Presented at the Conference on Swarming and C4ISR, Tysons Corner, VA, 3 Jan 2003

Making Swarming Happen

H. Van Dyke Parunak, Ph.D.
Chief Scientist, Altarum Institute
van.parunak@altarum.org

Abstract

The concept of swarming has been invoked to describe both human military tactics and the behavior of simple biological systems. Most research on the underlying mechanisms of swarming is in the biological community, and it is often not clear how these can be implemented in military systems. This paper reviews several architectures that have shown promise for generating swarming behavior in military systems, and briefly discusses how to measure and control such activity.

1 Introduction

The word "swarming" is currently in vogue to describe two widely different types of systems. Students of biological systems use it to describe decentralized self-organizing behavior in populations of (usually simple) animals [8, 9, 16, 40]. Table 1 lists a few examples that have been studied. Military historians use it to describe a battlefield tactic that involves decentralized, pulsed attacks [2, 22, 23, 33].

The link between these two uses of the word is not coincidental. Insect self-organization is robust, adaptive, and persistent, as anyone can attest who has tried to keep ants out of the kitchen or defeat a termite infestation, and military commanders would love to be able to inflict the frustration, discomfort, and demoralization that a swarm of bees can visit on their victims. The linkage between swarming and warfare is ancient. In the Bible, God promises to demoralize the indigenous population of Canaan before the invading Israel-

Table 1: Some Examples of Swarming in Nature

Swarming Behavior	Entities
Pattern Generation	Bacteria, Slime Mold
Path Formation	Ants
Nest Sorting	Ants
Cooperative Transport	Ants
Food Source Selection	Ants, Bees
Thermoregulation	Bees
Task Allocation	Wasps
Hive Construction	Bees, Wasps, Hornets, Termites
Synchronization	Fireflies
Feeding Aggregation	Bark Beetles
Web Construction	Spiders
Schooling	Fish
Flocking	Birds
Prey Surrounding	Wolves

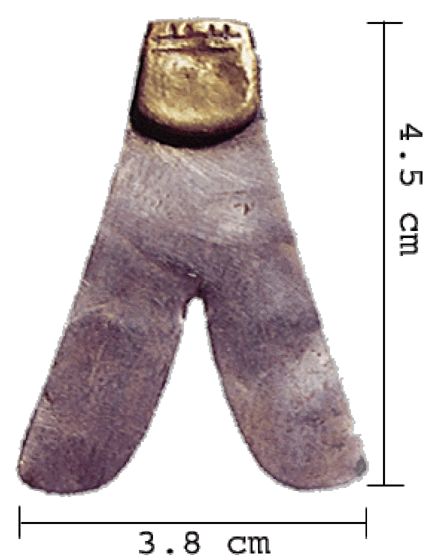

Figure 1: Egyptian "fly" medal for military heroes, 1550 BC (National Gallery of Art)

ites in the words, "I will send the hornet before you" (Exodus 23:28; cf. Deuteronomy :20; Joshua 24:12). In the eighteenth dynasty (1550 BC), the ancient Egyptians awarded military heroes a gold and silver medal in the form of a stylized fly (Figure 1) [32], and there are documented cases of the ancients' hurling hives of stinging insects against their enemies [39].

In spite of the military promise of swarming, little attention has been given to how to implement the mechanisms observed in biological communities into military systems. To the contrary, many conventional aspects of military C^3I, such as centralized command and assumptions about the availability of high-bandwidth battlefield communications, may actually make it more difficult to achieve swarming behavior. This paper bridges this gap by reporting on a number of architectures that have been used to produce swarming. Section 2 defines swarming more precisely and outlines its value to modern warfighting. Section 3 describes several architectures that support swarming, and Section 4 briefly reviews how one can measure and control swarming to ensure that the commander's intent is being achieved. Section 5 offers a summary and conclusion.

2 What is Swarming, and Why is it Desirable?

Numerous definitions have been proposed for swarming. We compare these definitions, offer a synthesis, and discuss some of the benefits of a swarming approach in modern warfare.

2.1 Definitions

Definitions of swarming have been proposed by insect ethologists, roboticists, and military historians. Of the many definitions that have been proposed, a few will illustrate the main themes.

Biologists studying swarming (e.g., [9]) define swarming as "distributed problem-solving devices inspired by collective behavior of social insect colonies and other animal societies."

The use of the term to describe artificial systems can be traced to Beni, Hackwood, and Wang in the late 1980's [4-7, 30, 31]. Their work focuses on populations of cellular robots, and they use the term to describe self-organization through local interactions. In the context of unpiloted air vehicles (UAV), Clough defines a swarm as a "collection of autonomous individuals relying on local sensing and reactive behaviors interacting such that a global behavior emerges from the interactions" [17]. He distinguishes swarming (resulting from reactive behaviors of simple homogeneous entities performing simple tasks) from the emergent behavior of heterogeneous teams of deliberative entities performing complex tasks.

Military historians focus less on the process of self-organization and more on the resulting organization itself: "the systematic pulsing of force and/or fire by dispersed, internetted units, so as to strike the adversary from all directions simultaneously" [2]; a "scheme of maneuver" consisting of "a convergent attack of several semi-autonomous (or autonomous) units on a target" [23].

For the purpose of this paper, we will define swarming as "useful self-organization of multiple entities through local interactions." This definition highlights elements of the others that have been suggested.

"Useful" emphasizes that we are interested in engineering systems that are answerable to a commander for their behavior. Some forms of self-organized behavior, such as riots and oscillation, might be interesting to a biologist, but undesirable in a robotic or military application. The military definition of swarming as convergent attack describes one form of useful behavior that an engineered military swarm should be able to produce. From a C4ISR perspective, we should expand the range of system-level behaviors that we consider useful. Swarming mechanisms can support many functions other than convergent attack, including maintaining communications networks, recognizing patterns in sensor arrays, and coordinating multi-phase missions.

Self-organization is most prominent in the robotic definitions, since the concern there is to distinguish swarming from conventional top-down control schemes. The military definition does not emphasize self-organization, perhaps because of a historic tradition of top-down centralized control. A system that can organize itself will be able to configure itself when deployed, helping to achieve the Objective Force requirements of deploying a brigade in 96 hours, a Division in 5 days, and five Divisions in 30 days [45]. It will be more robust in the fog and friction of war, and will adapt quickly to alternative missions. We do not require that the self-organization result from reactive rather than deliberative individual behavior. Thus our definition includes not only Clough's "swarms" but also his "teams," if they meet the other terms of the definition.

