Defence Research and Development Canada
Recherche et développement pour la défense Canada

DEFENCE R&D DÉFENSE

Expert Assessment of Human-Human Stigmergy

H. Van Dyke Parunak
Altarum Institute

Defence R&D Canada
CONTRACT REPORT
DRDC CR 2005-003
October 2005

Canada

Expert Assessment of Human-Human Stigmergy

July 11, 2005

H. Van Dyke Parunak, Ph.D.
Altarum Institute

Altarum
3520 Green Court, Suite 300 • Ann Arbor, Michigan 48105 • (734) 302-4600

Corporate Headquarters	Alexandria, VA	San Antonio, TX
734 • 302 • 4600	703 • 575 • 1200	210 • 832 • 3000

Expert Assessment of Human-Human Stigmergy

July 11, 2005

H. Van Dyke Parunak, Altarum Institute

1 Executive Summary

Human-Human Stigmergy is pervasive. A wide range of pre-computer social systems fit the pattern of stigmergic coordination, and have provided a rich set of metaphors on which a diverse set of computer-enabled systems for enabling human stigmergy have been constructed. It would be more difficult to show a functioning human institution that is not stigmergic, than it is to find examples of human stigmergy.

The reason that human-human stigmergy is so common can be understood from the growing body of experience in constructing large-scale distributed computing systems with resource-constrained elements. It has become clear that central control of such systems is not feasible, since resource-constrained components cannot cope with the large-scale, distributed aspects of such systems. The central insight of stigmergy is that coordination can be achieved by resource-constrained agents interacting locally in an environment. Two fundamental principles govern the success of this strategy.

1. No matter how large the environment grows, because agents interact only locally, their limited processing capabilities are not overwhelmed.
2. Through the dynamics of self-organization, local interactions can yield a coherent system-level outcome that provides the required control.

The essence of stigmergy is the coordination of bounded agents embedded in a (potentially unbounded) environment, whose state they both sense (to guide their actions) and modify (as a result of their actions). The ability of stigmergic systems to scale without overwhelming individual agents requires that the environment offer a topology in which agents are situated locally, and within which their actions and senses have a limited horizon. Previous studies of stigmergic coordination, inspired by insects and applied to software and robotic systems, have identified four classes of stigmergic interaction between the agent and the environment [67]. Stigmergy can be either marker-based or sematectonic, and (independently) either quantitative or qualitative. All four of these can be identified in human-human stigmergy.

Against this theoretical background, we classify a wide range of human stigmergic mechanisms, both pre-computational and computational. This survey shows that applications in the civilian world are both common and fairly mature. However, extending their effectiveness and transitioning them to military applications will require advances in several categories of enabling technology, including ubiquitous computing hardware and foundations, integrated social and distributed systems simulation technology, and technology for security and privacy.

As a way ahead for the Canadian Defense Organization, we recommend leveraging the wide range of commercial systems in order to deliver short-term benefits to the military, but casting these transition efforts in a framework in which each supports the exploration of one or more critical enabling technologies.

2 Technical Description of the Subject Area

"Stigmergy" is a term coined in the 1950's by the French entomologist Grassé [33] 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 notion that an agent's actions leave signs in the environment, signs that it and other agents sense and that determine their subsequent actions.

In spite of Grassé's professional preoccupation with insects, stigmergy is ubiquitous in the interactions of humans with one another. Our experience with a wide range of distributed systems suggests that it is the only way for members of a large distributed population to coordinate themselves with bounded computational resources. This assessment will summarize instances of stigmergic inter-human coordination both before and after the advent of the computer.

To set the context for this survey, we outline the basic architecture of stigmergy, then develop a taxonomy of stigmergic interactions that can be used to classify specific instances.

2.1 Architecture of Stigmergy

Figure 1 summarizes the basic components of a stigmergic system and their inrerrelations with one another.

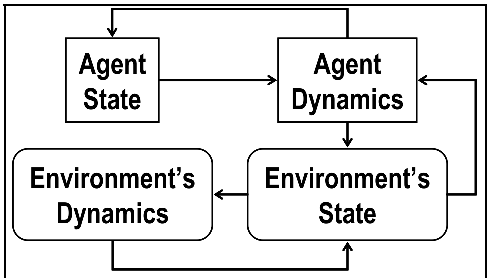

Figure 1: Basic Architecture of Stigmergy

The components of a stigmergic system are a population of agents, and an environment in which they are immersed.

Each agent has

- an internal state, which generally *is not* directly visible to other agents;
- sensors that give it access to some of the environment's state variables;
- actuators that enable it to change some of the environment's state variables;
- a program (its "dynamics") that maps from its current internal state and the readings of its sensors to changes in its internal state and commands given to its sensors and actuators.

The environment has

- a state, certain aspects of which generally *are* visible to the agents;
- a program (its "dynamics") that governs the evolution of its state over time.

The most important distinction between agents and the environment is that the internal state of agents is hidden, while the state of the environment is accessible to an agent with appropriate sensors. In most cases, a second distinction can be observed. Each agent is monolithic, a self-contained computational object with a well-defined boundary. Typically, the environment is not not monolithic, but is structured according to some topology. Some examples of environmental topologies include

- a Cartesian space (e.g., the surface of the earth);

- a graph structure (e.g., a telecommunications network or social organization);
- a list of disjoint categories (e.g., a list of topics, though these are usually organized into a graph by imposing an ontology).

When the environment is structured in this way, each agent is localized in the environment. That is, its sensors and actuators are confined to one region of the environment. If the agent is mobile, it can change location in the environment, but at any moment it is at one location. This localization of agents within the environment restricts the computational load imposed on the agents, and enables stigmergic systems to scale without exceeding the load on each agent.

While localization in a distributed environment keeps the computational load on each agent manageable, it does not ensure that a reasonable system-level behavior will emerge from the interactions of the agents. Critical support for this objective lies in the interaction of the dynamics of the agents with those of the environment. These dynamics are typically nonlinear individually, and their interactions are often nonlinear as well, resulting in a system that is susceptible to formal chaos. Far from being a disadvantage, such dynamics actually enable self-organization, since they permit the system to explore its state space efficiently. This exploration is a key ingredient of self-organization, other components of which are discussed in [75].

2.2 Varieties of Stigmergic Interaction

Four different varieties of stigmergic interaction between agent and environment 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”) [11]. 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”) [14]. As shown in Table 1, the two distinctions are orthogonal.

