Forthcoming, Proceedings of the First International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2002, Bologna, Italy)

Synthetic Pheromone Mechanisms for Coordination of Unmanned Vehicles

H. Van Dyke Parunak
ERIM
PO Box 134001
Ann Arbor, MI 48113-4001
+1 734 623 2509
vparunak@erim.org

Sven Brueckner
ERIM
PO Box 134001
Ann Arbor, MI 48113-4001
+1 734 623 2529
sbrueckner@erim.org

John Sauter
ERIM
PO Box 134001
Ann Arbor, MI 48113-4001
+1 734 623 2513
jsauter@erim.org

ABSTRACT

Agents guided by synthetic pheromones can imitate the stigmergetic dynamics of insects. The resulting software architecture is well suited to problems such as the control of unmanned robotic vehicles. We introduce the approach, describe the mechanisms we have developed, and summarize the technology's performance in a series of scenarios reflecting military command and control.

Keywords

Fine-grained agents, stigmergy, pheromones, biomimetics

1. INTRODUCTION

Many social insect species coordinate the activities of individuals in the colony without direct communication or complex reasoning. Instead, they deposit and sense chemical markers called "pheromones" in a shared physical environment that participates actively in the system's dynamics. The resulting coordination is robust and adaptive. Seeking such characteristics in engineered systems, we have developed a software runtime environment that uses synthetic pheromones (data structures inspired by the insect model) to coordinate computational agents using mechanisms similar to those of social insects.

We have applied synthetic pheromone mechanisms to the problem of controlling air combat missions, with special emphasis on unmanned air vehicles. [11]. In the course of our experimentation, we have developed several mechanisms that are promising for agent coordination in general. This report describes pheromone-based movement control as a variety of potential-field-based methods (Section 2), reviews the mechanisms we have developed (Section 3), and describes their performance in several air combat scenarios (Section 4).

2. POTENTIAL FIELDS VIA PHEROMONES

From an engineering perspective, pheromones are a particularly attractive way to construct a potential field that can guide coordinated physical movement.

2.1 Potential Fields

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 [14], 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 use a potential field to guide unmanned robotic vehicles (URV's) through the battlespace (Figure 1). In this scenario, robotic vehicles seek to destroy the tank farm, which is defended by two missile batteries. The vehicles climb a potential gradient

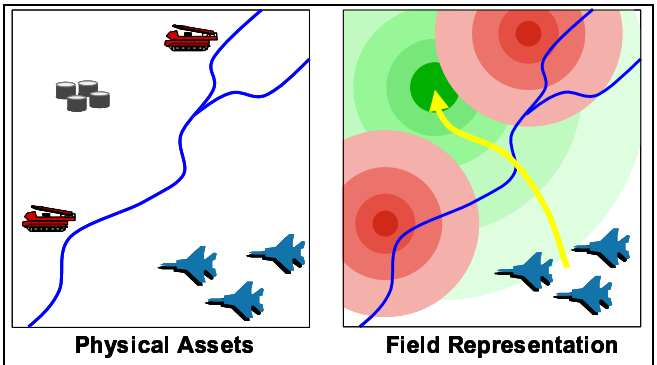

Figure 1: Potential Fields Corresponding to Physical Assets

AAMAS '02, July 15-19, 2002, Bologna, Italy.

centered on the tank farm while avoiding gradients centered on the threats. To be useful in warfighting, this field requires four characteristics (mnemonically, "4-D"):

Diverse.—It must fuse information of various types and from various sources, including targets to be approached, threats to be avoided, and the presence of other URV's with whom coordination is required.

Distributed.—Centralized processing of a potential field imposes bottlenecks in communications and processing, and generates localized vulnerabilities to attack. Ideally, the potential field should be stored close to where the information that it integrates is generated, and close to where it will be used.

Decentralized.—Efficiency and robustness also dictate that components of the system be able to make local decisions without requiring centralized control, ideally on the basis of nearest-neighbor interactions with one another.

Dynamic.—The battlespace is an uncertain and rapidly changing environment, and the methods and architecture used to construct and maintain the field must be able to incorporate such changes rapidly into the field.

An architecture inspired by insect pheromones satisfies these requirements, and can be applied to warfighting scenarios.

2.2 Synthetic Pheromones

Insects coordinate without direct communication, by sensing and depositing pheromones (chemical markers) in the environment [10]. For example, the networks of paths that they construct joining their nests with available food sources form minimum spanning trees [5], minimizing the energy ants expend in bringing food into the nest. This structure emerges as individual ants wander, depositing and sensing pheromones.

