. A., "Performance of Digital Pheromones for Swarming Vehicle Control," *Proceedings of Fourth International Joint Conference on Autonomous Agents and Multi-Agent Systems*, ACM, Utrecht, Netherlands, 2005, pp. 903-910, <http://www.newvectors.net/staff/parunakv/AAMAS05SwarmingDemo.pdf>
- ¹⁹Gaudio, P., Shargel, B., Bonabeau, E., and Clough, B.T., "Swarm Intelligence: A New C2 Paradigm with an Application to Control of Swarms of UAVs," *Proceedings of 8th ICCRTS Command and Control Research and Technology Symposium*, Washington, DC, 2003.
- ²⁰Walter, B., Sannier, A., Reiners, D., and Oliver, J., "UAV Swarm Control: Calculating Digital Pheromone Fields with the Gpu," *Journal of Defense Modeling and Simulation: Applications, Methodology, Technology*, Vol. 3, No. 3, 2006, pp. 167-176.
- ²¹Payton, D., Daily, M., Estowski, R., Howard, M., and Lee, C., "Pheromone Robotics," *Journal Autonomous Robots*, Vol. 11, No. 3, 2001, pp. 319-324.
- ²²Panait, L. and Luke, S., "Pheromone-Based Utility Model for Collaborative Foraging," *Proceedings of Third International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS'04)*, New York, NY, 2004, pp. 36-43.

²³Dasgupta, P., “Distributed Automatic Target Recognition Using Multiagent UAV Swarms,” *Proceedings of Fifth International Joint Conference on Autonomous Agents and Multiagent Systems*, Hakodate, Japan, 2006, pp. 479–481.

²⁴Grassé, 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.

²⁵Altshuler, Y., Wagner, I. A., and Bruckstein, A. M., “On Swarm Optimality in Dynamic and Symmetric Environments,” *Proceedings of Second International Conference on Informatics in Control, Automation and Robotics (ICINCO-MARS-05)*, Barcelona, Spain, 2005.

²⁶Brueckner, S., Dr.rer.nat., Humboldt University Berlin, 2000, <http://dochostrz.hu-berlin.de/dissertationen/brueckner-sven-2000-06-21/PDF/Brueckner.pdf>.

²⁷Parunak, H. V. D., Purcell, M., and O’Connell, R., “Digital Pheromones for Autonomous Coordination of Swarming UAV’s,” AIAA, Paper AIAA-2002-3446, May 2002.

²⁸SMDC-BL-AS, “Swarming Unmanned Aerial Vehicle (UAV) Limited Objective Experiment (Loe),” U.S. Army Space and Missile Defense Battlelab, Studies and Analysis Division, Huntsville, AL, Nov 2001.

²⁹Dubik, J., Richards, R., and Trinkle, G., “Joint Concept Development and Experimentation,” *Proceedings of Swarming: Network Enabled C4ISR*, ASD C3I, Tysons Corner, VA, 2003.

³⁰Parunak, H. V. D., Brueckner, S., and Odell, J. J., “Swarming Coordination of Multiple UAV’s for Collaborative Sensing,” AIAA, Paper AIAA-2003-6525, September 2003.

³¹Parunak, H. V. D. and Brueckner, S. A., “Stigmergic Learning for Self-Organizing Mobile Ad-Hoc Networks (Manet’s),” *Proceedings of Third International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS’04)*, ACM, Columbia University, NY, 2004, pp. 1324–1325, <http://www.newvectors.net/staff/parunakv/AAMAS04MANET.pdf>.

³²Weyns, D., Schelfhout, K., Holvoet, T., and Lefever, T., “Decentralized Control of E’GV Transportation Systems,” *Proceedings of 4th International Joint Conference on Autonomous Agents and Multi-Agent Systems, Industry Track*, Utrecht, Netherlands, 2005, pp. 67–74.

Fernando Figueroa
Associate Editor