AIAA 2002-3446

DIGITAL PHEROMONES FOR AUTONOMOUS COORDINATION OF SWARMING UAV'S

Dr. H. Van Dyke Parunak
Altarum
3520 Green Ct, Suite 300
Ann Arbor, MI 48105-1579
van.parunak@altarum.org

LCDR Michael Purcell
Fleet Composite Squadron SIX
1224 Pocahontas Street
Naval Air Station
Norfolk VA 23511-2414
purcellm@vc6.navy.mil

Mr. Robert O'Connell
Space and Missile Defense Command
106 Wynn Drive
Huntsville, AL 35807-3801
robert.oconnell@smdc.army.mil

Modern UAV's reduce the threat to human operators, but do not decrease the manpower requirements. Each aircraft requires a flight crew of one to three, so deploying large numbers of UAV's requires committing and coordinating many human warfighters. Insects perform impressive feats of coordination without direct inter-agent coordination, by sensing and depositing pheromones (chemical scent markers) in the environment [14]. We have developed a novel technology for coordinating the movements of multiple UAV's based on a computational analog of pheromone dynamics. The control logic is simple enough that it can be executed autonomously by a UAV, enabling a single human to monitor an entire swarm of UAV's. This paper describes the technology, its application to UAV coordination, and the results we have obtained.

INTRODUCTION

The current generation of UAV's reduces the threat to human operators, but does not decrease the manpower requirements. Each aircraft requires a flight crew of one to three people, so deploying large numbers of UAV's requires committing and coordinating many human warfighters. In addition to the cost of human operators,` this approach encounters unresolved challenges as to just how coordination can be achieved.

Insects perform impressive feats of coordination without direct inter-agent coordination, by sensing and depositing pheromones (chemical scent markers) in the environment [14]. We have developed a novel technology for coordinating the movements of multiple UAV's based on a computational analog of pheromone dynamics. The control logic is simple enough that it can be executed autonomously by a UAV, enabling a single human to monitor an entire swarm of UAV's. This paper describes the technology, its application to UAV coordination, and the results we have obtained.

The remainder of the paper has four sections.

1. We review four requirements for the command and control of swarming UAV's.
2. We describe digital pheromones [4], a refinement of robotic methods for guidance by potential fields [20].
3. We discuss how digital pheromones can be deployed in the real world, building on current and emerging architectures for Unattended Ground Sensors (UGS's).
4. Finally, we report results from two series of experiments with our mechanisms.

D^4 REQUIREMENTS FOR SWARMING UAV COMMAND AND CONTROL

Effective command and control 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 diverse in several ways. It must integrate diverse *functions*, including communications among the UAV's, command oversight, and information management to enable the UAV's to make reasonable decisions. It must handle information of diverse *kinds*, including imagery, vibration, chem/bio, ELINT, etc. This information may concern diverse *entities*, including a heterogeneous population of unmanned vehicles (air and ground), targets to be approached, threats to be avoided, and the presence of other friendly units with which coordination is required. It may also come from diverse *sources*, including local ground sensors, information from other nearby friendly units, and far-distant intelligence (e.g., national assets).

Distributed.—The US military is facing a serious shortfall in long-range communications bandwidth [1]. Warfighters cannot assume the availability of unlimited satcom channels. This limitation can be addressed by distributing the C^2 system physically over the battlespace. Distribution helps lower bandwidth in two ways.

1. The components of a distributed system are fairly close to one another, and can communicate with nearby neighbors using relatively low power, thus permitting bandwidth to be reused beyond a local horizon. Long-range communications can be handled by propagating messages through the network.
2. Much information about the battlespace has a strong geographical component, and is needed most by forces close to where the information is generated. A distributed system can store information close to where it is generated and close to where it is needed, rather than in some huge central repository. This strategy greatly reduces the need to move large bodies of information over long distances.

Decentralized.—We want UAV's to be able to make local decisions, within the scope of responsibility originally assigned to them, without detailed explicit commands from a central point, for several reasons.

1. The current (centralized) model of UAV control requires a team of two or three humans for each UAV. The manpower costs of this model make it impossible to field large swarms of UAV's.
2. The time delay for a central command to process data from UAV's and generate new commands is unacceptable in rapidly changing combat situations.
3. Centralized control generates "choke points" that impede system operation, in two ways. First, a single central control can be overloaded by data from many subordinate UAV's. Second, a central control facility presents the adversary with a single point of vulnerability that can disable the entire system with a single attack.

Dynamic.—The battlespace is an uncertain and rapidly changing environment. Red forces will try to change unexpectedly. Imperfect knowledge about Blue may lead to changed assessments. The nonlinear nature of warfare [12] may itself generate unexpected changes in the situation. The C^2 mechanisms must be able to change and adapt in such an environment.