- The paradigm for marker based stigmergy is the use of pheromones by certain social insects to coordinate their actions. Most insect species use a few dozen distinct pheromone “flavors,” and thus use qualitative as well as quantitative decision-making. In engineered systems, stigmergic markers can consist of actual physical chemicals deposited in a physical landscape, labeled scalar variables stored in a data structure whose topology reflects that of the problem (as in much of our UAV control work), or price signals in a marketplace [19]. The latter metaphor is particularly important in coordinating human interactions, as in our RAPPID system for collaborative design [78-80].
- In several insect behaviors, the structure of the domain itself provides sufficient signals for coordinating behavior, without the need for special markers. Ants cluster corpses in their cemeteries, using only the density of corpse distribution as a basis for their decisions

Table 1: 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

[9]. This is a quantitative decision, depending only on the distribution of a single type of object. Wasps decide where to add the next cell to their nests based on which of several templates best characterizes the current local shape of the nest, thus making a qualitative distinction. Sematectonic stigmergy is also illustrated by an algorithm that explains how wolves surround their prey [46], by being attracted to the prey while repelled by neighboring wolves. We have applied sematectonic stigmergy to coordination of multiple sensors [72, 76] and the assembly of intelligence information [97].

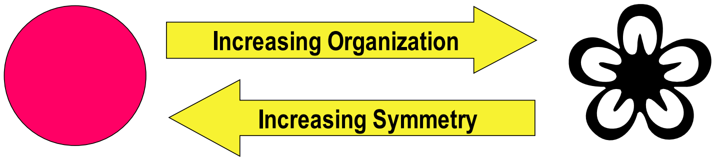

Figure 2: Symmetry vs. Organization.—Achieving high levels of organization requires breaking the symmetry among system components, a function that environmental noise supports.

A subset of stigmergic mechanisms, known as “coordination fields” or “co-fields” [49, 54, 55, 98], 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 [84].

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.

2.3 Characteristics of Stigmergy

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” [83]).
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 2), 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 [6]. 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.

2.4 Engineering Stigmergy

Two principles govern the engineering of stigmergic systems: their intimate relation with the environment, and the emergent nature of the system-level behavior to which they lead.

The Environment.—Stigmergy consists of coordination via interactions through a shared environment. The structure and dynamics of the environment offer two important mechanisms that engineers can manipulate in constructing a stigmergic system.

- The *structure* of the environment provides a sense of locality to agents that interact through it. Stigmergic agents are situated somewhere in the environment, and typically respond most immediately to other agents that are nearby in the topology imposed by the environment. As a result of this locality of interaction, the size of the overall system can grow considerably by extending the environment and adding more agents, without increasing the complexity or bandwidth of individual agents.
- The *dynamics* of the environment permit it to perform certain information processing tasks on behalf of the system. For example, in pheromone systems, the aggregation of pheromone deposits from different agents at a single location in the environment offers a form of information fusion; the diffusion of pheromones to nearby locations provides information sharing; and the evaporation of pheromones over time provides a highly efficient approach to truth maintenance, discarding obsolete information.

The stigmergic environment has just begun to emerge as the focus of research in its own right [100].

Emergent Behavior.—Stigmergic systems are often described as "self-organizing" and exhibiting "emergent behavior." We distinguish these two phenomena based on the difference between the horizontal concept of system boundary and the vertical concept of levels (Figure 3).

We define **Self-Organization** as organization Among elements *within a level*

- *Without information flow* across the boundary.

The second law of thermodynamics demands that there be *energy flow* across the boundary of any system whose organization increases over time. *Self*-organization requires that this energy flow not contain information. This definition depends critically on the location of the

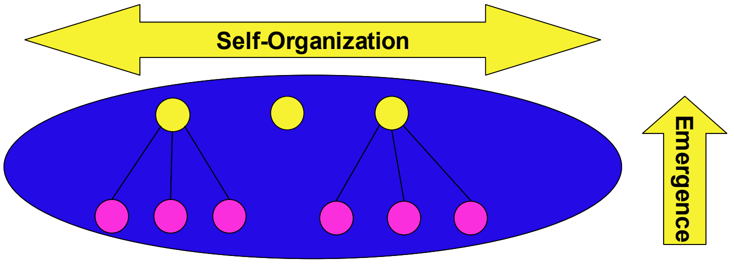

Figure 3: Comparing Self-Organization and Emergence

system boundary. If the boundary is moved, a system's character as self-organizing or not may change.

We define **Emergence** as a subcategory of self-organization (Figure 4). Emergence describes the appearance of structures at a *higher level* that are not explicitly represented in *lower-level* components. The reliance of stigmergic systems on locally available information makes it difficult for them to reason explicitly about higher-level structures, so emergence tends to be an important mechanism in swarming systems.

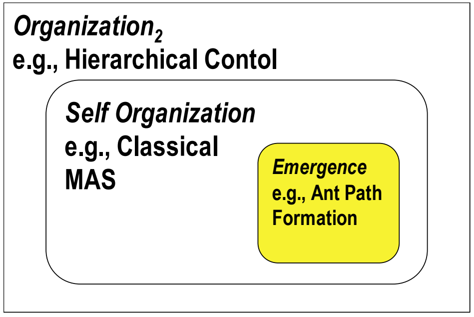

Figure 4: Emergence as a Subcategory of Self-Organization

Neither self-organization nor emergence is necessarily good. The formation of structures will correspond to a reduction in entropy, whether those structures support or frustrate the objectives of the system stakeholders. The fact that emergent structures can be pathological (as in the case of race conditions or herding behavior) may explain the apprehension with which some people view emergence. For example, Wooldridge and Jennings assert [105], "Emergent functionality is akin to chaos …." They urge engineers of agent systems to "severely restrict the way in which agents can interact with one-another … ensure that there are few channels of communication between agents … restrict the *way* in which agents interact" in order to reduce the likelihood of emergent behavior. A consequence of this restriction is that any desired system-level behavior must be explicitly represented in the lower-level components, a requirement that is difficult to meet if the system's requirements include responding gracefully to unanticipated changes in its environment. Our alternative approach is to develop principles for designing and developing systems whose emergent behavior is beneficial or at least benign.

This difference in vision leads to two distinct approaches to building multi-agent systems (Figure 5).

- Classical multi-agent systems achieve self-organization through deliberation among fairly sophisticated ("coarse-grained") agents. Emergent systems can use much simpler reactive agents.
- Reasoning in emergent self-organization is often non-symbolic, while classical systems are usually symbolic.
- Because of the need for representing system-level behavior explicitly at all layers, non-emergent systems are best suited for exploiting well-known environments. The ability of emergent systems to produce new behaviors is appropriate for more exploratory problems.

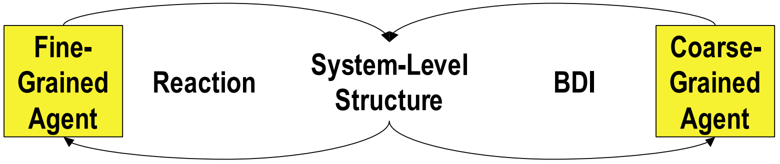

Emergence	Deliberation
Non-Symbolic	Symbolic
Exploration	Exploitation

Figure 5: Two Families of Self-Organizing Systems

A front line of our current research is understanding how to hybridize these two families of systems.