The notion of **multiple entities** is common to all the earlier definitions, and indeed is intrinsic to the common-sense use of the term. A major motivator for swarming in military operations is the proliferation of robotic platforms, such as vehicles and sensor systems. Although these systems are often referred to as "unmanned," in current practice it would be more accurate to describe them as "remotely manned." The flight crew for a Predator consists of two people. Housing them in a control van rather than on board the flying platform considerably reduces their risk, but does not reduce the manpower requirements for fielding the vehicle. A major promise of swarming is multiplying the number of platforms that a single warfighter can effectively control.

Our focus on **local interactions** has two motivations: a need and a promise. The *need* is a growing concern about the availability of long-range high-bandwidth links on the battlefield. The *promise* is the observation that local interactions suffice to maintain long-range coordination in biological systems, so that we ought to be able to reverse-engineer the underlying mechanisms for use in military systems.

2.2 Desirability

Swarming as defined above is appropriate for problems with four characteristics, which we summarize mnemonically as D^4: Diverse, Distributed, Decentralized, and Dynamic.

Diverse.—Swarming handles a number of forms of diversity, beginning with the multiple *platforms* in the swarm itself. It can integrate diverse *functions*, including communications among platforms, command oversight, and information management to enable the platforms to make reasonable decisions. It can 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. Warfighters cannot assume the availability of unlimited satcom channels [1], placing a premium on innovative processing mechanisms as a way of reducing bandwidth requirements [18]. Distributing the C^2 system physically over the battlespace addresses this problem in three ways.

1. The "local interaction" nature of swarming means that entities need communicate mainly with nearby neighbors using relatively low power, thus permitting bandwidth to be reused beyond a local horizon.
2. The distributed entities in a swarm can themselves form a communication network that propagates messages long distances through multiple short-range hops [27].
3. 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 central repository. This strategy greatly reduces the need to move large bodies of information over long distances.

Decentralized.—The self-organizing capability of a swarm, based on the local autonomy of individual entities, reduces the need for detailed centralized C^2. There are several motives for allowing members of a swarm to make local decisions, within the scope of responsibility originally assigned to them, without detailed explicit commands from a central point.

- The current (centralized) model of robotic control requires a team of two or three humans for each entity. The manpower costs of this model make it prohibitive to field large numbers of elements.
- The time delay for a central command to process data from sensing elements and generate new commands is unacceptable in rapidly changing combat situations.
- 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 elements. 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 [35] may itself generate unexpected changes in the situation. The self-organizing ability of swarms enables them to respond autonomously to such changes.

3 How can we Generate Swarming?

There are three major approaches to the command and control of multiple robotic entities, which can be distinguished on the basis of the location in the architecture at which intelligent decisions are concentrated (Table 2). Centralized command and control, which is not swarming under our definition, treats the centralized commander as the main locus of intelligence. Classical AI mechanisms seek to endow the individual entity with local intelligence, while stigmergic mechanisms generate system-level intelligent behavior through the interactions among entities that individually may not exhibit high levels of intelligence.

Table 2: Approaches to Controlling Multiple Entities

Approach	Locus of Intelligence
Centralized C^2	Distinguished agent ("commander")
Classical AI	Each individual agent
Stigmergy	Interactions among agents

3.1 Centralized C^2

The classic model of centralized C^2 envisions decisions being made by the central commander and then propagated through a hierarchy of subordinates for refinement and execution. Most intelligence is concentrated in the central commander, not in the other entities in the system. This model satisfies only the first and third elements of our definition, and so does not constitute swarming.

- The resulting behavior may be **useful**, although the time delays associated with central decision-making may compromise its performance.
- It certainly is not **self-organizing**. Rather, the organization is defined by the central authority and imposed top-down on the elements.
- It can coordinate **multiple entities**, although limitations in human span of attention require multiple layers of hierarchy as the number of entities increases.
- **Local communications** are not generally sufficient for such schemes. In fact, much of the emphasis on long-range high-bandwidth communications capacity in modern military systems can be traced to the desire to support centralized command.

3.2 Intelligent Agents

Most research in multi-agent systems (e.g., [52]) seeks to endow the individual agent with intelligence. At some level, most of these systems can be described as a finite-state machine (FSM). Work by the team at JHU-APL [50, 51] is an important example of how these systems can be applied to unpiloted vehicles.

The basic idea of an FSM is that the agent at any time is in one of a number of distinct states, and moves from one state to another based on events that it experiences or conditions that it detects in the environment. Figure 2 is an example of a fragment of an actual UAV FSM. This fragment records four states in which the UAV may be, indicated by non-italic type, and six conditions under which it may change state, in italics. For example, if the UAV is in the "Global Search" state and receives a message requesting it to search a particular area, it enters the "Goto Search Point" state, which invokes the flight control system to move the UAV to the waypoint specified in the search request message. When the UAV reaches that waypoint, it enters the "Local Search" state, which now tells the avionics to perform local search operations. If two transitions out of a state are enabled, the choice between them is made nondeterministically.

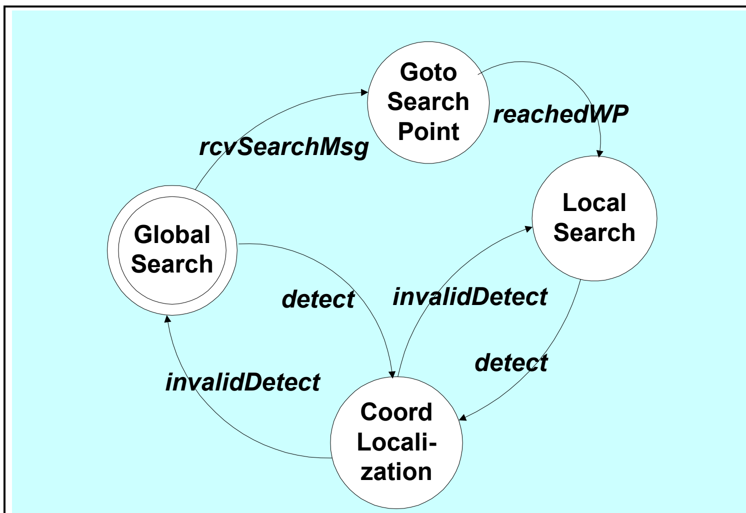

Figure 2: Fragment of a UAV FSM (from [51])

As this example shows, interactions among entities are one class of events that can trigger a transition from one state to another. Such interactions may be of different types: one entity may *detect* the presence of another (e.g., to avoid collision), *command* another to carry out an action (thus implementing a centralized C^2 scheme), or *negotiate* with other entities. A population of individually intelligent agents can function in different modes, including central command and what we have defined as swarming. Swarming will result if the interactions are local and if resulting interplay of the individual FSM's leads to self-organization.