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. Traditional 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 detecting inconsistencies that result from changes in the domain. 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 the 4-D characteristics:

Diverse.—Ants can respond to combinations of pheromones, thus modifying their reaction to multiple inputs at the same time.

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) [2]. Thus new information is quickly integrated into the field, while obsolete information is automatically forgotten, through pheromone evaporation.

An implementation of synthetic pheromones has two components: the *environment* (which maintains the pheromone field and performs aggregation, evaporation, and diffusion), and the *walkers* (which deposit and react to the field maintained by the environment). Our implementation has two corresponding species of agents. A set of *place agents* with a Neighbor relation defining adjacency makes up the environment, and each walker is represented by a *walker agent*.

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 or by propagation from a neighboring place), 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, are described elsewhere [2]. If the strength of the pheromone at a location drops below a threshold, the software no longer processes that pheromone, and it disappears.

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 JFACC, 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 as a function of their flavors, and deposit pheromones into the place. It can also determine from the place agent the relative strength of a given flavor at the place and at each of its neighbors. A walker moves from one place to another by spinning a roulette wheel whose segments are weighted according to this set of strengths.

Such techniques can play chess [4] and do combinatorial optimization [1], and we have applied them to manufacturing [2] and military C^2 [11].

3. BASIC MECHANISMS

We have explored several basic mechanisms essential to the engineering deployment of pheromone mechanisms. These fall into three broad categories: combinations of multiple pheromones, using history in movement decisions, and ghost agents. Some of the results discussed in this section are expounded at more length in other publications, but are drawn together here so that they can be more readily considered as an integrated system.

3.1 Pheromone Vocabulary

There are two ways in which the pheromone vocabulary can be multiplied. First, different flavors may reflect different features of the environment (e.g., Red (hostile) air defenses, Blue (friendly) bases). These flavors have different *semantics*. Second, different flavors with the same semantics (e.g., all generated by the same

feature) may differ in their evaporation or propagation rate or threshold, thus having different *dynamics*.

Pheromones with Different Semantics.—We explored the effect of increasing the semantics of a pheromone vocabulary in the context of the classic missionary-cannibal problem [12]. Three missionaries and three cannibals are together on one bank of a river, with a dugout canoe capable of carrying only one or two people. If at any time the cannibals outnumber the missionaries on either bank of the river, they will eat them. The problem is to plan a sequence of moves that gets all six people safely across the river.

At each decision epoch, only those agents on the bank with the boat make a movement decision. Each such agent decides whether to move by evaluating a personal choice function that returns a real number between 0 and 1, evaluating a random variable uniformly distributed on [0,1], and comparing these two values. If the random number is less than the value of the choice function, the agent volunteers to move. The actual riders in the boat are chosen randomly from the list of candidates.

The details of the agent's decision are embedded in its choice function, which is a function of the levels of the available pheromones. In principle, each individual agent could have its own choice function, but in our experiments all Missionaries share one choice function and all Cannibals share another.

We explore the performance of the system for various combinations of three pheromones: a bank pheromone that tells agents where they are, an undifferentiated population pheromone deposited by both Missionaries and Cannibals, and distinctive Missionary and Cannibal pheromones. Our performance metric is the number of steps necessary for the system to move the agents from one bank to the other. Because of the stochastic nature of the decisions, different runs often yield different numbers of steps, and we report the median run length over 100 runs.

Figure 2 shows the result for one series of experiments, comparing three different pheromone configurations. "R" indicates the performance for agents executing a random walk. "BSmart" shows the performance when the agents have access only to a pheromone indicating which bank they inhabit (thus one bit of information). The performance at "BPSmart" results from telling them in addition the total population on their bank. Since there are six possible populations on either of two banks, the information available is $\log_2(2*6) = 3.58$. "BMCSmart" reflects the performance when missionaries and cannibals deposit distinct pheromones. There are 4*4 possible equilibrium values on each bank, but no agent will ever sense the combination {0,0}, so the total information available is $\text{Log}_2(2*(4*4-1)) = 4.91$. Figure 2 shows that log performance is linear in information content, so performance is exponential in information content.