POTENTIAL FIELDS AND DIGITAL PHEROMONES

Digital pheromones are an extension of the use of potential fields to guide robotic entities.

Potential-based movement systems are inspired by electrostatics. The (vector) electric field $\vec{E}(\vec{r})$ at a point in space is defined as the force felt by a unit charge at that point. We define a (scalar) potential field $\varphi_{21} = -\int_{P_1}^{P_2} \vec{E} \bullet d\vec{r}$ by integrating this vector field from an arbitrary reference point to each point in the space. Conversely, the field may be expressed as the gradient of the potential, $\vec{E} = -\nabla\varphi$, and a massless charged particle will move through space along this gradient.

In electrostatics, the field is generated by the physical distribution of charges according to Coulomb's law. Einstein's extension of the formalism to gravity leads to a gravitational field generated by the physical distribution of mass. Thus the movement of a massive charged particle will follow a composition of two fields.

The notion of movement guided by a potential gradient has been applied to other situations in which the field is generated, not by natural physical phenomena, but by synthetic constructs. A parade example is robot navigation [20], which automatically maps from a given distribution of targets and obstacles to a movement plan. In such applications, the designer of the field is not limited to two components of the field (electrostatic and gravitational), but can include many different fields to represent different classes of targets and obstacles.

We are interested in using a potential field to guide UAV's through the battlespace (Figure 1). In this scenario, the UAV's seek to destroy the tank farm, which is defended by two missile batteries. The vehicles climb a potential gradient centered on the tank farm while avoiding gradients centered on the threats.

How does a potential field satisfy the D^4 requirements?

Diverse.—Yes. One of the strengths of traditional potential field methods is their ability to incorporate multiple information types (for example, repulsive threats and attractive targets as in Figure 1).

Distributed.—No. Traditional potential field methods require the data underlying the field to be collected to a single point for centralized processing.

Decentralized.—Yes. Assets moving under the

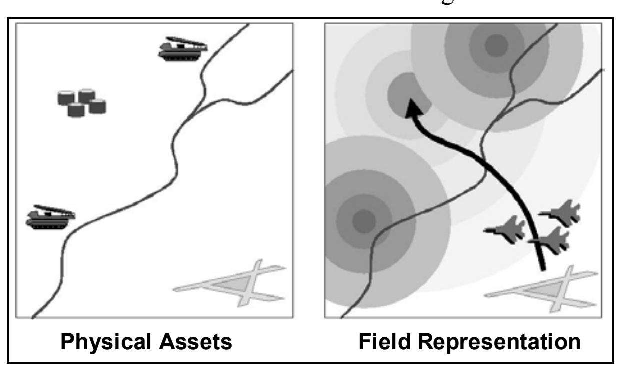

Figure 1: Potential Fields Corresponding to Physical Assets

influence of a potential field make local decisions, based on the field strength in their immediate vicinity.

Dynamic.—No. Conventional potential fields are computed in their entirety before the robot begins moving.

Digital pheromones provide a way to construct a potential field dynamically and in a distributed way, thus satisfying all of the D^4 requirements.

Insects perform impressive feats of coordination without direct inter-agent coordination, by sensing and depositing pheromones (chemical scent markers) in the environment [14]. For example, ants construct networks of paths that connect their nests with available food sources. Mathematically, these networks form minimum spanning trees [10], minimizing the energy ants expend in bringing food into the nest. Graph theory offers algorithms for computing minimum spanning trees, but ants do not use conventional algorithms. Instead, this globally optimal structure emerges as individual ants wander, preferentially following food pheromones and dropping nest pheromones if they are not holding food, and following nest pheromones while dropping food pheromones if they are holding food.

The real world provides three operations on chemical pheromones that support purposive insect actions.

- It *aggregates* deposits from individual agents, fusing information across multiple agents and through time.
- It *evaporates* pheromones over time. This dynamic is an innovative alternative to traditional truth maintenance. Traditionally, knowledge bases remember everything they are told unless they have a reason to forget something, and expend large amounts of computation in the NP-complete problem of reviewing their holdings to detect inconsistencies that result from changes in the domain being modeled. Ants immediately begin to forget everything they learn, unless it is continually reinforced. Thus inconsistencies automatically remove themselves within a known period.
- It *diffuses* pheromones to nearby places, disseminating information for access by nearby agents.

The pheromone field constructed by the ants in the environment is in fact a potential field that guides their movements. Unlike many potential fields used in conventional robotics applications, it satisfies all of the D^4 requirements.

Diverse.—Ants can respond to combinations of pheromones, thus integrating the effect of multiple inputs.

Distributed.—The potential field is generated by pheromone deposits that are stored throughout the environment. These deposits do their work close to where they are generated, and are used primarily by ants that are near them.