Many refinements of FSM's hold promise for swarming military systems. Three in particular are the use of nondeterminism within states, the subsumption architecture, and market models.

- As described in [50, 51], the behavior of a platform while in a given state is deterministic, and the system exhibits nondeterminism only in deciding among multiple eligible transitions. Adding nondeterminism within states, and refining its application in state transitions, can make FSM's more robust, as discussed in Section 3.3.2 below.
- In the subsumption architecture [11, 34], the states are arranged in a series of levels. Lower-level states define default behaviors that may be overridden by higher-level states in the presence of specific stimuli. The RAND PRAWN model of swarming UAV's [26] uses this architecture.
- Market models [20, 49] structure the negotiations among entities in economic terms, building on an early agent protocol known as the "contract net" [19]. Agents responsible for high-level tasks publish descriptions of subtasks, for which other agents bid according to their capabilities. The market dynamics rationalize the allocation of tasks across agents.

FSM's and other architectures with individually intelligent agents offer a number of benefits. The states of each vehicle map directly onto features of the domain in a way that human users can easily understand, and the interactions of agents are also understandable. A set of well-chosen states and transitions can be configured to address a variety of tasks and missions. These architectures can control all aspects of the system, not just platform motion, and can be applied to heterogeneous platforms.

At the same time, these architectures pose some notable challenges. Because the states are defined at a high conceptual level, a FSM can be brittle, and its performance can deteriorate when its predefined collection of states and transitions does not map neatly onto the environment. (Adding nondeterminism can reduce this brittleness.) Defining the correct set of states and transitions is a non-trivial knowledge engineering task. The architecture encourages commands between entities rather than negotiation, leading to rigid task assignment (although market models can be more flexible). As the number of platforms increases, coordination becomes increasingly difficult, leading to scaling problems with large populations.

3.3 Stigmergic Systems

"Stigmergy" is a term coined in the 1950's by the French biologist Grassé [29] to describe a broad class of multi-agent coordination mechanisms that rely on information exchange through a shared environment. The term is formed from the Greek words *stigma* "sign" and *ergon* "action," and captures the

Table 3: Varieties of Stigmergy

	Marker-Based	Sematectonic
Quantitative	Gradient following in a single pheromone field	Ant cemetery clustering
Qualitative	Decisions based on combinations of pheromones	Wasp nest construction

notion that an agent's actions leave signs in the environment, signs that it and other agents sense and that determine their subsequent actions. Different varieties of stigmergy can be distinguished. One distinction concerns whether the signs consist of special markers that agents deposit in the environment ("marker-based stigmergy") or whether agents base their actions on the current state of the solution ("sematectonic stigmergy"). Another distinction focuses on whether the environmental signals are a single scalar quantity, analogous to a potential field ("quantitative stigmergy") or whether they form a set of discrete options ("qualitative stigmergy"). As shown in Table 3, the two distinctions are orthogonal.

Whatever the details of the interaction, 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 the centralized model) nor in each individual agent (the intelligent agent model), but in the interactions among the agents and the shared dynamical environment.

Stigmergic mechanisms have a number of attractive features for military systems.

Simplicity.—The logic for individual agents is much simpler than for an individually intelligent agent. This simplicity has three collateral benefits.

1. The agents are easier to program and prove correct at the level of individual behavior.
2. They can run on extremely small platforms (such as microchip-based "smart dust" [44]).
3. They can be trained with genetic algorithms or particle-swarm methods rather than requiring detailed knowledge engineering.

Scalable.—Stigmergic mechanisms scale well to large numbers of entities. In fact, unlike many intelligent agent approaches, stigmergy *requires* multiple entities to function, and performance typically improves as the number of entities increases.

Robustness.—Because stigmergic deployments favor large numbers of entities that are continuously organizing themselves, the system's performance is robust against the loss of a few individuals. The simplicity and low expense of each individual means that such losses can be tolerated economically.

Environmental Integration.—Explicit use of the environment in agent interactions means that environmental dynamics are directly integrated into the system's control, and in fact can enhance system performance. A system's level of organization is inversely related to its symmetry (Figure 3), and a critical function in achieving self-organization in any system made up of large numbers of similar elements is breaking the natural symmetries among them [3]. Environmental noise is usually a threat to conventional control strategies, but stigmergic systems exploit it as a natural way to break symmetries among the entities and enable them to self-organize.

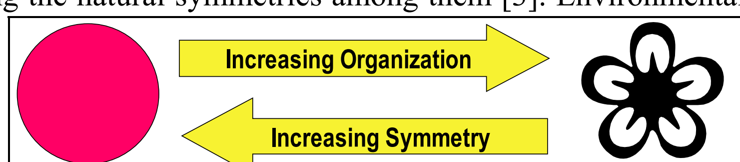

Figure 3: Symmetry vs. Organization.—Achieving high levels of organization requires breaking the symmetry among system components, a function that environmental noise supports.

A subset of stigmergic mechanisms, known as "coordination fields" or "co-fields" [36-38, 51], consists of quantitative stigmergy (scalars mapped to the problem topology). The scalar field is generated by a combination of attracting and repelling components, and the agents follow gradients in this field, thus tending to avoid repellers and approach attractors. Such techniques have an extended history in controlling individual robots [46]. Among numerous instances of this approach to modeling and controlling swarms of multiple entities, three are illustrative, summarized in Table 4 and discussed in the next three subsections.

Table 4: Examples of Co-Fields for Military Swarming

Sandia Laboratories	Particle Swarms	Simulated Coulomb forces among the particles
Johns Hopkins APL	Extensions to FSM's	Gradients along contact vectors in space-time
Altarum	Digital Phero-mones	Agents deposit and sense digital pheromones in computational infrastructure that supports Aggregation, Evaporation, Diffusion

3.3.1 Sandia: Particle Swarms

The particle swarm approach to modeling swarms is based on an analogy between swarming entities and physical particles: For example, one can view a swarm of robots as particles in a gas:

- They have momentum.