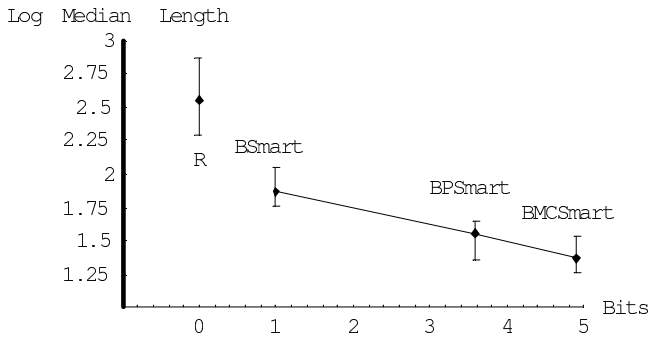

Figure 2: Performance and Pheromone Vocabulary (Error bars show inter-quartile spread)

In these experiments, the agent's choice function explicitly takes into account the levels of the different pheromones. An alternative approach, used in our air combat applications, computes a weighted function of the various input pheromones to create a single "net pheromone" whose gradient walkers then follow. In this case, the basic pheromone flavors are:

- *RTarget*: emitted by a red (hostile) target.
- *GTarget*: emitted by a blue (friendly) agent who has encountered a red target and is returning to base.
- *GNest*: emitted by a blue agent who has left the base and is seeking a target.
- *RThreat*: emitted by a red threat (e.g., missile battery)

In addition, we provide the blue agent with *Dist*, an estimate of how far away the target is.

Initially, we experimented with an equation of the form

$$\frac{\theta \cdot RTarget + \gamma \cdot GTarget + \beta}{\alpha \cdot RThreat + \delta \cdot Dist + \beta}$$

where α, β, γ, δ, and θ are tuning factors, easily manipulated in a genetic algorithm or particle swarm optimization [15]. β avoids singularities when other terms are 0. This form attracts blue agents to targets or to the trails of other blue agents who have found targets, avoids threats, and seeks to minimize distance to the target. While yielding reasonable performance, this equation left some performance gaps. Manual manipulation of the equation yielded the alternative form

$$\frac{\theta \cdot RTarget + \gamma \cdot GTarget + \beta}{(\rho \cdot GNest + \beta)(Dist + \varphi)^{(\delta + \alpha(RThreat+1))} + \beta},$$

which gives much improved performance. While more complex, this latter equation could be discovered by genetic programming.

Pheromones with Different Dynamics.—Another technique involving multiple pheromones uses pheromones with the same semantics but differing dynamics (e.g., rates of evaporation *E* and propagation *F* and threshold *S*) [3]. To motivate this mechanism, consider the distribution of pheromone sources shown in Figure 3. Each source (or background 0) is at one cell of a hexagonal grid.

We are interested in the guidance that the pheromone field offers a walker at a given place. Let f_i be the pheromone strength at place *i*. The guidance g_j available to a walker at place *j* is

$$g_j = \underset{i \in \{j\} \cup N(j)}{Max} \left(f_i \Big/ \sum f_i \right)$$

Guidance thus ranges from 0 (if all accessible places have the same pheromone strength) to 1 (if only one place has pheromone and all the others have none).

Figure 4 shows the distribution of guidance (white = 1, black = 0) for two

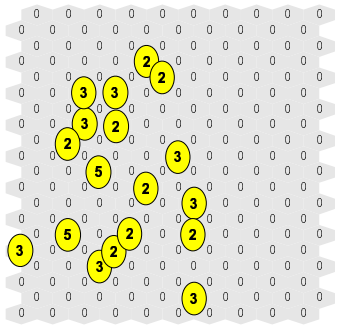

Figure 3: Test Distribution of Pheromone Sources

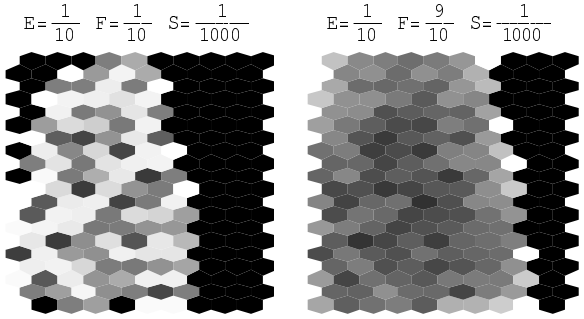

Figure 4: Guidance fields for low (left) and high (right) propagation parameters

different propagation parameters *F*. When *F* is low (left plot), most places in the target-rich region at the left of the figure have high guidance, but the pheromones do not propagate across the targetless right side of the figure, yielding a broad "valley" with low guidance. When *F* is high (right plot), propagation merges signals from individual sources, yielding low guidance in the target-rich region but a much narrower valley on the right. Thus high propagation gives good long-range guidance but poor short-range guidance, while low propagation gives good short-range guidance but poor long-range guidance.

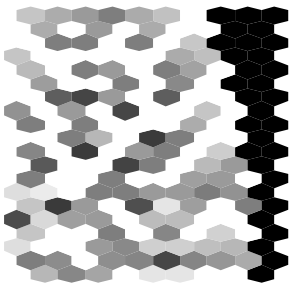

Figure 5: Guidance field with multiple pheromones

A reasonable resolution is to have each source deposit multiple pheromones with different dynamics. A walker picks its next step first by measuring the guidance available from each flavor, then computing its movement based on the pheromone with highest guidance. Figure 5 shows the guidance field from six flavors with different dynamics, yielding both high guidance in the target area, and propagation of pheromones across most of the eastern valley.