Decentralized.—Both ant behavior and pheromone field maintenance are decentralized. Ants interact only with the pheromones in their immediate vicinity, by making deposits and reading the local strength of the pheromone field. Because diffusion falls off rapidly with distance, deposits contribute to the field only in their immediate vicinity.

Dynamic.—Under continuous reinforcement, the pheromone field strength stabilizes rapidly, as a concave function of time (proportional to $\int_0^t E^\tau d\tau$ where $E \in (0,1)$ is the evaporation rate) [4]. Thus new information is quickly integrated into the field, while obsolete information is automatically forgotten, through pheromone evaporation.

ADAPTIV's synthetic pheromones are digital data structures, not chemicals. These data structures live in a network of *place agents,* which represent regions of the battlespace. All place agents can run on a single computer for simulation purposes, but in actual deployment each place agent might run on an enhanced unattended ground sensor (UGS) placed in the battlespace by air drops or artillery and responsible for any location to which it is closer than any other UGS. ADAPTIV refers to such an enhanced UGS as a HOST (Hostility Observation and Sensing Terminal). Each place agent is a neighbor to a limited set of other place agents, those that are responsible for adjacent regions of space, and it exchanges local information with them. In addition to place agents, ADAPTIV includes *walker agents,* analogous to the ants, representing Blue resources such as UAV's. Walker agents move through the battlespace by interacting with the place agent for each region that they visit. Place agents and walker agents are software entities, while HOST's and UAV's are the hardware in which they run.

Each place agent maintains a scalar variable corresponding to each pheromone flavor. It augments this variable when it receives additional pheromones of the same flavor (whether by deposit from a walker agent, from its own sensors, or by propagation from a neighboring place agent). It also evaporates the variable over time, and propagates pheromones of the same flavor to neighboring place agents based on the current strength of the pheromone. The underlying mathematics of the field developed by such a network of places, including critical stability theorems, rest on two fundamental equations [4]. The parameters in both are:

- $P = \{p_i\}$ = set of place agents
- $N: P \to P$ = neighbor relation between place agents. Thus the place agents form an asymmetric multigraph.
- $s(t,p)$ = pheromone strength at time t and place agent p
- $r(t,p)$ = external input at time t to place agent p
- $q(t,p)$ = propagated input at time t to place agent p
- $E \in (0,1)$ = evaporation parameter
- $F \in [0,1)$ = propagation parameter

The first equation describes the evolution of the strength of a single pheromone flavor at a given place agent:

$$s(t+1, p) = E * s(t, p) + r(t, p) + q(t, p)$$

The first term of this equation reflects evaporation of pheromone, the second reflects external deposit of new pheromone, and the third reflects propagation from neighboring place agents. The second fundamental equation describes that propagation:

$$q(t+1, p) = \sum_{p \in N(p')} \frac{F}{|N(p')|} \left(r(t, p') + q(t, p') \right)$$

This equation states that each place agent that has p as a neighbor propagates a proportion of its newly-acquired pheromone to p in each time period, the proportion depending on the parameter F and the total number of neighbors.

Using these equations, one can demonstrate several critical stability and convergence theorems [4], including:

- *Local Stability*: The strength of the output propagated from any set of place agents to their neighbors at $t + 1$ is strictly less than the strength of the aggregate input (external plus propagated) to those place agents at t.
- *Propagated Stability*: There exists a fixed upper limit to the aggregated sum of all propagated inputs at an arbitrary place agent if a one-time and one-place external input is assumed.
- *Global Stability*: The pheromone strength in any place agent is bounded.

Our implementations include another parameter, the pheromone *threshold S*. If the strength of the pheromone at a place agent drops below this threshold, the software no longer processes that pheromone, and it disappears from the system.

In principle, there are no restrictions on the graph of place agents. In physical movement problems, each place agent is responsible for a region of physical space, and the graph of place agents represents adjacency among these regions. There are different ways in which place agents can be assigned to space. In this LOE, we tile the physical space with hexagons, each representing a place agent with six neighbors.

A walker agent inhabits one place agent at any given time. It can read the current strength of pheromones at that place agent as a function of their flavors, and deposit pheromones into the place agent. It can also determine from the place agent the relative strength of a given flavor at the place agent and at each of its neighbors. A walker agent combines the strengths of the different pheromones that it senses using a control equation. Then it moves from one place agent to another by spinning a roulette wheel whose segments are weighted according to the output of this control equation.

The potential of insect models for multi-agent coordination and control in practical applications is receiving increasing attention [2, 14]. Both of these reviews cover other mechanisms in addition to pheromones. Important theoretical discussions with simple applications are described in [6, 8], and [7] shows how these techniques can play a credible game of chess. The most elaborate model for synthetic pheromones is [25], which pursues the analogy to the point of constructing a synthetic chemistry through which agents interact with multi-typed pheromone traces.