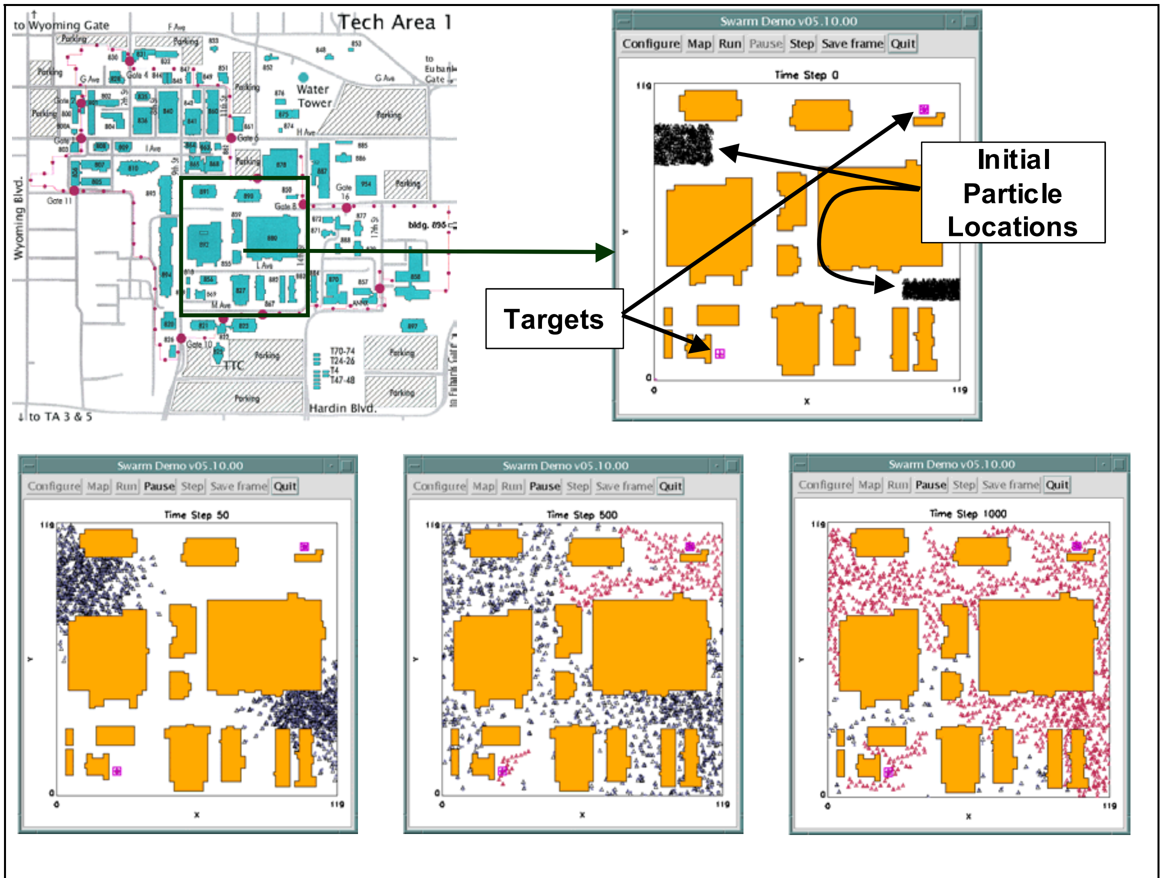

Figure 4: Ballistic simulation of a swarm in a complex urban environment.

- They respond to potential fields in the environment (including real potentials like gravity, drag, and propulsion, and artificial ones like repulsion from collisions and attraction to targets).
- When they hit another robot or a wall, they bounce as in an elastic collision, responding to a simulated Coulomb field.

Table 5: Modeling Swarms with Particle Physics Codes

Problem	Simulation Code
UAV's	Particle-in-Cell (PIC)
UGV in Simple Environment	Lattice Gas Automata
UGV in Complex Urban Environment	Ballistic Code
	Particle-in-Cell (PIC)

Robots bounce around the environment until they find something interesting, then stop. Sandia has explored a variety of such models, leveraging existing software codes for physical particles (Table 5) [24]. This approach offers several advantages:

- Existing particle simulation codes are benchmarked.
- The statistical mechanics of finite-sized particles are well documented and understood.
- These codes are highly efficient: some scale as O(n).
- Single and parallel computer implementations are available.
- The codes can support heterogeneous particles.

Figure 4 shows the result of such a simulation in a complex urban environment. The vehicles start in two large clumps along the edges of the figure, and their task is to find the two targets at the upper right and lower left. Vehicles move following ballistic laws until they learn of the location of a target, then they stop (and turn a lighter shade in the display). When a vehicle learns the location of a target (either by running into it, or by hearing from another vehicle), it broadcasts the location to others. Though the communication range of an individual vehicle is very small, eventually the knowledge propagates through the entire community.

3.3.2 JHU APL: Extensions to FSM's

The JHU-APL team is applying co-fields to the coordination and control of UAV's [51]. Their approach has two important features.

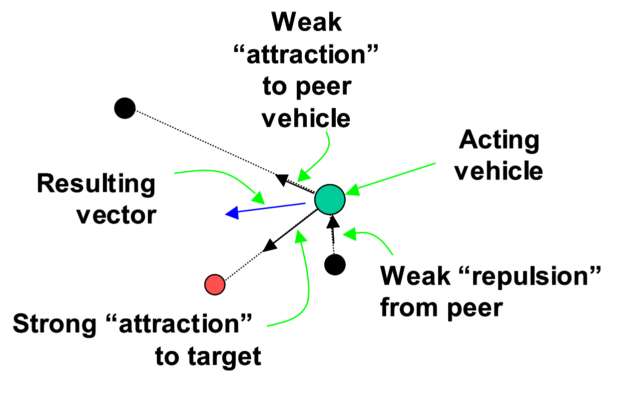

Figure 5: Local estimation of co-field from propagated influence vectors.

First, they use co-fields in a hybrid architecture, jointly with FSM's. FSM's provide high-level control of entities, and co-fields define the detailed behavior of an entity while it is in a single state. They are also being explored as a way to resolve ambiguous transitions between states. This hybrid architecture illustrates an important point that can be generalized across the techniques summarized in this paper: swarming technologies are not necessarily mutually exclusive, but can often be combined in complementary ways to solve a problem.

Second, they do not attempt to generate a co-field covering the entire problem space, but rather propagate influence vectors among neighboring vehicles. A vehicle sums the influence vectors from its neighbors with that resulting from its own sensors, and follows the resultant (Figure 5). This technique is an example of how a technique inspired by natural systems can be abstracted and simplified so that its implementation is structurally more amenable to digital computation than the biological counterpart.

3.3.3 Altarum: Digital Pheromones

Our own research has concentrated on applications of co-fields modeled rather closely on the pheromone fields that many social insects use to coordinate their behavior. We have developed a formal model of the essentials of these fields, and applied them to a variety of problems.

The real world provides three continuous processes 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 in artificial intelligence. 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.

These dynamics can be modeled in a system of difference equations across a network of "places" at which agents can reside and in which they deposit and sense increments to scalar variables that serve as "digital pheromones," and these equations are provably stable and convergent [12]. They form the basis for a "pheromone infrastructure" that can support swarming for various C4ISR functions, including path planning and coordination for unpiloted vehicles, and pattern recognition in a distributed sensor network.