3.2 History

A walker's movement through the graph of places should balance several factors. A strong field gradient enables deterministic hill climbing that the walker should exploit. However, a weak gradient may result from noise in the system. In this case, it does not provide reliable guidance. We would prefer that the walker continue moving in the general direction of its previous steps if there is one, and otherwise that it explore more broadly.

To balance deterministic hill climbing and stochastic exploration, the walker moves from one place to another by spinning a roulette wheel whose segments are weighted according to the relative strengths of a pheromone flavor (or weighted combination of flavors) in the place and its neighbors. The mapping function from relative pheromone weight to segment width determines the degree of stochasticity in the walker's behavior. If s_i' is the perceived pheromone concentration at place *i*, the normalized weight p_i at that place is $p_i = s_i' \Big/ \sum_{j=1} s_j'$, where the summation ranges over place *i* and its neighbors, and the probability p_i' that the walker will move to that place is $p_i' = e^{\beta * p_i} \Big/ \sum_{j=1} e^{\beta * p_j}$. The parameter β determines the degree of stochasticity in the walker's movement. On a hex grid, when $\beta <$ 4, selection probabilities are more similar than the pheromone strengths would indicate, favoring exploration, while $\beta >$ 5 tends to emphasize stronger gradients, favoring exploitation.

To balance hill climbing against previous direction, we assign momentum to the walker. Models of actual ant behavior usually restrict the ant's ability to smell pheromones to some angle on either side of its current orientation. In our implementation, this technique takes the form of multiplying each segment in the walker's roulette wheel by a weight that is strongest in the direction the walker is currently heading, and weakest in the direction from which it has just come.

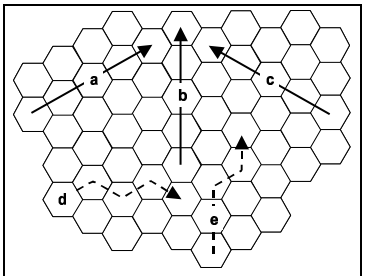

Figure 6: Path anisotropy in a hex lattice

Such a momentum works well if the walker is moving over continuous space. However, representing (continuous) space as a (discrete) graph of place agents can introduce anisotropies that confuse a simple momentum computation. Figure 6 shows five geodesics on a hex lattice. Trajectories a, b, and c maintain a constant heading, but trajectories d and e experience local direction changes while executing a shortest path across the lattice. A straightforward momentum function will interfere undesirably with these necessary changes of direction. To avoid this problem, each walker maintains an exponentially-weighted moving average of its past headings and modulates the relative strengths of the pheromones in its vicinity by a measure of the angular alignment between each candidate place and the current value of the heading history.

3.3 Ghost Agents

So far, we have distinguished stationary *place agents* (corresponding to regions of the problem space, and forming a graph structure representing the connectivity of that space) from *walker agents* (mobile agents that are associated with one place agent at a time and move among them according to the edges in the place network). For some purposes, it is useful to further refine the concept of walkers into two species.

The walker associated with a single physical robot is its *avatar*. In Hindu mythology, the term refers to an incarnation of a deity, hence, an embodiment or manifestation of an idea or greater reality. In our system, an avatar is the manifestation in our system of the greater reality (ground truth in the battlespace). A physical entity has only one avatar, which travels with the physical entity that it represents. It moves from one place agent to another only when its parent entity moves physically from one region to

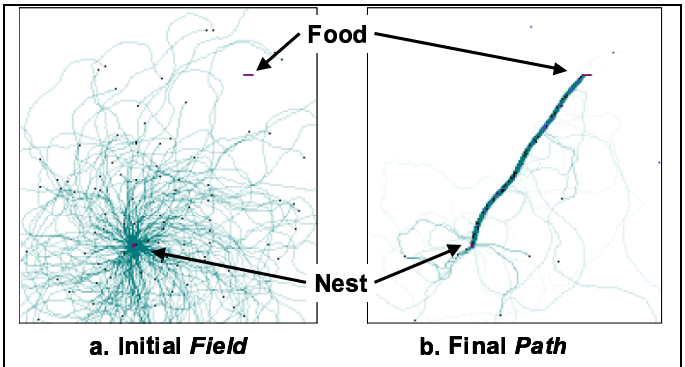

Figure 7: Path Condensation

another. Thus its speed is limited by the physical speed of its associated entity.