In developing a disciplined approach to self-organization and emergent behavior, it is helpful to draw on the antecedent of statistical physics, which offers a mature, mathematically precise account of how the properties of matter at the human scale emerge from interactions among much smaller units (atoms and molecules). We have been successful in exploiting a number of concepts from statistical physics in understanding, designing, and controlling stigmergic systems. These concepts include entropy [5, 37, 71], universality [74], phase shifts [73], and formal analysis of overall system convergence [77].

Environment:
- Topology:
- State:
- Dynamics:

Agents:
- Sensor:
- Actuator:
- Dynamics:

Emergent system behavior:

Figure 6: Template for analysis of examples

3 Present Status of Technology

We summarize instances of stigmergic human coordination in two broad categories: those that do not rely on digital computers (though they may be enhanced by them), and those that are distinctively creatures of the computer age. For each category, we document an example, and analyze it in terms of the theoretical categories developed in the previous section, using the template shown in Figure 6. We first identify the environment, describe its topology, the state variables that it supports, and any internal dynamics. Then we identify the agents, discuss how they sense and modify the environment (with attention to the main distinction between sema(tectonic) and marker(-based) stigmergic interaction), and summarize their internal dynamics. Finally, we describe the overall emergent system behavior that the stigmergy achieves.

Environment: Vegetated terrain
- Topology: 2D manifold
- State:
    - Degree of ground cover
    - Obstacles
- Dynamics:
    - Trodden vegetation dies
    - Vegetation regrows on untrodden areas

Agents: People (pedestrians or in vehicles)
- Sensor:
    - Sema: smoothness to path
    - Sema: direction to destination
    - Marker: road signs
- Actuator:
    - Sema: direction of next step
    - Marker: pave the path
    - Marker: set road signs
- Dynamics: optimize smoothness and direction

Emergent system behavior: globally marked paths

Figure 7: Trail Formation

3.1 Pre-Computational

Humans have long coordinated their activities through non-computational environments, though in many cases these mechanisms can be enhanced with the use of computers.

3.1.1 Movement Coordination

Humans have always needed to move about in their environment, and have drawn on stigmergic mechanisms to coordinate both the trails along which they travel and the moment-by-moment decisions as to which trail to choose.

3.1.1.1 Trail Formation

The simplest and most primitive trail formation mechanisms rely on sematectonic stigmergy, in which humans wear down the vegetation on frequently-traveled routes, and grass regrows if an old path is not used (Figure 7). There is an extensive literature on mathematical models for such path formation by "active walkers" [39].

While trails can form entirely with sematectonic stigmergy, humans tend to enhance them with markers. My college campus was notorious for laying pavement along bare tracks through the grass, turning emergent trails into permanent ones. A common modern example is traffic signs, often maintained locally by a variety of jurisdictions, but Native Americans also used artificial markers, including unnaturally bent trees [25] and petroglyphs [57], to mark trails. While such mechanisms can make trails easier to follow, they also render them less dynamic. It takes longer for a concrete path to crumble than for grass to regrow over an unused dirt path.

Environment: Trail network
- Topology: Graph
- State:
  - Congestion
  - Toll fees
- Dynamics: Convey traffic across edges of the graph

Agents: People (pedestrians or in vehicles)
- Sensor:
  - Sema: current congestion on edge
  - Sema: direction to destination
  - Marker: requested toll
- Actuator:
  - Sema: choice of edges at each node
  - Marker: pay toll
- Dynamics: optimize speed and economy

Emergent system behavior: more balanced load on different paths

Figure 8: Traffic Flow

3.1.1.2 Traffic Flow

Humans not only generate trails stigmergically, but also allocate their movement across alternative routes, both sematectonically and using markers (tolls) (Figure 8). Note that this behavior builds on the product of a previous stigmergic activity (the formation of the trails themselves).

3.1.2 Market Systems

Adam Smith's "invisible hand" is an example of the self-organizing potential of a stigmergic system, and have the benefit of an immense body of formal study [94].

Two varieties of open markets have been identified, and both can be described as stigmergic, though at different levels.

In an auction-based market (such as a stock market or commodities exchange), the marketplace provides an environment that integrates the bids of individual buyers and sellers interact, using currency as markers (Figure 9). The markets for individual products are distinct, so the topology of the environment is strictly

Environment: Product-specific exchanges
- Topology: Categories (though often linked by product dependencies, e.g., steel in autos)
- State: Current bid and ask prices
- Dynamics: integrate offers to compute prices that clear the market

Agents: Buyers and sellers
- Sensor (Marker): Current prices
- Actuator (Marker): State own bid or ask price
- Dynamics: Maximize time integral of revenue over expense

Emergent system behavior: globally optimum allocation of resources

Figure 9: Auction Market Systems

speaking a set of disjoint categories. However, different products are linked by other processes in the economy (for example, the markets for steel and automobiles), leading to an implicit graph structure.

A less-studied alternative to auction systems is the pattern of Edgeworth barter, in which buyers and sellers interact directly with each other in repeated pairwise transactions, without the benefit of an auctioneer [3] (Figure 10). In this system, the environment consists entirely of the graph of dependencies among different products induced by patterns of joint use, since it is through these dependencies that individual transactions have an effect on one another.

Environment: Pairwise encounters
- Topology: Spatial distribution of traders
- State: Current locations of traders
- Dynamics: Support mixing of traders

Agents: Traders
- Sensor (Marker): Bid and ask price for single entity
- Actuator (Marker): State own ask or bid price
- Dynamics: Integrate information over successive exchanges

Emergent system behavior: In repeated trades, balances supply and demand

Figure 10: Edgeworth Barter Markets

3.1.3 Elections

An election can be viewed as a market in which candidates and issues are the commodities and votes correspond to currency. In a single-issue or single-office election, the dynamics are quite simple. However, many elections involve a series of issues or offices, often linked to one another through political platforms. The ultimate outcome in terms of governance depends on the policies advocated by a candidate. A voter may favor some of those policies and oppose others, but rarely has the opportunity of voting for or against individual policies. Voting for a candidate can be compared to buying a complex product with multiple attributes. For example, in an automobile, one may want high fuel economy, good off-road performance, and low maintenance, but usually must make compromises.

If an election can be compared to a market, representative government corresponds to an economy. Elected officials themselves participate in numerous legislative actions, including both explicit votes and implicit agreements, to pursue their platforms. The electorate often chooses a set of representatives in anticipation of the subsequent legislative give-and-take, providing for a balance of power in the overall structure.

Viewed in this way, the topography within which elections take place is a graph with different colors of nodes and edges. Some nodes represent candidates, others represent policies, and still others represent representative divisions (such as geographical regions). Edges link candidates to their regions and the issues they support or oppose. Votes for candidates propagate to the issues associated with them (Figure 11).