Path Planning.—Ants construct networks of paths that connect their nests with available food sources. Mathematically, these networks form minimum spanning trees [28], 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.

We have adapted this algorithm to integrate ISR into a co-field that then guides unpiloted vehicles away from threats and toward targets [43]. The battlespace is divided into small adjoining regions, or "places," each managed by a "place agent" that maintains the digital pheromones associated with that place and serves as a point of coordination for vehicles in that region. The

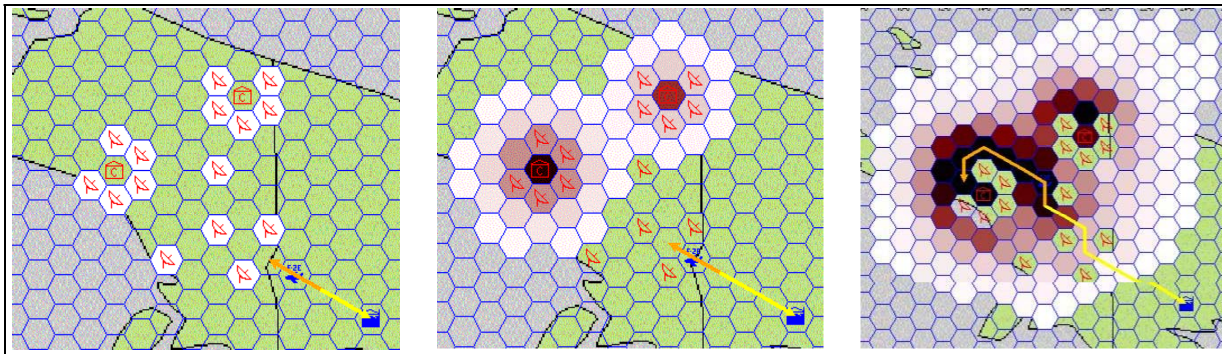

Red AD Pheromones | Red HQ Pheromones | Path Pheromones from Blue Ghosts

Figure 6: Digital Pheromones for Path Planning.

network of place agents can execute on a sensor network distributed physically in the battlespace, onboard individual vehicles, or on a single computer at a mission command center. When a Red entity is detected, a model of it in the form of a software agent is initiated in the place occupied by the Red entity, and this agent deposits pheromones of an appropriate flavor indicating the presence of the entity. The agent can also model any expected behaviors of the Red entity, such as movement to other regions. Blue agents respond to these pheromones, avoiding those that represent threats and approaching those that represent targets, and depositing their own pheromones to coordinate among themselves. (The distinction between threat and target may depend on the Blue entity in question: a SEAD resource would be attracted to SAM's that might repel other resources.) The emergence of paths depends on the interaction of a large number of Blue entities. If the population of physical resources is limited, a large population of software-only "ghost agents" swarms through the pheromone landscape to build up paths that the physical Blue agents then follow. Figure 6 shows repulsive and attractive Red pheromones, and the resulting co-field laid down by Blue ghost agents that forms a path for a strike package to follow. This mechanism can discriminate targets based on proximity or priority, and can plan sophisticated approaches to highly-protected targets, approaches that centralized optimizers are unable to derive.

Vehicle Coordination.—The algorithms developed in our path planning work were incorporated into a limited-objective experiment conducted by SMDC for J9 in 2001 [21, 48]. In this application, up to 100 UAV's coordinated their activities through digital pheromones. UAV's that had not detected a target deposited a pheromone that repelled other UAV's, thus ensuring distribution of the swarm over the battlespace. When a UAV detected a target, it deposited an attractive pheromone, drawing in nearby vehicles to join it in the attack. This capability enabled the deployment of many more vehicles without an increase in human oversight, and yielded significant improvements in performance over the baseline, including a 3x improvement in Red systems detected, a 9x improvement in the system exchange ratio, and a 11x improvement in the percentage of Red systems killed.

Pattern Recognition.—The Army's vision for the Future Combat System includes extensive use of networks of sensors deployed in the battlespace. Conventional exploitation of such a network pipes the data to a central location for processing, an architecture that imposes a high communication load, delays response, and offers adversaries a single point of vulnerability. We have demonstrated an alternative approach in which pattern recognition is distributed throughout the sensor network, enabling individual sensors to recognize when they are part of a larger pattern [15]. The swarming agents are not physical, but purely computational, and move between neighboring sensors using only local communications. Figure 7a shows an example distribution of sensors (a 70x70 grid). With a global view, we can quickly identify the sensors with high readings (plotted as white), but individual sensors do not have this perspective and cannot be sure whether they are high or low. One species of swarming agents compares each sensor's readings with a summary of what it has seen on other sensors to estimate whether the current sensor is exceptional, and deposits search pheromones (Figure 7b) to attract its colleagues to confirm its assessment. Each agent has seen a different subset of the other sensors, so a high accumulation of find pheromone on a sensor (Figure 7c) indicates that the sensor really is high in comparison with the rest of the network, and it can call for appropriate intervention. A second species of agents moves over the sensors both spatially and (through stored histories of recent measurements) chronologically. The movement of this species is not random, but embodies a spatio-temporal pattern, and its pheromone deposits highlight sensors that are related through this pattern (in Figure 7d, an orientation from SW to NE).

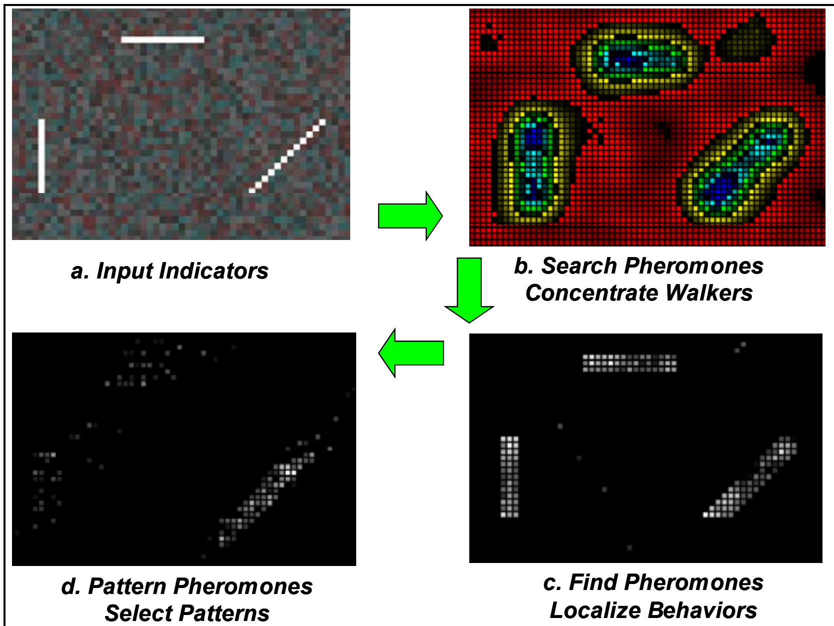

Figure 7: Multiple species of software agents swarming over a sensor network can enable the network to detect patterns without centralizing the data.