One avatar may send out many unembodied walkers, or *ghosts*. Ghosts move as fast as the network among place agents can carry them. Because they are more numerous than physical entities and their associated avatars, they can do "what-if" explorations that physical entities could not afford, and generate emergent behavior by their interactions. Because they move faster than physical entities and their avatars, they can look ahead to plan an avatar's next steps.

Of particular interest to robotic applications is the emergence of discrete paths in the pheromone field as many ants concurrently read and reinforce it. For example, Figure 7a shows the pheromone field deposited by a swarm of ants wandering out from their nest (at the lower left of the figure) in search of food (at the upper right). Initially, the field is roughly circularly symmetrical, and serves to guide food-bearing ants back home. Once some ants find the food and begin returning home, this field rapidly collapses into a path (Figure 7b).

At first glance, this dynamic [6] violates second-law tendencies to increasing disorder in systems consisting of many components. Left to themselves, large populations tend to disorder, not organization. Natural systems can organize at the macro level because their actions are coupled to a flow field at a micro level. Agents *perceive* and orient themselves to the flow field and reinforce that field by their *rational action*, as shown by the solid lines in Figure 8 [8]. Metaphorically, they drain unwanted entropy from the macro level (where organization is desired) to the micro level (where disorder is tolerated).

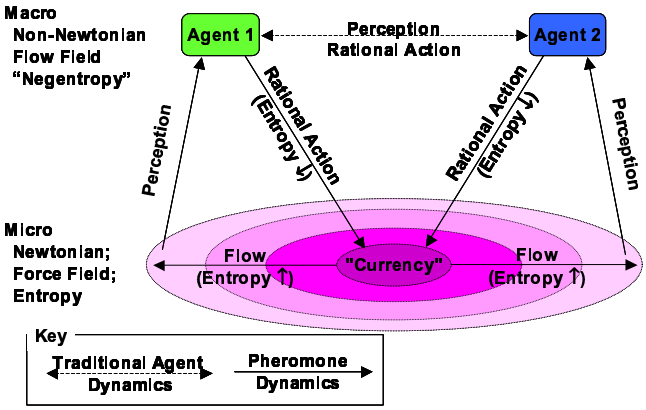

Figure 8: Coordination through Environmental Fields

Traditional coordination mechanisms ignore the micro level completely, as agents perceive and act directly on one another (dashed line in Figure 8). We link agents through the environment so that perception and action serve both to coordinate multiple agents and to control overall disorder.

We validate this mechanism of emergent coordination explicitly through experiments that compute the entropy over time of the pheromone molecules at the micro level and the agents at the macro level [13]. The increase in entropy at the micro level (through Brownian motion of pheromone molecules) more than balances the decrease in entropy experienced by walkers following the pheromone gradient.

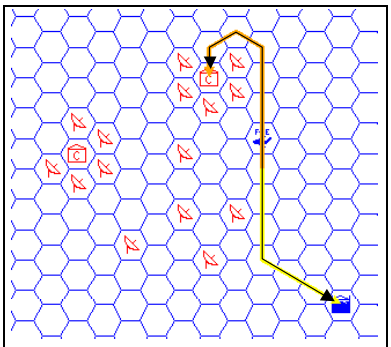

Figure 9: Path to the nearer of two targets.

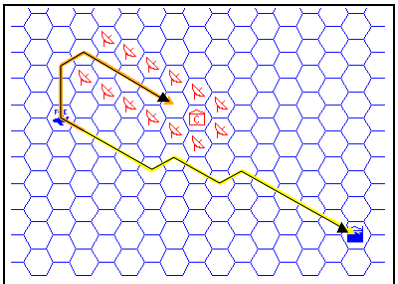

Figure 10: Threading a gauntlet.

The path emergence illustrated in Figure 7 is the result of interactions among many walkers. Each walker's behavior is highly stochastic, performing a real-time Monte Carlo search of its local vicinity, and contributing to the emergence of a long-range path. In engineering applications, it may not be feasible to ask hundreds of physical robots to explore the domain in this manner, nor is it necessary. As an avatar moves, it continuously sends out ghosts. The interaction of the ghosts forms the path, which is being constantly revised to accommodate dynamic changes in the environment.

Our experiments show this path formation dynamic to be extremely robust and adaptive. Figure 9 shows the formation of a path from a friendly airbase (lower right) to the nearer of two targets (the house-shaped icons), avoiding threats (the radar icons). If we increase the strength of the left-hand target to twice that of the closer target, the path will lead there instead. Figure 10 shows a path to a target protected by a gauntlet of threats, a configuration that resists classical potential field methods.