The most mature practical use of pheromone techniques is in routing telecommunications packets (e.g., [3, 11]). Application of these techniques to moving physical entities can be traced to Altarum's Cascade system [17, 18], a self-routing modular material handling system for manufacturing. That work appealed to a neural backpropagation model rather than a pheromone model for its antecedents. However, the accumulation of weights on frequently-activated links through backpropagation has many formal similarities to the accumulation of pheromones on well-traveled paths. The application of these techniques to routing and load balancing is being extended in the ESPRIT MASCADA project [19], and is developed most fully in a Ph.D. dissertation by one of Altarum's scientists [4]. Steels proposed similar mechanisms for coordinating small robots used in exploring remote planets [23]. Dorigo and colleagues [5, 9, 13] have applied these mechanisms to a range of optimization problems including the traveling salesperson problem and the quadratic assignment problem. The ADAPTIV system itself, for military C^2, is described in detail in [15].

We have developed a number of techniques to facilitate the application of synthetic pheromones in D^4 domains, including multiple pheromones with differing semantics and dynamics, giving agents a sense of history, using multiple software "ghost agents" to develop paths for physical resources, visualizing the run-time behavior of the agents, and evolving agent behavior as the agents run. These techniques are discussed in more detail in [16, 21].

IMPLEMENTING THE VISION

This section discusses how synthetic pheromones would be deployed in combat, and three different simulation environments: the ADAPTIV software platform, the adaptation to EADTB used in the LOE outlined below, and a recommended configuration to provide greater flexibility for future experimentation.

Deployment.—The most natural implementation of ADAPTIV as a combat system builds on a network of enhanced unattended ground sensors (HOST,s) physically distributed throughout the battlespace. Each HOST includes a small computer that stores synthetic pheromones. and provides processing for the place agent located at that HOST. Walker agents can run either on HOST's or on computers on board UAV's. (Walker agents can also represent other forms of robotic vehicles, or manned assets, or even special operations forces, all of which could take advantage of ADAPTIV's D^4 C^2 framework.) The population of HOST's can be deployed by air drops or artillery, flown in as UAV's, or placed by SOF. To reduce their vulnerability to enemy attack, they may be camouflaged to blend in with the surrounding terrain (e.g., rocks in desert or mountain terrain, plants in jungle terrain). Alternatively, they may be packaged to resemble anti-personnel mines or bomblets from a cluster munition, protecting them even in full view. Their local communications protocols can reconfigure the network even if some of the HOST's are removed or become inoperative. Alternatively, the place agents can be grouped together on computers at various locations (e.g., regional AWACS and ground HQ's), with a corresponding reduction in the degree of distribution of the system, but still providing the other three D's.

ADAPTIV Platform.—The ADAPTIV software system developed in the DARPA JFACC program is based on the SWARM agent simulation platform [24]. This platform is especially designed to support swarming agents, including the ability to define different subpopulations of agents, flexible interaction mechanisms, data logging and real-time querying interfaces, efficient execution for large numbers of concurrent agents, and a variety of scheduling mechanisms. As discussed elsewhere [16, 21], ADAPTIV supports a wide range of pheromone mechanisms, including pheromone aggregation, propagation, and evaporation, multiple pheromones with differing semantics and dynamics, history-sensitive agent behavior, multiple software "ghost agents," visualization of agent behavior, and evolution of agent behavior as the agents run.

"As-Is" LOE Implementation.—For the purposes of the LOE described later in this paper, a subset of the ADAPTIV algorithms has been implemented in EADTB. This subset does not include such features as ghost agents, pheromone visualization, or agent evolution. Its control equation for combining pheromone flavors is much simpler than that used in ADAPTIV. It uses only two pheromones, an attractive pheromone leading to targets and a repulsive pheromone deposited by other UAV's, and simply computes the quotient attractive/repulsive. The computational overhead of the full suite of ADAPTIV mechanisms is considerable, and there is a limit to how many of them can be embedded directly into EADTB without adversely affecting its performance.