4 How can we Measure and Control Swarming?

The mechanisms outlined in the previous section can enable populations of software or hardware entities to self-organize through local interactions, but to be useful, human overseers must be able to measure their performance and control their actions. This section briefly discusses approaches to these important functions.

4.1 Measurement

We have defined swarming as "useful self-organization of multiple entities through local interactions." The terms in this definition offer a useful template for measuring the performance of a swarm. The criteria of "multiple entities" and "local interactions" identify independent variables that characterize the kind of swarm being considered, while the notion of "useful self-organization" leads to several dependent variables. Because of the nonlinearities involved in both individual agent behavior and the interactions among agents, the values of the dependent variables can change discontinuously as the independent variables are adjusted, and qualification of a swarm requires careful study of such "phase shifts." An example of such a study is [42].

Multiple entities.—Sometimes mechanisms that work satisfactorily for small numbers of entities do not scale well as the population increases. In other cases, there may be a critical minimum population below which the swarm will not function. (The latter condition motivated our use of "ghost agents" in the path planning example discussed above.) In evaluating swarms, it is crucial to study how the performance varies with population.

Local interactions.—Another set of variables under the direct control of the implementer of a swarm is the nature of local interactions among swarm members. This interaction may be varied along a number of dimensions, including mode (direct messaging, either point-to-point or broadcast, or sensing), range, and bandwidth.

Measures of Usefulness.—The measures used to assess the usefulness of a swarm are drawn directly from the MOE's and MOP's appropriate to the application problem. For example, in the LOE exploring the use of swarming UAV's for SEAD reported above, the MOE's were Percent of Red assets detected, Percent reduction of successful TBM launches, Percent of Red assets destroyed (by type), Percent of Blue assets destroyed, Time first Red units are detected, and System Exchange Ratio (SER). The MOP's were Number of Red targets detected by type and time, Number of Red targets nominated by type, Number of Red targets attacked by type, Number of Red targets destroyed by type, Number of Blue assets destroyed by type, and Number of TBM launches.

Measures of Self-Organization.—Some of the benefits of swarming are difficult to measure directly, but are directly correlated with the degree to which a swarm can organize itself. For example, directly assessing a swarm's robustness to unexpected perturbations would require a very large suite of experiments, but our confidence in this robustness can be strengthened if we can measure its self-organizing capabilities. We have found a variety of measures derived from statistical physics to be useful indicators of self-organization, including measures of entropy over the messages exchanged by agents, their spatial distribution, or the behavioral options open to them at any moment [13]. Frequently, local measures of these quantities permit us to deduce the global state of the swarm, a crucial capability for managing a distributed system [41]. It has recently been suggested that a Lebesgue measure of the portion of the swarm's space of behaviors that is dominated by the Pareto frontier might also be a useful measure of self-organization [25].

4.2 Control

The "self-organizing" aspect of a swarm implies that its global behavior emerges as it executes, and may vary in details from one run to the next because of changes in the environment. Detailed moment-by-moment control of the swarm would damp out this self-organization and sacrifice many of the benefits of swarming technology. However, swarming does not imply anarchy. Swarms can be controlled without sacrificing their power in two ways: by shaping the envelope of the swarm's emergent behavior, and by managing by exception.

Envelope Shaping.—While the details of a swarm's behavior may vary from one run to the next, those variations often are constrained to an envelope that depends on the configuration of the swarm. An illustration of this distinction can be seen in the Roessler attractor from chaos theory (Figure 8). This figure is a plot in three-dimensional phase space of a set of differential equations in their chaotic regime. The line that twists through this figure indicates the trajectory of this system, a trajectory that is so intertwined that arbitrarily small differences in initial conditions can lead to widely varying outcomes. For instance, if the system starts at location "A," it is in principle impossible to predict whether at a specified future time it will be at location B or location C. However, in spite of its detailed unpredictability, the system is confined to a highly structured envelope, and it is impossible for it to visit the point D.

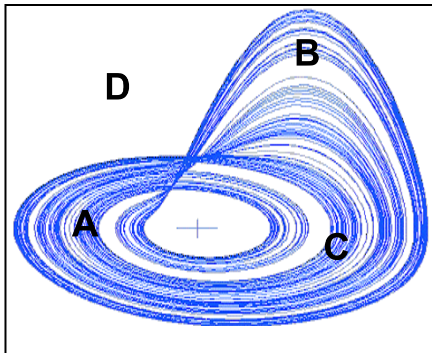

Figure 8: The Roessler Attractor

To shape a swarm's envelope, it is exercised in simulation, and human overseers evaluate its performance, rewarding appropriate behavior and punishing inappropriate behavior. Evolutionary or particle swarm methods then adjust the behaviors of individual swarm members so that desirable behavior increases and undesirable behavior decreases [10, 47]. The process adjusts the envelope of the system's behavior so that undesirable regions are avoided. Incidentally, these techniques enable swarms to be trained rather than designed, an approach that reduces the need for specialized software skills on the part of the warfighter. Evolution can also be used to explore the behavioral space of a swarm in much greater detail than exhaustive simulation would permit, by selectively altering later simulation runs based on the results of earlier ones [14].

Managing by Exception.—Once a swarm has been launched, human overseers can observe its emerging behavior and intervene on an exception basis. For example, a swarm with kill capability can autonomously detect a target and configure itself for attack, then apply for human permission to execute. Digital pheromones are especially amenable to human direction. Graphic marks on a map can be translated directly into pheromone deposits that modify the emergent behavior of the swarm in real-time (Figure 9). A path being formed by the system can be blocked or a whole region excluded; the priority of individual targets and threats can be adjusted; segments of paths can be explicitly designated; and bounds can be placed on performance metrics. The important point is that human intervention is on an exception basis. Routine operation proceeds without detailed human control, freeing human warfighters to concentrate on more strategic concerns and calling their attention to situations where their judgment is required.