When ghost agents choose between two targets, they cannot tell whether one target's pheromone is stronger because it is depositing at a higher rate, or because it is nearer than the other target. We explore the balance between these factors by setting up two targets T_1 and T_2 diametrically opposite one another from the ghosts' origin, with varying ratios of distance and strength. Then we compute the percentage p_1 of runs (out of a total of 45) that form a path to T_1 rather than to T_2. Figure 11 plots of this probability as a function of the strength and distance ratios. The dots represent experimental observations, between which other values are linear interpolations. Most of the plot is dominated by regions in which p_1 is either 1 or 0. The region within which both strength and distance play an active role in target selection is

relatively narrow. As both ratios grow, the difference in distance overwhelms the difference in strength.

Another important trade in understanding the behavior of ghost agents is between time and distance. When they are far from a target, ghost agents execute a random walk. Closer to the target, they can sense the target's pheromone field, and climb its gradient. One might expect that the number of steps required to reach a target would increase precipitously as the distance between a ghost's origin and its target grows. In fact, the transition is quite well behaved (Figure 12).

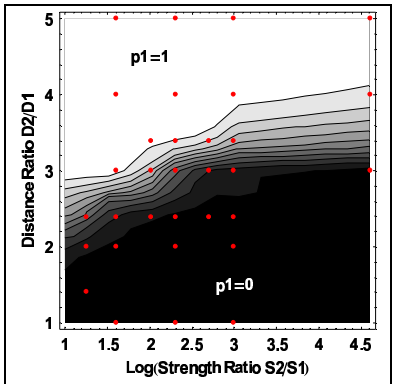

Figure 11: Trade between target strength and distance

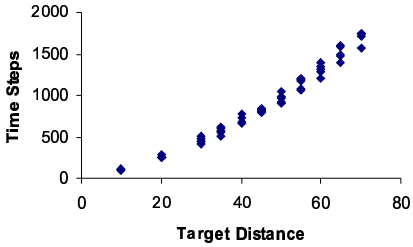

Figure 12: Time for first ghost to reach target

4. OPERATIONAL SCENARIOS

We have demonstrated these mechanisms in military air operations in four increasingly sophisticated scenarios.

4.1 SEADy Storm

SEADy Storm [7] is a war game used to explore technologies for controlling air tasking orders. The battlespace is a hexagonal grid of sectors, each 50 km across). Friendly (Blue) forces defend a region in the lower left against invading Red forces that occupy most of the field. Red's playing pieces include ground troops (GT's) that are trying to invade the Blue territory, and air defense units (AD's, surface-to-air missile launchers) that protect the GT's from Blue attack. Blue has bombers (BMB's) that try to stop the GT's before they reach the blue territory, and fighters tasked with suppressing enemy air defenses (SEAD's).

Each class of unit has a set of commands from which it periodically chooses. Ground-based units (GT and AD) choose a new command every 12 hours, while air units (BMB and SEAD) choose every five minutes, reflecting the time it would take the resource to cross a sector. The commands fall into three categories (Table 1). GT cannot attack Blue forces, but can damage BMB's if they attack GT.

Blue can attack AD and GT when they are moving or attacking, and AD may attack any Blue forces that are not moving or waiting. Each unit has a strength that is reduced by combat.

	Move	Attack	Wait
AD	Relocate	Fire (on any Blue aircraft)	Hide Deceive
GT	Advance		Hide
SEAD	NewSectors	AttackAD	Rest
BMB	NewSectors	AttackAD AttackGT	Rest

Table 1. Unit Commands in SEADy Storm.

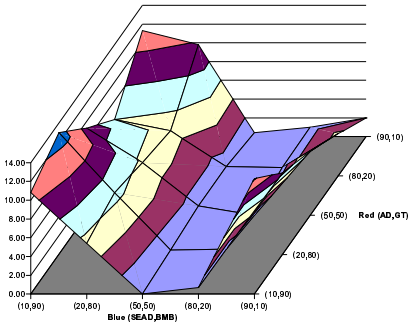

Figure 13: Red Strength in Blue as a function of force composition

The strength of the battling units, together with nine outcome rules, determine the outcome of such engagements. Informally, the first five rules are:

1. Fatigue: The farther Blue flies, the weaker it gets.
2. Deception: Blue strength decreases for each AD in the same sector that is hiding.
3. Maintenance: Blue strength decreases if units do not rest on a regular basis.
4. Surprise: The effectiveness of an AD attack doubles the first shift after the unit does something other than attack.
5. Cover: BMB losses are greater if the BMB is not accompanied by enough SEAD.

Rules 6-9 specify the percentage losses in strength for the units engaged in a battle, based on the command they are currently executing. For example, Rule 9 states: "If BMB does "AttackGT" and GT does "Advance": a GT unit loses 10% for each BMB unit per shift; a BMB unit loses 2% per GT unit per shift."