"To-Be" Experimental Implementation.—The results of experimentation so far show impressive benefits for swarming UAV's, even under the constraints imposed by embedding the synthetic pheromone logic in EADTB. A much more powerful approach would use ADAPTIV alongside EADTB or a similar combat simulator. At a minimum, offline use of the full ADAPTIV system would enable more effective tuning of the current subset of ADAPTIV algorithms embedded in EADTB. In our experience, different scenarios require different agent configurations (e.g., different structure or weights in the control equation) for optimal performance, and ADAPTIV is a much more flexible environment for testing different configurations than is EADTB. Ideally, ADAPTIV and EADTB could communicate in simulated real time via HLA (Figure 2). ADAPTIV would model the performance of the pheromone system (place agents and walker agents), while EADTB would model the

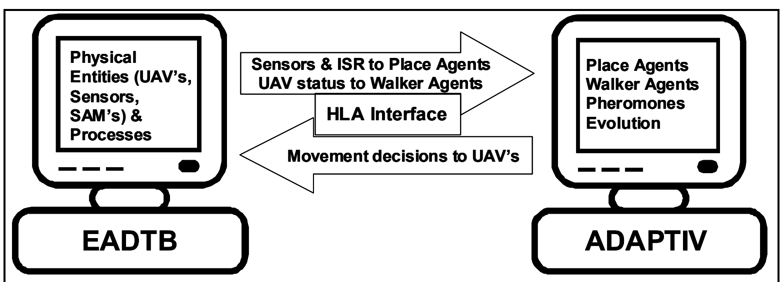

Figure 2: Recommended "To-Be" Experimental Implementation

dynamics of the physical world. Messages from EADTB to ADAPTIV would provide sensory input to the place agents and report location changes of the UAV's assigned to walker agents, while ADAPTIV would communicate to EADTB course changes for UAV's based on the pheromone-driven logic of walker agents. This configuration provides the best of both systems: high-fidelity detailed modeling of the physics of combat in EADTB, and efficient execution of the full spectrum of synthetic pheromone techniques in ADAPTIV.

EXPERIMENTAL RESULTS

Two series of experiments demonstrate the value of the ADAPTIV approach. The first series was conducted under the DARPA JFACC program using the ADAPTIV testbed. The second was conducted under the auspices of the J9/Joint Experimentation division of JFCOM, using EADTB.

ADAPTIV Experiments

The ADAPTIV research program included numerous experiments that addressed various aspects of digital pheromones [16]. The experimental environment is a uniform grid of place agents, each with six nearest neighbors, yielding a hexagonal grid of regions. Each region is nominally 50 km in diameter. Regions can be occupied by a variety of entities, including

- ARBTN: armored battalion
- INBTN: infantry battalion
- ADA: air defense (artillery)
- BRDG: bridge
- ROADX: road crossing
- FA: fueling area
- REGHQ: regional headquarters
- ADHQ: air defense headquarters
- CORHQ: corps headquaters
- F2E: good performance fighter jet
- F4: better performance fighter jet
- F5: best performance fighter jet
- B100: big bomber
- B101: medium-big bomber
- B102: medium bomber
- OIL: oil field
- BLDG-S: building/structure
- BUNKER: bunker
- H-1: helicopter
- SAM-FL: surface-to-air missile: fixed launcher
- SAM-ML: surface-to-air missile: mobile launcher
- TOWER1: air field tower
- TRKVEH: track vehicle
- TRUCKA: truck
- ARMOR: armored vehicle

Two sets of experiments in this environment are particularly relevant for UAV C^2: path formation and response to dynamic change. In both sets, Blue was exclusively an air force (entities F2E, F4, F5, B100, B101, B102, and H-1), while Red was exclusively a ground force.

Path planning.—The path planning experiments studied the ability of digital pheromones to plan paths over different configurations of threats (e.g., ADA) and targets (typically CORHQ), as illustrated in Figure 3. In each case, the display shows the status of the path while mission package is en route to the target, at the location indicated by the aircraft icon, and if a disruption were to occur (e.g., a pop-up SAM), the path would be replanned dynamically. In Figure 3a, the path forms to the nearer of two targets of equal priority. Figure 3b shows the path when the western target is given a higher priority than the northern one. Figure 3c shows that the digital pheromones can find an entrance into a heavily guarded area, while Figure 3d shows that when no entrance exists, the algorithm finds a shortest path.