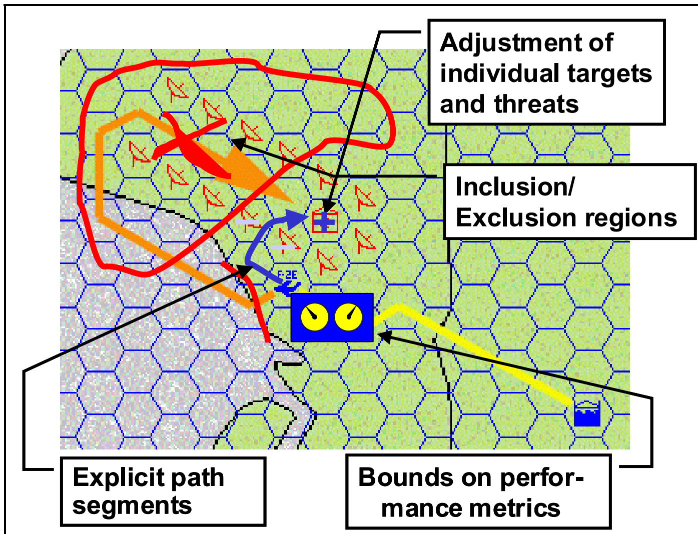

Figure 9: Human Direction of a Digital Pheromone System

5 Conclusion

Swarming is an ancient military vision that has been emulated by human warriors for centuries. Recent growth in our understanding of the mechanisms underlying natural swarming, and advances in information technology, enable us to construct, configure, and deploy swarms of unmanned systems.

These systems are useful not only for the attack scenarios that have dominated history, but also for planning, communications management, and sensor processing, among other applications. These systems are measurable and controllable, and offer an important opportunity for military technologists to increase adaptability, robustness, and lethality, while reducing human manpower requirements and risk.

Acknowledgements

Portions of this work are supported by Project Alpha in J9 Joint Experimentation, under the direction of Russell Richards and Gary Trinkle, through a contract with Rome Laboratories under Carl DeFranco. Some of the underlying technologies were developed under various projects sponsored by DARPA. The opinions expressed in this paper are solely those of the author, and do not necessarily reflect the views of the sponsoring agencies or the US Department of Defense. I am grateful to my colleagues at Altarum Institute for their contributions to the work on digital pheromones, and to Frank Weiskopf and David Scheidt of JHU APL and John Feddema of Sandia National Laboratories for clarifying their work on various swarming architectures.

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] J. Arquilla and D. Ronfeldt. Swarming and the Future of Conflict. DB-311, RAND, Santa Monica, CA, 2000. URL http://www.rand.org/publications/DB/DB311.

[3] P. Ball. *The Self-Made Tapestry: Pattern Formation in Nature*. Princeton, NJ, Princeton University Press, 1996.

[4] G. Beni. The Concept of Cellular Robotic System. In *Proceedings of IEEE Int. Symp. on Intelligent Control*, Los Alamitos, CA, pages 57-62, IEEE Computer Society Press, 1988, 1988.

[5] G. Beni and S. Hackwood. Stationary Waves in Cyclic Swarms. In *Proceedings of IEEE Int. Symp. on Intelligent Control*, Los Alamitos, CA, pages 234-242, IEEE Computer Society Press, 1992, 1992.

[6] G. Beni and J. Wang. Swarm Intelligence. In *Proceedings of Seventh Annual Meeting of the Robotics Society of Japan*, Tokyo, pages 425-428, RSJ Press, 1989, 1989.

[7] G. Beni and J. Wang. Theoretical Problems for the Realization of Distributed Robotic Systems. In *Proceedings of IEEE International Conference on Robotic and Automation*, Los Alamitos, CA, pages 1914-1920, IEEE Computer Society Press, 1991, 1991.

[8] E. Bonabeau. Swarm Intelligence. In *Proceedings of Swarming: Network Enabled C4ISR*, Tysons Corner, VA, ASD C3I, 2003.

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

[10] L. Booker. Learning Tactics for Swarming Entities. In *Proceedings of Swarming: Network Enabled C4ISR*, Tysons Corner, VA, ASD C3I, 2003.

[11] R. A. Brooks. A Robust Layered Control System for a Mobile Robot. *IEEE Journal of Robotics and Automation*, RA-2(1 (March)):14-23, 1986.

[12] 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.

[13] S. Brueckner and H. V. D. Parunak. Information-Driven Phase Changes in Multi-Agent Coordination. In *Proceedings of Autonomous Agents and Multi-Agent Systems (AAMAS 2003)*, Melbourne, Australia, pages (forthcoming), 2003. URL http://www.erim.org/~vparunak/AAMAS03InfoPhaseChange.pdf.

[14] S. Brueckner and H. V. D. Parunak. Resource-Aware Exploration of Emergent Dynamics of Simulated Systems. In *Proceedings of Autonomous Agents and Multi-Agent Systems (AAMAS 2003)*, Melbourne, Australia, pages (forthcoming), 2003. URL http://www.erim.org/~vparunak/AAMAS03APSE.pdf.

[15] S. A. Brueckner and H. V. D. Parunak. Swarming Agents for Distributed Pattern Detection and Classification. In *Proceedings of Workshop on Ubiquitous Computing, AAMAS 2002*, Bologna, Italy, pages (forthcoming), 2002. URL http://www.erim.org/~vparunak/PatternDetection01.pdf.

[16] S. Camazine, J.-L. Deneubourg, N. R. Franks, J. Sneyd, G. Theraulaz, and E. Bonabeau. *Self-Organization in Biological Systems*. Princeton, NJ, Princeton University Press, 2001.

[17] B. Clough. Emergent Behavior (Swarming): Tool Kit for Building UAV Autonomy. In *Proceedings of Swarming: Network Enabled C4ISR*, Tysons Corner, VA, ASD C3I, 2003.

[18] C. Croom, Jr. C4ISR Warfighter Integration. In *Proceedings of Swarming: Network Enabled C4ISR*, Tysons Corner, VA, ASD C3I, 2003.

[19] R. Davis and R. G. Smith. Negotiation as a metaphor for distributed problem solving. *Artificial Intelligence*, 20:63-109, 1983.

[20] M. B. Dias and A. Stentz. A Free Market Architecture for Distributed Control of a Multi-robot System. In *Proceedings of The 6th International Conference on Intelligent Autonomous Systems (IAS)*, Venice, Italy, 2000. URL http://www.frc.ri.cmu.edu/projects/colony/DiasStentz.ps.

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

[22] S. J. A. Edwards. Swarming on the Battlefield: past, Present, and Future. MR-1100-OSD, RAND, Santa Monica, CA, 2000.

[23] S. J. A. Edwards. Military History of Swarming. In *Proceedings of Swarming: Network Enabled C4ISR*, Tysons Corner, VA, ASD C3I, 2003.