The primary parameter explored in the experiments reported here is the proportion of SEAD in the Blue military, and of AD in the Red military. Each side began with a 100 units, each with unit strength, and 10%, 20%, 50%, 80%, or 90% of SEAD or AD. The uneven spacing reflects a basic statistical intuition that interesting behaviors tend to be concentrated toward the extremes of percentage-based parameters. In current military doctrine, 50% is an upper limit on both AD and SEAD. We explore higher values simply to characterize the behavioral space of our mechanisms.)

The central outcome is total Red strength in Blue territory at the end of the run (Figure 13). The landscape shows several interesting features, including

- a "valley" of Blue dominance for all Red ratios when Blue SEAD is between 50% and 80%, with slightly increasing Red success as the AD proportion increases;
- clear Red dominance for lower SEAD/BMB ratios, decreasing as SEAD increases;
- a surprising increase in Red success for the high SEAD and low AD levels.

Figure 14 compares the population of Red in Blue territory as a function of red and blue force composition for three different Blue control strategies. In the top plot, Blue does not use pheromones at all. The variations are due to the intrinsic dynamics of the combat, yielding a narrow valley up the center of the plot where Red's population is 3 or less (the criterion for Blue victory). When Blue uses pheromones to seek out Red targets and threats (middle plot, shown in profile in Figure 13), the wider valley reflects improved Blue performance. In the bottom figure,

when Blue uses pheromones to avoid threats and approach targets, the valley with the lowest Red population is about the same area but of a very different shape than in the previous case, but the next level of Red occupation (4-6) is much larger, showing a reduction in higher levels of Red occupation.

A detailed discussion of the dynamics of this scenario and effects when we change the modeling formalism is available at [9].

4.2 CyberStorm

At the next level of sophistication, we expand the range of unit types. Red now has armored and infantry battalions, air defense units, distinct headquarters types for regiments, air defense, and the entire corps, and fueling stations. Blue has three types of fighters and two types of bombers. The environment includes bridges and road crossings (which speed the movement of ground units that encounter them) and oil fields (which Red seeks to attack and Blue seeks to protect). Combat outcome is based on the percentage survival of the oil fields.

Using this enriched environment, we have explored a variety of issues around blue decision-making. In these experiments (as in SEADy Storm), Blue resources move directly in response to Red pheromones, without using ghosts. Our experiments show that reasonable numbers of Blue resources cannot sample the pheromone field adequately to overcome the stochasticity inherent in the domain. As a result, outcomes vary widely with random seeds. These experiments demonstrated the need for ghost agents to sample the primary pheromone field at a statistically more significant level, and preprocess it for use by Blue avatars and the physical resources with which they are associated.

4.3 Super Cyber Storm

We exercised the ghost agents on a third model of the domain, which includes a significantly wider range of entity types, combat resolution on the basis of individual weapon type rather than unit type, more realistic dependencies among entities (for example, the effectiveness of Red air defense now depends on the status of other Red air defense units), and most importantly, a "pop-up" Red capability that lets us increase greatly the range of changes in Red's visibility as a scenario unfolds. This environment permits us to assess the effectiveness of ghost-based pheromones in dealing with pop-up threats.

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 a 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.

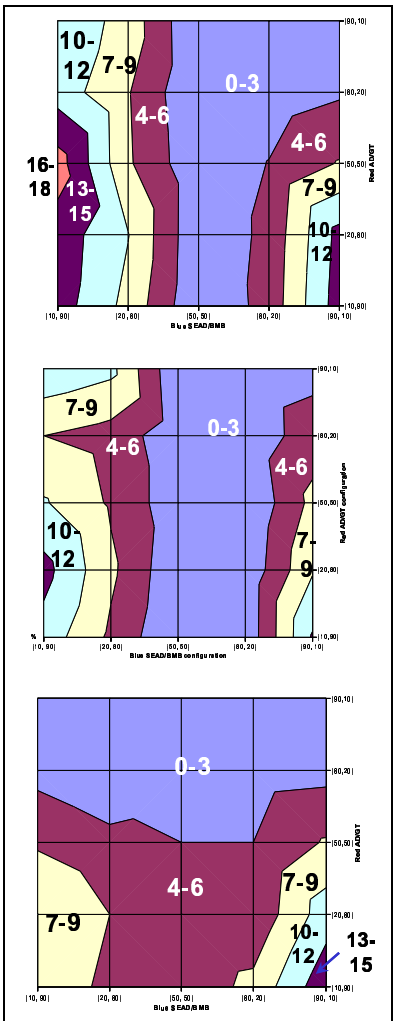

Figure 14: Population of Red in Blue without pheromones (top), with pheromones attracted to both targets and threats (center), and attracted to targets but repelled by threats (bottom)

Figure 15 shows the medians over five runs of Red and Blue total unit strengths for 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.