Environment: Network of districts, candidates, and issues
- Topology: Graph
- State: Connectivity of links
- Dynamics: Population of districts; support for candidates and issues

Agents: Voters
- Sensor (Sema): Affiliation of candidates with issues
- Actuator (Marker): Vote
- Dynamics: Choose policies indirectly through candidates

Emergent system behavior: Set of policies aligned with voter interests

Figure 11: Elections

3.1.4 Document Editing

Joint authorship has always been a stigmergic activity, mediated by the emerging document itself. Each author is stimulated by what previous authors have written to add main-line content or marginal comments. The dynamics of this process have been greatly enhanced by sophisticated word processing software that includes specific facilities for review, comment, and tracking multiple authors.

While a document may seem to be a static entity, internal semantic relations can change as a result of individual modifications. As a trivial example, consider a document on a sensitive issue. A later reviser adds a tendentious definition to the first page of the document. The structural integrity of the document has the effect of propagating the semantics of this definition to later sections, potentially changing their meaning (Figure 12).

Environment: A document
- Topology: Linear or (if structured) hierarchical
- State: Current content, both mainline and marginal comments
- Dynamics: Internal semantic propagation

Agents: Writers and editors
- Sensor (Sema): Current state of the document
- Actuator:
  - Sema: New content
  - Marker: Strike-outs, highlighting
- Dynamics: Adjust content to modulate ambiguity, tune an argument, or advocate a particular position

Emergent system behavior: Expression of jointly held consciousness

Figure 12: Document Editing

An extension of document editing is the development of knowledge in a scientific community. Each paper that is published contributes knowledge that other researchers can use in stimulating their ideas, modifying their research directions, and deriving their results (not to mention the more mundane results of achieving tenure for the authors and thus prolonging their ability to contribute to the field) [60].

3.1.5 Status Boards

Many social settings use a publicly visible display to coordinate activity (Figure 13). Examples include “in-out” boards indicating which staff are currently in the office, situation boards used in military battle management, and a wide range of bulletin boards advertising items for sale, employment opportunities, services offered, etc. The emergent behavior mediated by the board depends on its theme. “For sale” boards generate market encounters, usually of the Edgeworth variety. Situation boards enable tactical coordination of military forces. “In-out” boards enable more effective direction of inquiries and assignments to staff who are immediately available.

Environment: Modifiable public display
- Topology: 2-dimensional surface
- State: Current contents
- Dynamics: Most current material obscures or replaces older material

Agents: People visiting the board
- Sensor (Sema): Contents of the board
- Actuator (Sema): New postings
- Dynamics: Add new postings, remove or regroup old ones

Emergent system behavior: Depends on theme of the board (see text)

Figure 13: Status Board

3.1.6 Viral Marketing

Viral marketing describes “any strategy that encourages individuals to pass on a marketing message to others” [104]. Before the advent of computers, this form of communication was known as “word-of-mouth,” and in noncommercial venues was known as “rumors.” The speed of digital communication has made it particularly powerful. One classic example is free email such as Hotmail: every message sent by a user includes a system-generated tagline that encourages recipients to get their own Hotmail account, thus opening themselves to follow-on advertisements. Another, less blatant example is the use of planted participants in chat rooms to generate “buzz” in favor of a new music album, film, or book.

The heart of these strategies is the propagation of a message along the social network of participants (Figure 14), analogous to the propagation of a disease (thus the name). Like disease propagation, viral marketing depends on several critical factors [81], including

- the susceptibility of the members of the network to the message (if the message is intrinsically uninteresting, it will die out);
- the connectivity of the infected individuals (if the product of connectivity times susceptibility falls below 1, the infection will die out);
- the structure of the social network (an infection will die out in a lattice if the percentage of infected individuals falls below a certain threshold, but can persist in a power-law network no matter how small the percentage of infected individuals).

Environment: Social network
- Topology: Graph
- State: Connectivity of participants
- Dynamics: Communication

Agents: Members of network
- Sensor (Sema): Content heard from others
- Actuator (Sema): Repeating content
- Dynamics: Spread of message

Emergent system behavior: Number of individuals who have heard the message increases rapidly.

Figure 14: Viral Marketing

3.2 Computational

The advent of information technology has extended the applicability of stigmergic mechanisms for human coordination, by augmenting human abilities for sensing, communication, and information processing. Each of the examples in this section can be viewed as a descendant of one or more of the pre-computational examples discussed in the previous section.

3.2.1 Intelligent Transportation Systems

The application of computer technology to human movement coordination (Section 3.1.1) has produced the burgeoning field of “Intelligent Transportation Systems” (ITS, also known as “Intelligent Highway Systems,” IHS), with extensive government attention [24, 90],

Environment: Highway network
- Topology: Graph
- State:
  - Locations and velocities of vehicles
  - Timing of control signals
- Dynamics: Enable movement of vehicles from one place to another

Agents: Vehicles
- Sensor:
  - Marker: Signals
  - Sema: Local congestion
- Actuator (Sema): Route choices
- Dynamics: Obey signals; minimize local congestion; make progress toward destination

Emergent system behavior: Increased throughput, reduced collisions

Figure 15: Intelligent Highway System

independent business groups [44, 45], and dedicated research societies and journals [42].

While the field as a whole includes all modes of transportation, our analysis (Figure 15) focuses on highway systems. Computational mechanisms in this domain rely on the existence of road networks that in most cases were established using pre-digital stigmergy. Computational enhancements include

- roadbed sensors for real-time estimates of traffic density and velocity;
- improved signal systems (including not only traffic lights but also digital signs) for providing feedback to motorists;
- advanced algorithms [47, 106] for controlling signals on the basis of sensed traffic.

3.2.2 Collaboration Environments

Recognition of the potential of stigmergy for promoting coordination in human organizations [21] has led to a proliferation of systems that support human collaboration in one way or another. It is helpful to discuss these under two broad headings: content (technologies that enable a community to assemble knowledge structures that exceed their individual expertise) and process (technologies that enable members of a community to act in coordination with one another). In many ways, these systems can be viewed as a digital extension of the "Status Boards" discussed in Section 3.1.5 above.

3.2.2.1 Content

Content-oriented collaboration environments can be viewed as digital libraries (a domain with an extensive research literature [23], organizational infrastructure [22, 43], and public-sector support [17, 63]). Their function is to store, index, and provide access to shared materials. Because of the ease with which digital materials can be authored and distributed, these environments can be modified by the same communities to which they provide information, thus closing the stigmergic feedback loop. The computer adds three successive layers of functionality to the traditional library, discussed in the following three sections:

1. Enhanced storage and interlinking of materials, increasingly by the user community
2. Automatic ranking of materials, based on utilization by others
3. Dynamic distribution and sharing of content.

3.2.2.1.1 Content Storage and Linking

Most computer users today are frequent users of the World-Wide Web (WWW), a worldwide network of interlinked documents to which anyone can add. The notion of such a web of information was initially proposed by Vannevar Bush in 1945 [13]. There were numerous attempts at implementation, but the approach that finally took hold was that developed by Berners-Lee [7], based on a common linking protocol (http) embedded in simple text files.