[24] J. T. Feddema, R. Robinett, D. Schoenwald, C. Lewis, J. Wagner, and E. Parker. Analysis and Control for Distributed Cooperative Systems. *Presentation at Swarming Entities – Joint C4ISR DSC Study Plan Conference*, Johns Hopkins University Applied Physics Laboratory, Laurel, MD, 2002.

[25] M. Fleischer. Foundations of Swarm Intelligence: From Principles to Practice. In *Proceedings of Swarming: Network Enabled C4ISR*, Tysons Corner, VA, ASD C3I, 2003.

[26] D. Frelinger, J. Kvitky, and W. Stanley. Proliferated Autonomous Weapons: An Example of Cooperative Behavior. DB-239-AF, RAND Corporation, 1998.

[27] M. Gerla. Swarm Communications in the ONR-Minuteman Project. In *Proceedings of Swarming: Network Enabled C4ISR*, Tysons Corner, VA, ASD C3I, 2003.

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

[29] P.-P. Grassé. 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*, 6:41-84, 1959.

[30] S. Hackwood and G. Beni. Self-Organizing Sensors by Deterministic Annealing. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robot and Systems*, Los Alamitos, CA, pages 1177-1183, IEEE Computer Society Press, 1991, 1991.

[31] S. Hackwood and G. Beni. Self-organization of Sensors for Swarm Intelligence. In *Proceedings of IEEE Int. Conf. on Robotics and Automation*, pages 819-29, 1992.

[32] E. Hornung and B. M. Bryan, Editors. *The Quest for Immortality: Treasures of Ancient Egypt*. Washington, DC, National Gallery of Art, 2002.

[33] D. Inbody. Swarming: Historical Observations and Conclusions. In *Proceedings of Swarming: Network Enabled C4ISR*, Tysons Corner, VA, ASD C3I, 2003.

[34] C. R. Kube and H. Zhang. Collective Robotics: From Social Insects to Robots. *Adaptive Behavior*, 2(2):189-219, 1993. URL ftp://ftp.cs.ualberta.ca/pub/kube/ab93.pdf.

[35] 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.

[36] L. Leonardi, M. Mamei, and F. Zambonelli. Co-Fields: Towards a Unifying Model for Swarm Intelligence. DISMI-UNIMO-3-2002, University of Modena and Reggio Emilia, Modena, Italy, 2002. URL http://polaris.ing.unimo.it/didattica/curriculum/marco/Web-Co-Fields/stuff/Swarm.pdf.

[37] M. Mamei, F. Zambonelli, and L. Leonardi. A Physically Grounded Approach to Coordinate Movements in a Team. In *Proceedings of First International Workshop on Mobile Teamwork (at ICDCS)*, Vienna, Austria, IEEE CS Press, 2002. URL http://polaris.ing.unimo.it/Zambonelli/PDF/Teamwork.pdf.

[38] M. Mamei, F. Zambonelli, and L. Leonardi. Distributed Motion Coordination with Co-Fields: A Case Study in Urban Traffic Management. In *Proceedings of 6th IEEE Symposium on Autonomous Decentralized Systems (ISADS 2003)*, Pisa, Italy, IEEE CS Press, 2003. URL http://polaris.ing.unimo.it/Zambonelli/PDF/isads.pdf.

[39] E. Neufeld. Insects as Warfare Agents in the Ancient Near East. *Orientalia*, 49(1):30-57, 1980.

[40] 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.

[41] H. V. D. Parunak, S. Brueckner, R. Matthews, and J. Sauter. How to Calm Hyperactive Agents. In *Proceedings of Autonomous Agents and Multi-Agent Systems (AAMAS 2003)*, Melbourne, Australia, pages (forthcoming), 2003. URL http://www.erim.org/~vparunak/AAMAS03Ritalin.pdf.

[42] H. V. D. Parunak, S. Brueckner, J. Sauter, and R. Savit. Effort Profiles in Multi-Agent Resource Allocation. In *Proceedings of Autonomous Agents and Multi-Agent Systems (AAMAS02)*, Bologna, Italy, pages 248-255, 2002. URL www.erim.org/~vparunak/AAMAS02Effort.pdf.

[43] H. V. D. Parunak, M. Purcell, and R. O'Connell. Digital Pheromones for Autonomous Coordination of Swarming UAV's. In *Proceedings of First AIAA Unmanned Aerospace Vehicles, Systems,Technologies, and Operations Conference*, Norfolk, VA, AIAA, 2002. URL www.erim.org/~vparunak/AIAA02.pdf.

[44] K. Pister. Smart Dust: Autonomous sensing and communication in a cubic millimeter. 2001. Web Page, http://robotics.eecs.berkeley.edu/~pister/SmartDust/.

[45] J. M. Riggs. The Objective Force. In *Proceedings of Swarming: Network Enabled C4ISR*, Tysons Corner, VA, ASD C3I, 2003.

[46] 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.

[47] J. A. Sauter, R. Matthews, H. V. D. Parunak, and S. Brueckner. Evolving Adaptive Pheromone Path Planning Mechanisms. In *Proceedings of Autonomous Agents and Multi-Agent Systems (AAMAS02)*, Bologna, Italy, pages 434-440, 2002. URL www.erim.org/~vparunak/AAMAS02Evolution.pdf.

[48] 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. URL https://home.je.jfcom.mil/QuickPlace/experimentation/PageLibrary85256AB1003BBEA7.nsf/h_0036FB98FFD2ACCA85256AB2004161B0/D7680995272C266B85256B20004E1BF0/?OpenDocument.

[49] S. Thayer, B. Digney, M. B. Dias, A. Stentz, B. Nabbe, and M. Hebert. Distributed Robotic Mapping of Extreme Environments. In *Proceedings of SPIE: Mobile Robots XV and Telemanipulator and Telepresence Technologies VII*, 2000. URL http://www.frc.ri.cmu.edu/projects/colony/pdf_docs/thayer_scott_2000_1.pdf.

[50] F. Weiskopf, T. Gion, D. Elkiss, H. Gilreath, J. Bruzek, and R. Bamberger. Control of Cooperative, Autonomous Unmanned Aerial Vehicles. In *Proceedings of First AIAA Technical Conference and Workshop on UAV, Systems, Technologies, and Operations*, pages AIAA Paper 2002-3444, 2002.

[51] F. Weiskopf and D. Scheidt. Cooperative Autonomous UAV Team. *Presentation at Swarming Entities – Joint C4ISR DSC Study Plan Conference*, Johns Hopkins University Applied Physics Laboratory, Laurel, MD, 2002.

[52] M. Wooldridge. *An Introduction to MultiAgent Systems*. Chichester, UK, John Wiley, 2002.