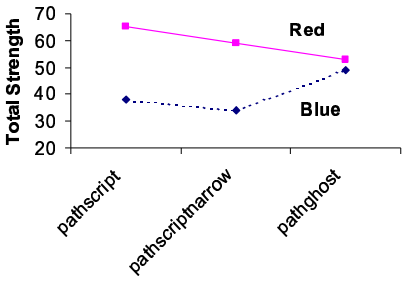

Figure 15: Effect of Ghost Agents in Super Cyber Storm

4.4 Swarming UAV Experiment

Recently, these algorithms have been applied successfully to an experiment on the effectiveness of swarming UAV's (unmanned air vehicles) in suppressing antiaircraft threats in a wargame simulation conducted by the U.S. military. The pheromone approach shows significant performance improvements over the baseline. The public report has not been released, but will be by the time the final version of this paper is available, and details will be included in the publication version.

5. SUMMARY

Synthetic pheromones are a powerful mechanism for controlling the movement of agents through space. They provide the elegance of potential field methods, with particular support for integrating diverse information sources, processing information in a completely distributed and decentralized environment, and coping with dynamic changes in the landscape. In exploring successively complex military scenarios, we have developed a toolkit of methods and mechanisms, including pheromone vocabularies that vary in both semantics and dynamics, mechanisms for incorporating agent momentum into movement decisions, ghost agents to preprocess the pheromone field and reduce stochasticity at the level of physical resources, and visualization mechanisms to enable human stakeholders to understand and monitor the emergent behavior of the system.

6. ACKNOWLEDGMENTS

This work is supported in part by the DARPA JFACC program under contract F30602-99-C-0202 to ERIM CEC, under DARPA PM's COL D. McCorry and MAJ S. Heise, Chief Technologist Steve Morse, and Rome Labs COTR's C. Defranco and T. Busch. 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 Defense Advanced Research Projects Agency or the US Government. Realization of these concepts owes much to our research team, including E. Feibush, E. Greene, O. Gilmore, R. Matthews, and M. Nandula. Portions of this technology are covered by US and international patents pending.

7. REFERENCES

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

[2] S. Brueckner. *Return from the Ant: Synthetic Ecosystems for Manufacturing Control*. Thesis at Humboldt University Berlin, Department of Computer Science, 2000.

[3] S. Brueckner and H. V. D. Parunak. Multiple Pheromones for Improved Guidance. In *Proceedings of Symposium on Advanced Enterprise Control*, 2000.

[4] 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)*, pages 13-27, Springer, 1995.

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

[6] D. Helbing, F. Schweitzer, J. Keltsch, and P. Molnár. Active Walker Model for the Formation of Human and Animal Trail Systems. Institute of Theoretical Physics, Stuttgart, Germany, 1998. URL http://xxx.lanl.gov/ps/cond-mat/9806097.

[7] A. Kott. SEADy Storm. 2000. Word file, http://www.erim.org/cec/projects/adaptiv/SeadyStorm-v1.2.doc.

[8] P. N. Kugler and M. T. Turvey. *Information, Natural Law, and the Self-Assembly of Rhythmic Movement*. Lawrence Erlbaum, 1987.

[9] H. V. Parunak, S. Brueckner, J. Sauter, and R. Matthews. Distinguishing Environmental and Agent Dynamics: A Case Study in Abstraction and Alternative Modeling Technologies. In A. Omicini, R. Tolksdorf, and F. Zambonelli, Editors, *Engineering Societies in the Agents' World (ESAW'00)*, vol. LNAI 1972, *Lecture Notes in Artificial Intelligence*, pages 19-33. Springer, Berlin, 2000.

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

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

[12] H. V. D. Parunak and S. Brueckner. Ant-Like Missionaries and Cannibals: Synthetic Pheromones for Distributed Motion Control. In *Proceedings of Fourth International Conference on Autonomous Agents (Agents 2000)*, pages 467-474, 2000.

[13] H. V. D. Parunak and S. Brueckner. Entropy and Self-Organization in Multi-Agent Systems. In *Proceedings of International Conference on Autonomous Agents*, pages (forthcoming), 2001.

[14] 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.

[15] J. Sauter, H. V. D. Parunak, S. Brueckner, and R. Matthews. Tuning Synthetic Pheromones With Evolutionary Computing. In *Proceedings of EcoMAS 2001*, 2001.