The strength of the WWW is its open character, enabling it to grow rapidly. As sites link to one another, the web becomes a framework for self-organizing communities [96] (Figure 16). In fact, one of the first references to appear in response to a Google search on "stigmergy" (as of the date of writing) is a discussion of the collaborate effect of web logs, or blogs [34]. In the military domain, the US Army has made effective use of the emergent character of the WWW to share knowledge and experience among soldiers [91].

For some commercial purposes private file storage and sharing mechanisms are preferred (and often built on top of the WWW). One widely marketed example is Groove [35], which provides a common file repository and a variety of tools for project management to support distributed project teams.

Originally, the only way to modify the WWW was to add a new document that contained links to documents already there. A given document could only be changed by its author. Recent technical developments (most prominently, Wiki [51]) enable the maintenance of web pages that can be edited by anyone with the appropriate access. A prominent example of the potential of this approach is the WikiPedia [102], an encyclopedia of over half a million articles (many of very high quality), maintained entirely by the users. (By way of comparison, the Encyclopaedia Britannica contains about 120,000 articles.) The Wiki technology is an example of tools for social bookmarking [38], by which people can share not only their documents but also their annotations on the documents of others. This dynamic extends to entire libraries the dynamics of shared document editing discussed in Section 3.1.4 above.

Environment: Networked Computers
- Topology: Categories (but linked through cross references)
- State: Current collection of articles
- Dynamics:
  - Web link: Addressing
  - Search engine: Indexing, Ranking, Summarizing

Agents: People
- Sensor:
  - Sema: Entries
  - Marker: Linking Scores (Google)
- Actuator (Sema): Post a document
- Dynamics: Find desired information

Emergent system behavior: Maintain rationalized system of interrelated information

Figure 16: World-Wide Web

3.2.2.1.2 Site Ranking

In a shared information system such as the WWW or a Wiki, computers store the content, make it available to users, and facilitate changes and additions, but in the original form of such systems the interpretation of the resulting network is based on the human user's perception. Computers can augment this perception, to provide the user with a richer view of the network than would otherwise be possible. For example, the same openness to growth that makes the WWW so powerful can also make it overwhelming. A keyword search can return thousands of documents ("stigmergy," in spite of its rarity, returns 16,500 under Google at the time of writing), far more than a human user can effectively use. A common example of this enhancement is the ranking of web sites.

For example, much of the benefit of the Google indexing system lies in its PageRank^{TM} algorithm

Environment: Network of hyperlinked documents
- Topology: Graph
- State: Connectivity of the graph
- Dynamics: Maintenance of indices (via web spiders)

Agents: Creators of links in pages
- Sensor (Sema): Content of other relevant pages
- Actuator (Marker): Insert links to relevant pages
- Dynamics: Seeks to maximize connectivity of own pages to relevant pages

Emergent system behavior: Pages are assigned ranks that guide agents in finding useful material.

Figure 17: Google Page Rank

[10, 32, 64] (Figure 17), which assigns each web page a rank based on how many pages point to it, weighted by the ranks of those pages and the number of pages to which they point. This recursive computation is far beyond the ability of a human to calculate, but can readily be performed by computers, and enables users to find valuable material far more easily than would otherwise be the case. A page's rank under this algorithm is a clear example of an emergent phenomenon generated stigmergically. It guides web page authors seeking material to reference in their own pages, and in turn is modulated by the links that those authors insert to the material they find. Authors who wish to promote their own pages often seek to subvert the stigmergic nature of the system and control their page ranking. Much of Google's effort is devoted to blocking such subversion.

Another system for helping users find useful material is the PackHunter system [82]. Users deposit digital pheromones on a map of their web browsing activity, leading other users with similar pheromone patterns to sites likely to be of interest to them. Where PageRank™ focuses on the structure of the network as constructed by its authors, PackHunter takes into account the actions of users. It can be viewed as a form of collaborative filtering (Section 3.2.3) applied to web pages.

3.2.2.1.3 Peer-to-Peer Computing

In the WWW, every document lives on some computer, which must be active to make the content available. If a document is very popular, the computer's individual bandwidth may be overwhelmed, making access slow or impossible. The computer hosting the document is called a "server," and computers that seek to access it are "clients."

An alternative strategy, peer-to-peer (P2P) computing, seeks to do away with the client-server distinction. In this strategy, content moves dynamically from one computer to another, and may exist in several places at the same time. When a user seeks a document, the network dynamically retrieves it from the nearest available machine. No single server has a monopoly on any document, and if one source for a document is heavily loaded, others can supply the demand.

Environment: Computer network
- Topology: Graph
- State: Semantic signature of users at each computer
- Dynamics: Maintain semantic signatures based on documents generated and accessed

Agents: Users at nodes of the network
- Sensor (Sema): Review of accessed documents
- Actuator:
  - Sema: Documents and queries generated
  - Marker: Rewards sent to relevant documents
- Dynamics: User interactively generates and reviews documents

Emergent system behavior: Documents find their way to nodes where they are likely to be of most value

Figure 18: PARTNER (Smart P2P Document Distribution)

Current implementations of P2P, such as Gnutella [31], serve mainly as ways to share content without the need for a central file server, and are popular mechanisms for private users interested in sharing media. LOCKSS [52] applies the P2P approach to preserving digital media by distributing them over multiple machines, and is oriented to the library community. BitTorrent [8] provides a general-purpose peer-to-peer distribution system. In principle, such a framework would be an excellent environment for "smart" information that finds its way to users based on an emergent model of their interests, as in Altarum's PARTNER technology [68] (Figure 18) and in a research project currently underway at the Université de Tours in France [60].

3.2.2.2 Process

The collaboration systems described so far concentrate on making digital content readily accessible to users. A further level of collaborative support is represented by systems that help users manage the processes of their work. We briefly describe several examples, beginning with the most mature.

3.2.2.2.1 BPM: ActionWorks

There is an established market in the commercial world for workflow or business process management (BPM) systems. These systems help organizations define and follow standard processes to ensure uniformity of performance. A premier example is the ActionWorks system [2] from Action Technologies [1] (Figure 19). This system analyzes all workflows as built on a basic four-step cycle involving a Customer for whom work is being done and a Performer who does the work.

1. The Customer *prepares* a plan of the work to be done and issues a request.
2. The Customer and Performer *negotiate* the terms of the work.
3. The Performer *performs* the work and reports completion.
4. The Customer evaluates the work and either *accepts* it or identifies what remains to be done.

Each of these steps can in turn be broken down into further cycles until the company's entire business process has been analyzed. The resulting network forms an environment that supports stigmergic interactions among workers, who receive and give local signals concerning the state of their own responsibilities.