The configuration in Figure 3c has also been submitted

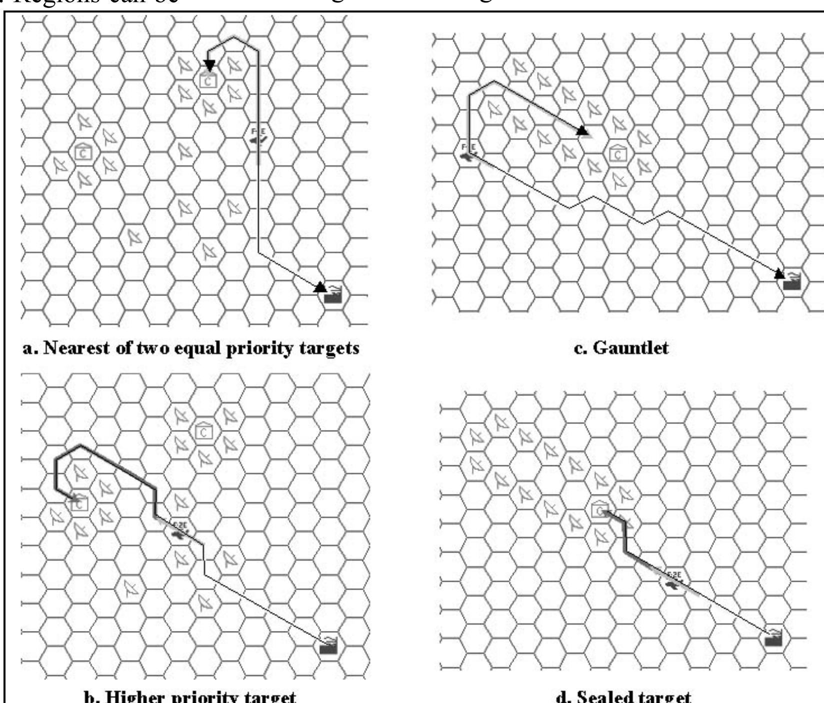

Figure 3: Path planning.—Digital pheromones are able to plan the appropriate path in real time, as the mission unfolds, for varying configuration.

to a conventional centralized optimization procedure, which is unable to find a safe path to the target unless the experimenter establishes an interim waypoint at the entrance to the gauntlet. ADAPTIV is able to solve this configuration because it computes the path as the mission package flies, thus gathering information from several different perspectives.

Dynamic change.—To assess the capability of digital pheromones to address a changing environment, we use a more complex scenario, in which 87 Blue units fly 181 missions against 116 Red units.. First, we make all Red threats visible and stationary, and let the ghosts plan paths to the target for each of 181 offensive Blue missions against an entrenched Red force. We compute these paths using two different propagation parameters for Red threat pheromones, one that permits paths to fly relatively close to the threats, and another that keeps paths relatively far from the threats. Then we turn on Red movement and hiding behaviors, and compare the outcome of two sets of runs. In one set, Blue does not use ghosts or pheromones at all, but simply flies each mission on its precomputed path. This mode of operation corresponds to traditional pre-planned flight itineraries, except that our pre-planned paths, based on complete knowledge of Red's locations at the time of planning, are superior to those that could be constructed in any real conflict. In the other set of runs, Blue ignores precomputed paths and relies on ghosts to form paths for its missions dynamically. We assess the outcome of each run by the total remaining strength of Blue and Red assets at the end of the set of missions.

Figure 4 shows the medians of Red and Blue total unit strengths remaining at the end of the mission. The points plotted are medians of five runs. We report three configurations. In "pathscript," each mission flies the path precomputed for it using a high Red propagation parameter, leaving a conservative margin around Red threats. In "pathscriptnarrow," Blue again flies precomputed paths, this time using paths computed with a lower Red propagation parameter, and permitting Blue to come closer to Red locations. These less conservative paths lead to increased combat between Blue aircraft and Red threats, and both Red and Blue losses increase compared with "pathscript." In "pathghost," Blue missions ignore precomputed paths and send out ghosts to compute their paths dynamically as the mission unfolds. In this mode of operation, Blue's losses are least, since it can now avoid pop-up Red threats. As a result, it can deliver more weaponry to its assigned targets, increasing Red's losses in comparison with the other two scenarios.

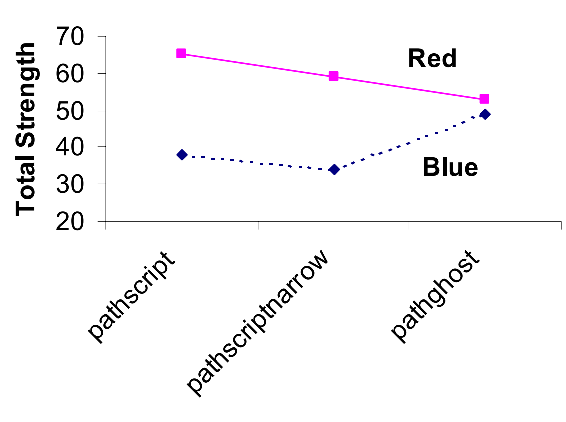

Figure 4: Performance in Dynamic Environment

J9 LOE

The U.S. Army Space and Missile Defense Battlelab, in support of the Joint Forces Command, used a subset of the ADAPTIV algorithms in a limited-objective experiment to determine the effects of employing affordable Swarming UAV's against an enemy's mobile strategic Surface to Air Missiles (SAM's; SA20's) utilizing an anti-access strategy [22]. The study considered four cases. The base case was drawn from a previous JFCOM study, Unified Vision 00, which utilized Global Hawks as UAV's. The comparison cases envisioned a swarm of smaller UAV's, with flight characteristics typical of a LOCASS-type platform. The study cases were 1) UAV's with sensors only, 2) UAV's with both sensors & munitions, and 3) UAV's with sensors/ munitions/jammers. The munitions on armed UAV's were deployed by flying the UAV into the target, thus sacrificing the UAV. The matrix also included excursions for each of the study cases that varied the quantities (10, 50,100) of UAV's in the swarm. The base case and study case excursions resulted in a total of 10 excursions with 10 runs each. The results were then analyzed for statistically supported comparisons across several measures of effectiveness (MOE's). The results for UAV's with sensors, weapons, and jammers were the same as those for UAV's with only sensors and weapons.