Environment: Model of interlocking Prepare-Negotiate-Perform-Accept cycles
- Topology: Graph
- State: Identify of Customer and Performer for each cycle, and current state of the cycle
- Dynamics: Propagate information about the state of each cycle to its sub- and super-cycles

Agents: Workers
- Sensor (Marker): State of current cycle and component subcycles
- Actuator (Marker): Report state of current work package
- Dynamics: Seek to move along cycles in which one is either a Customer or a Performer

Emergent system behavior: Coordinated execution of an overall workflow without missing or duplicative action.

Figure 19: ActionWorks BPM Framework

Environment: Web site
- Topology: Graph
- State: Information about products offered, seller identity, reputation, and conditions, buyer identity, reputation, and bid, state of the overall auction process
- Dynamics: Maintains and publishes current state of the auction; determines the winner; notifies participants.

Agents: Buyers and sellers
- Sensor (Marker): Current prices
- Actuator (Marker): State own bid or ask price
- Dynamics: Maximize time integral of revenue over expense

Emergent system behavior: globally optimum allocation of resources

Figure 20: On-Line Auctions

3.2.2.2.2 On-Line Auctions

On-line auctions such as eBay [26] provide a standardized process that guides sellers and buyers in finding one another, engaging in bidding, and concluding deals. The process being automated is essentially a Walrasian auction (Section 3.1.2). The overall system includes WWW structures for organizing products offered for sale and bids offered, a reputation system for enforcing honesty in transactions, and time-based mechanisms for managing the flow of an actual auction (Figure 20). Similar mechanisms are provided by Amazon in support of the network of used book sellers that advertise through its website, or the Yahoo merchants network.

Environment: Set of markets on interfaces
- Topology: Graph reflecting product structure
- State: Current bid and ask prices
- Dynamics: Compute prices that clear the markets

Agents: Designers
- Sensor (marker): Current prices
- Actuator (marker): State bid or ask
- Dynamics: Maximize time integral of revenue over expense

Emergent system behavior: Balanced assignments to interacting design variables

Figure 21: RAPPID

3.2.2.2.3 Market-Based Design: RAPPID

Altarum's RAPPID technology for distributed electromechanical design [78-80] (Figure 21), developed under the DARPA RaDEO/MADE program, helps designers reach agreement on design specifications at the intersections. RAPPID is based on a generalization of Walrasian market model for coordination (Section 3.1.2). A market exists for each interface parameter (for example, the torque or RPM of a shaft connecting a motor and a transmission), and the goods being traded are the assignments to those parameters. The prices manipulated in the markets are either catalog costs for actual components or "play money" that designers are allocated by the customer and must spend to get the functionality they require from other designers.

3.2.2.2.4 Battle Plan Adjustment: Coordinators

A current DARPA program, Coordinators [95], focuses on the task of helping fielded military units adapt their mission plans as the situation around them changes. Each unit has a networked computer or personal digital assistant (PDA). Agents representing each unit and running on their computer negotiate with one another to determine the interactions among tasks, the impact of the unfolding battle, and possible adaptive changes such as task timings, task assignments, or adoption of pre-planned contingencies.

One analyst has briefly described guerilla operations from a stigmergic perspective [85].

Environment: Communication network
- Topology: Graph
- State: Current state of the plan; representation of extraneous events
- Dynamics: Negotiation of plan adjustments among agents

Agents: Warfighters
- Sensor (Sema): Learn of constraints from other warfighters, and recommendations from the agents
- Actuator (Sema): Represent state of the battle and current constraints in to the system
- Dynamics: Report constraints and preferences

Emergent system behavior: Adjusted battle plan that balances requirements of the entire team

Figure 22: DARPA Coordinators

3.2.3 Recommender Systems

A "recommender system" or "recommendation system" attempts to predict items (such as books, movies, or music) that a user may find interesting, based on the user's profile. Such systems are usually implemented using collaborative filtering [40]. The system (the stigmergic environment) collects a large number of profiles on different users. Each profile is a vector over the universe of items for which recommendations are being made, and the magnitude of each element in a given user's profile indicates the relative attractiveness of that item for that user. To make a recommendation, the system first finds other users whose profiles are similar to that of the user for whom the recommendation is intended (the recommendee). Then it identifies elements in their profiles with high scores, elements for which the recommendee has registered no score.

Environment: Collection of user profiles
- Topology: Colored graph (user and item nodes with links indicating preferences between users and items, and similarities between users)
- State: Degree of preference of each user for each item; degree of similarity among users
- Dynamics: Maintain preference and similarity scores

Agents: Purchasers
- Sensor (Sema): Description of recommended purchase
- Actuator (Sema): Purchase an item
- Dynamics: Spend money on items most likely to be of interest

Emergent system behavior: Identify items most likely to be of interest to the user

Figure 23: Recommender System

The big challenge for recommender systems is collecting and maintaining a collection of user profiles. Users typically find it onerous to rank their preferences explicitly. However, their preferences can often be deduced as a side-effect of other actions they take. Perhaps the best known recommender system is that used by Amazon to recommend books to its buyers. This system functions stigmergically. Every purchase of an item is taken as evidence that the purchaser has a high level of interest in that item, and is registered in that user's profile. Thus individual actions (purchases) leave signs in the environment (the collection of profiles) that are integrated to provide feedback to the individual (what book might be a reasonable next acquisition).

3.2.4 Scheduling and Planning

Many problems of great operational importance can be cast as scheduling and planning problems. Abstractly, these problems concern the allocation of scarce resources to a set of tasks over time. Much of the research in the field of Operations Research is devoted to formal algorithms for solving such problems, with special emphasis on guaranteeing the optimality of the solution.

Environment: Factory conveyance structure
- Topology: Graph
- State: Connectivity between workstations; capabilities and load profile of each workstation
- Dynamics: Maintain load profiles on workstations

Agents: Workstations, Jobs
- Sensor (Marker): Market bids, pheromone levels
- Actuator (Market): Make bids; deposit pheromones; (jobs) select next operation
- Dynamics: Maximize local performance (utilization for a workstation; delivery time for a job)

Emergent system behavior: Reduced overall production time and increased throughput

Figure 24: Manufacturing Scheduling and Planning

A large body of centralized algorithms has been developed for solving these problems. However, many versions of the problem are NP-complete [28], meaning that for an instance of reasonable size, the time required to compute the solution is too long to complete by the time the plan is needed. To address these problems, a number of agent-oriented solutions have been proposed. While these methods vary considerably in their details, they tend to share two features. First, they do not guarantee optimal solutions, but use heuristics to obtain “good-enough” solutions in reasonable time. Second, a common heuristic is for the agents to restrict their interactions to other agents that are near them in some problem-specific topology. Thus these systems often qualify as stigmergic, and frequently draw on recognized stigmergic mechanisms, such as market or pheromone systems. We can illustrate these systems with two examples from the domain of manufacturing planning and scheduling.

The AARIA system [69, 70] is based on the model of the factory as a marketplace [4, 65]. The participants in the market are the workstations in the factory that can change the state of material, and agents representing each job that moves through the factory. A job agent negotiates with workstations to perform the operations that it requires in the appropriate sequence, using market mechanisms as already discussed in Section 3.1.2.

An important innovation in AARIA is the use of a loading profile or “dance card” on each workstation agent [18]. This profile aggregates the expected load on the workstation over time. Job agents search this profile to find relatively unoccupied times when they can be executed, and augment it when they book a reservation on a workstation. Thus the profile behaves like a pheromone over time, leveling the load that each workstation experiences.

The pheromone approach is taken even further in the system described in [11]. Each job sends out a swarm of ghost agents that explore alternative possible routings and record their findings in the form of digital pheromones in a graph representing the factory’s conveyance system. The actual job then follows the emergent pheromone trace. Another set of dynamics from insects, based on task allocation in wasps, has also been exploited in support of manufacturing planning and scheduling [16].

Yet another leading effort in stigmergic scheduling and planning applied to manufacturing is the work of the PROSA group at the Katholieke Universiteit Leuven in Belgium [93]. The acronym PROSA stands for “Product, Resource, Order, Staff Architecture,” and identifies the main software agents (products, resources, and orders) that interact stigmergically to guide the factory.

David Scheidt of Johns Hopkins University’s Advanced Physics Lab reports work on using markets for shipboard power and fluids distribution at JHU/APL, Rockwell Automation, Nutech, and Icosystems [88], but these systems are not documented in the open literature. Work on market mechanisms for power distribution in residential and industrial settings has been published [108, 109]

4 Future Developments

Further technical enablement of human-human stigmergy depends on developments in two areas: enabling technologies and specific applications.

4.1 Enabling Technologies

Computation-enhanced human-human stigmergy depends on significant advances in enabling technologies, including ubiquitous computing, theoretical foundations, simulation, and privacy and security.

4.1.1 Computing Hardware

The environment is the backbone of any stigmergic system. As computing technology becomes smaller, less costly, and lower in power requirements, it can be distributed more widely throughout the environment, increasing the potential for stigmergic interactions.

Embedded sensors in the environment will enable passive tracking of the locations of mobile entities. Roadway sensors that can detect vehicles by their ferrous signature are already a critical component of intelligent transportation systems (Section 3.2.1). RFID [101] provides a way to make less readily sensed objects (including humans) visible, and can provide support for digital pheromones [53]. A new generation of microsensors include limited computational and communications capability. The University of Berkeley is a leading center in the development of this hardware [83].

Today, we think of human interactions with computers as deliberate and explicit, relying on keyboards, display screens, and pointing devices. As sensors, processing, and communications become more tightly integrated with the environment, computers will become invisible [89], a movement that has been described as “ubiquitous” or “pervasive” computing. Humans will interact with computers via ordinary objects. (A current example is the modern automobile. Most drivers are unaware of the half-dozen or so computers they are manipulating as they drive down the road.) To enhance the transparency of this interaction, human interface devices will be critical. These include

- Heads-up displays to merge computer-generated information unobtrusively with the user’s normal field of vision,
- RFID technologies to track a person’s location and physiological state without explicit action on the part of the subject,
- Haptic technologies to guide the user through touch and feel.

An important area of research will be the development of power sources for pervasive computers. The smaller and more numerous computing elements become, the less power each will require, but the more impractical conventional power sources (such as batteries) become. In some cases (such as passive RFID), many passive sensors can receive the power they need from a few active readers. More generally, mechanisms such as parasitic power extraction from ambient RF or thermal noise will be come critical to the operation of such systems.

4.1.2 Foundations

Engineering requirements for ubiquitous computing are radically different from those that support the development, deployment, and maintenance of traditional computer systems [112], and will rely on a new body of theory and software engineering practices that can cope with the large numbers of processing elements, their physical distribution, the nonlinearity of their dynamics, and the nature of emergent behavior. Leading centers for the development of such software engineering methods include the Altarum Institute [67, 75] and the Università di Modena e Reggio Emilia in Italy with its vision of spraying computers onto the physical environment [111].

An important basis for this new theory is likely to be statistical physics, a mature quantitative science that is concerned with the emergence of macro-level system characteristics (such as pressure and temperature) from the behaviors of micro-level elements (atoms and molecules). Some research has been done on transferring insights from statistical physics to self-organizing systems [12, 50, 71, 74, 77, 86], but the field merits much more systematic exploration.

Applying formal methods will require adopting simplified models of the phenomena in question. One important such model in the case of resource allocation is the minority game [73, 86, 87]. Altarum has been a leader in developing this analogy. Another class of models is derived from the new science of network structure [62], where Mark Newman of the University of Michigan is a recognized leader [61]. Understanding better the dynamics of processes constrained to such networks is vital in understanding the spread of stigmergic information, such as the effectiveness of viral marketing.

The recent institution of the Workshop on Environments for Multi-Agent Systems (E4MAS [20, 100]) enables a necessary focus on the environment in which agents exist as a first-class engineering object in constructing and maintaining such systems. Recognizing the environment in this way not only facilitates system engineering, but is in fact essential to avoid inconsistencies that otherwise hamper the design [58]. Researchers from Altarum, Leuven, and Montpellier are the organizers of this series of workshops.

4.1.3 Integrated System-Social Simulation

Because of their emergent character, stigmergic systems require extensive use of simulation for their design and analysis.

Most simulation platforms used historically to study stigmergy in animals or robotic systems focus on support for large populations of fairly simple agents and the presence of an active environment. The leading simulation platforms in this area are Swarm [48] from the Santa Fe Institute, RePast [92] from the University of Chicago, and NetLogo [103] from Northwestern University.

There is a growing tradition in social simulation that takes account of the richness of human behavior. A leading center for social simulation is at the University of Surrey in the UK [29, 30]. These systems typically do not support large populations or environmentally-mediated interactions. Advances in human-human stigmergy will require simulation tools that combine the features of these two classes of existing systems.

4.1.4 Security and Privacy

One of the strengths of stigmergy is that information deposited by one agent can be retrieved and acted upon by other agents that visit the same location in the environment where the deposit was made. This openness of information poses a challenge from the point of view of security and privacy. Adversarial applications (such as commerce or warfare) require guarantees that adversaries will not be able to learn a system's intentions by eavesdropping, or to disrupt its operation by inserting malicious information. More generally, western society rightly places a premium on the rights of the individual human to protect personal information from broad dissemination.

In some respects, stigmergy is intrinsically more secure than current alternatives. Conventional systems for command and control involve messages whose contents provide considerable semantic detail about plans and objectives. Stigmergic messages, by way of contrast, are often numeric (e.g., a digital pheromone deposit). Such messages make sense only in the context of the entire system, making it much more difficult for an adversary either to interpret an intercepted message or to craft a spurious one to achieve a desired disruption. In addition, the stochasticity implicit in many stigmergic designs means that the detailed behavior of the system is unpredictable even to the user, making it even more obscure to an adversary without knowledge of its context.