- Percent of Red assets detected: all swarming cases significantly outperformed the base case, and larger swarms significantly outperformed smaller ones. Sensor-only cases were slightly better than cases with multi-function UAV's, presumably because the population of armed UAV's decreases over the run as some UAV's function as weapons. The greatest difference was 30% detection (base case) vs. 95% detection (100 sensor-only UAV's).
- Percent reduction of successful TBM launches: no significant difference from base case.

- Percent of Red assets destroyed (by type): the smallest swarm of sensor-only UAV's did not significantly outperform the base case, but all other swarms did. Larger swarms significantly outperformed smaller ones, and swarms in which UAV's were armed outperformed those in which UAV's carried only sensors. The greatest difference for each category of Red asset is between the base case and a swarm of 100 armed UAV's. The differences are 25% destroyed vs. 63% for TBM TEL's, 5% vs 56% for tombstone radars, and 6% vs. 68% for SAM TEL's.
- Percent of Blue assets destroyed: no significant difference from base case. (UAV's are considered expendable, and not counted among blue assets for the purpose of this statistic.) When measured as a percentage of total missions flown, this metric drops slightly for larger swarms and for armed UAV's compared with unarmed ones, but the differences are still within the margin of error of the experiment.
- System Exchange Ratio (SER): the base case had a SER of 0.51, indicating that Blue lost twice as many assets as Red. All swarms except the 10-unit sensor-only swarm significantly outperformed the baseline. Larger swarms outperformed smaller ones, and armed UAV's outperformed unarmed ones. The best SER, for a swarm of 100 armed UAV's, was 4.56.

It has been observed that the improved performance in these scenarios is due to the increased number of sensors deployed in the battlespace, not to the use of a swarming algorithm per se. However, no competing algorithm can coordinate a hundred UAV's effectively. Current C^2 mechanisms require multiple human operators per UAV, and coordination across such a team poses formidable problems. The swarming approach is valuable precisely because it does not require a large cadre of operators.

CONCLUSIONS

Swarming techniques inspired by insect pheromones offer a powerful mechanism for coordinating unmanned vehicles such as UAV's. These mechanisms can support diverse functions and information sources. They distribute computation in a way that enables systems to scale well, and decentralize it, reducing vulnerability to attack or system overload, and they adapt well to dynamic changes in the environment. Experiments with these mechanisms show considerable promise, and encourage their exploration in more sophisticated war-gaming environments.

ACKNOWLEDGEMENTS

The swarming UAV's in the JFCOM UAV LOE coordinate their actions using techniques developed by Altarum, a not-for-profit research and development organization in Ann Arbor, MI, under DARPA's JFACC program (program managers COL Dan McCorry USAF and MAJ Sharon Heise USAF). The technical team included Sven Brueckner, Olga Gilmore, Robert Matthews, Murali Nandula, and John Sauter. These techniques are embodied in a software toolkit called ADAPTIV (Adaptive control of Distributed Agents through Pheromone Techniques and Interactive Visualization). Further information about these techniques and ADAPTIV is available from Dr. Van Parunak, Altarum's Chief Scientist, at van.parunak@altarum.org, or from the project website at www.erim.org/cec/projects/adaptiv/. The LOE described above was funded by J9 and DARPA.

REFERENCES

[1] D. Anhalt. The Changing Nature of Commercial Satcom and its Impact on US Military Advantage. *Satellite 2001*, Office of Net Assessment, Washington, DC, 2001.

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

[3] E. Bonabeau, F. Henaux, S. Guérin, D. Snyers, P. Kuntz, and G. Theraulaz. Routing in Telecommunications Networks with "Smart" Ant-Like Agents. Santa Fe Institute, Santa Fe, NM, 1998.

[4] S. Brueckner. *Return from the Ant: Synthetic Ecosystems for Manufacturing Control*. Dr.rer.nat. Thesis at Humboldt University Berlin, Department of Computer Science, 2000. URL http://dochost.rz.hu-berlin.de/dissertationen/brueckner-sven-2000-06-21/PDF/Brueckner.pdf.

[5] M. Dorigo, V. Maniezzo, and A. Colorni. The Ant System: Optimization by a Colony of Cooperating Agents. *IEEE Transactions on Systems, Man, and Cybernetics, Part B*, 26(1):1-13, 1996.

[6] A. Drogoul. *De la Simulation Multi-Agents à la Résolution Collective de Problèmes*. Thesis at University of Paris IV, 1993.

[7] A. Drogoul. When Ants Play Chess (Or Can Strategies Emerge from Tactical Behaviors? In *Proceedings of Fifth European Workshop on Modelling Autonomous Agents in a Multi-Agent World (MAAMAW '93)*, Neuchâtel, CH, pages 13-27, Springer, 1995.

[8] J. Ferber. *Multi-Agent Systems: An Introduction to Distributed Artificial Intelligence*. Harlow, UK, Addison Wesley Longman, 1999.

[9] L. M. Gambardella and M. Dorigo. HAS-SOP: Hybrid Ant System for the Sequential Ordering Problem. Istituto Dalle Molle di Studi sull'Intelligenza Artificiale (IDSIA), Lugano, Switzerland, 1997. URL ftp://ftp.idsia.ch/pub/luca/papers/SOP-tr-idsia-11-97.ps.gz.

[10] S. Goss, S. Aron, J. L. Deneubourg, and J. M. Pasteels. Self-organized Shortcuts in the Argentine Ant. *Naturwissenschaften*, 76:579-581, 1989.

[11] M. Heusse, S. Guérin, D. Snyers, and P. Kuntz. Adaptive Agent-Driven Routing and Load Balancing in Communication Networks. *Advances in Complex Systems*, 1:234-257, 1998.

[12] M. K. Lauren. Describing Rates of Interaction between Multiple Autonomous Entities: An Example Using Combat Modelling. 2001. PDF File, http://aps.arxiv.org/pdf/nlin/0109024.

[13] V. Maniezzo, A. Colorni, and M. Dorigo. The Ant System Applied to the Quadratic Assignment Problem. IRIDIA, Université Libre de Bruxelles, Bruxelles, Belgium, 1994. URL ftp://iridia.ulb.ac.be/pub/dorigo/tec.reps/TR.03-ANT-QAP.ps.gz.

[14] H. V. D. Parunak. 'Go to the Ant': Engineering Principles from Natural Agent Systems. *Annals of Operations Research*, 75:69-101, 1997. URL http://www.erim.org/~vparunak/gotoant.pdf.

[15] H. V. D. Parunak. Adaptive control of Distributed Agents through Pheromone Techniques and Interactive Visualization. 2000. Web page, www.erim.org/cec/projects/adaptiv/.

[16] H. V. D. Parunak, S. A. Brueckner, J. Sauter, and J. Posdamer. Mechanisms and Military Applications for Synthetic Pheromones. In *Proceedings of Workshop on Autonomy Oriented Computation*, Montreal, Canada, 2001. URL http://www.erim.org/~vparunak/aoc2001.pdf.

[17] H. V. D. Parunak, J. Kindrick, and B. Irish. Material Handling: A Conservative Domain for Neural Connectivity and Propagation. In *Proceedings of Sixth National Conference on Artificial Intelligence*, Seattle, WA, pages 307-311, American Association for Artificial Intelligence, 1987. URL http://www.erim.org/~vparunak/AAAI87.pdf.

[18] H. V. D. Parunak, J. Kindrick, and B. W. Irish. A Connectionist Model for Material Handling. *Robotics & Computer-Integrated Manufacturing*, 4(3/4):643-654, 1988. URL http://www.erim.org/~vparunak/RCIM88.pdf.

[19] P. Peeters, P. Valckenaers, J. Wyns, and S. Brueckner. Manufacturing Control Algorithm and Architecture. In *Proceedings of Second International Workshop on Intelligent Manufacturing Systems*, Leuven, Belgium, pages 877-888, K.U. Leuven, 1999.

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

[21] J. Sauter, H. V. D. Parunak, S. A. Brueckner, and R. Matthews. Tuning Synthetic Pheromones With Evolutionary Computing. In *Proceedings of Genetic and Evolutionary Computation Conference Workshop Program, 2001*, San Francisco, CA, pages 321-324, 2001. URL http://www.erim.org/~vparunak/ECOMAS2001.pdf.

[22] SMDC-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, 2001.

[23] L. Steels. Cooperation between Distributed Agents through Self-Organization. Vrije Universiteit Brussel AI Laboratory, 1989.

[24] Swarm Development Group. Swarm Development Group. 2001. Web page, www.swarm.org.

[25] T. White and B. Pagurek. Towards Multi-Swarm Problem Solving in Networks. In *Proceedings of Third International Conference on Multi-Agent Systems (ICMAS'98)*, Paris, FR, pages 333-340, IEEE Computer Society, 1998.