Nevertheless, because of their use of the environment for coordination, responsible deployment of stigmergic systems will require new advances in security technology. Leading research in security for highly distributed systems is being pursued at many institutions, including IBM [15] and the University of New Mexico [27], and is published in the IEEE Distributed Systems Online e-journal [41] and workshops such as the International Workshop on Security in Distributed Computing Systems [36] and the International Workshop on Security in Systems and Networks [107].

4.2 Applications

The scope of possible applications of human-human stigmergy is difficult to bound, but it may be helpful to define several orthogonal dimensions along which one can categorize possible applications. We discuss three such dimensions: the problem domain, the system life cycle, and particular stigmergic mechanisms. Every combination of a category from each of these dimensions defines a possible application space to be explored, and we give some examples.

4.2.1 Problem Domain

What is the field or domain that is being addressed? Examples of high importance to the military include

- Logistics: facilitating the movement of materiel through the procurement supply chain and out to the warfighter in the field.
- Battle Management: coordinating multiple warfighters in the face of the fog and friction of war; providing an effective command and control loop between echelons.
- Intelligence Analysis: filtering massive amounts of data for relevant information; "connecting the dots" to identify critical information in time to take action on it; facilitating collaboration among intelligence analysts.
- Psychological Operations: changing the perceptions and expectations of the adversary, and in particular of civilians who may be undecided about which side to support.

4.2.2 Life Cycle

Operations in a domain typically move through a life cycle that includes, at a minimum, planning, execution, and after-action analysis.

4.2.3 Mechanisms

A handful of basic mechanisms arise repeatedly in the examples of human-human stigmergy that we have collected.

Resource allocation draws largely on market-inspired metaphors.

Knowledge formation enables the integration of contributions from many individuals into a coherent system of information.

Process management manages individual activities so that they reinforce an overall process map.

Movement coordination defines trails through a physical or cognitive space, and manages the traffic over them.

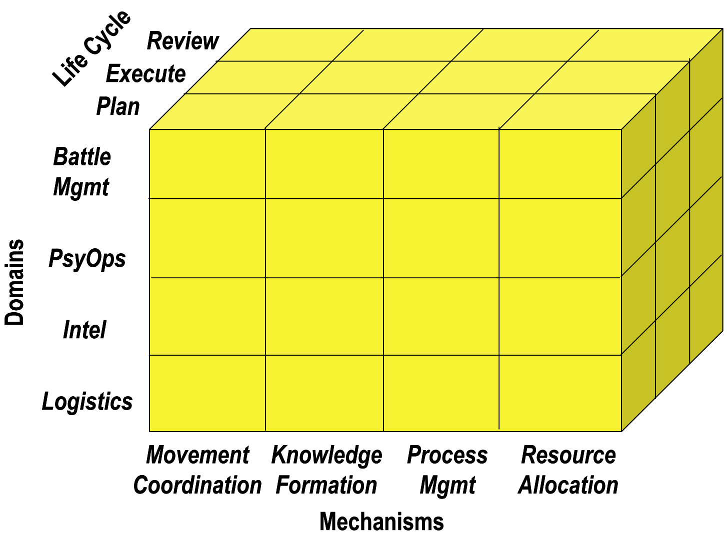

Figure 25: Application Space for Human-Human Stigmergy

4.2.4 Examples

Figure 25 summarizes the application space for human-human stigmergy, as outlined in the past three sections. While each of these dimensions could be refined and extended, these categories are sufficient to generate a good list of promising application areas. Each subcube in Figure 25 is a candidate application. In this section, we expand several by way of example. In each case, we indicate the element from the life cycle, the domain, and the mechanism.

4.2.4.1 Execute x Logistics x Movement Coordination

During the execution phase of logistics, materiel must move efficiently to users in the face of unexpected changes in supply, demand, and transport capacity. Stigmergic mechanisms for movement coordination (Sections 3.1.1 and 3.2.1) could move pallets to appropriate staging locations, generate appropriate field-deployable kits, and manage the movement of kits from the kitting location to the warfighter in the field. Such a system for real-time self-organizing logistics would focus human effort on establishing the requirements and priorities, while relieving humans from the mundane details of getting material to the right place at the right time.

4.2.4.2 Review x Battle Management x Knowledge Formation

A prototype of this application is currently active as Army Knowledge Online [91]. Presently, its main capability is content storage and linking, but it could be enhanced with mechanisms for site ranking (Section 3.2.2.1.2) and (in support of troops in the field) peer-to-peer distribution mechanisms (Section 3.2.2.1.3). In particular, the PARTNER vision for self-routing information [68] would be a great help in moving information to the people most likely to need it, before they ask for it.

4.2.4.3 Plan x Intel x Process Management

IPB (Intelligence Preparation of the Battlefield) is a labor-intensive process that requires coordination among a large community of sources, analysts, and operators. The techniques discussed in Section 3.2.2.2 could streamline this process and help it to converge more rapidly, while ensuring that essential processes are not overlooked. The "any-time" nature of stigmergic mechanisms could yield an IPB system that interfaces seamlessly with ongoing intelligence operations during combat and with after-action assessment activities.

4.2.4.4 Execute x PsyOps x Knowledge Formation

Understanding the dynamics of viral marketing (Section 3.1.6) may enable more effective application of techniques of collaborative knowledge formation (Section 3.2.2) to the adversary or relevant civilian populations in a campaign of "viral propaganda."

5 Technology Readiness Assessment

We have reviewed both (computer) applications of human-human stigmergy (from Section 3.2)and enabling technologies (from Section 4.1) against two readiness scales: the

Table 2: Two Systems for Assessing Technology Readiness

US Army / NASA (space systems)	Parunak (industrial systems)
1. Basic principles observed & reported 2. Tech concept / application formulated.	1. Modeled: architectural description or theoretical analysis
3. Analytical and experimental critical function / characteristic proof of concept	2. Emulated: simulation
4. Component / breadboard validation in lab	3. Prototype: domain HW under laboratory conditions
5. Component / breadboard validation in relevant environment 6. System / subsystem model of prototype demo in relevant environment 7. System prototype demo in operational environment 8. Actual system flight qualified thru test and demo	4. Pilot: commercial environment, off critical path
9. Actual system flight proven in missions	5. Production: commercial environment, on critical path
<<no counterpart in TRL system>>	6. Product: sold and supported for use by others

Table 3: Technology Readiness of Applications

Applications	TRL	Parunak
Intelligent Highway Systems	3-9	2-6
Content Collaboration: Linking (WWW, Wiki, Groove, List Servers), Ranking (Google), Sharing (Gnutella, LOCKSS, BitTorrent, …)	9	6
Process Collaboration: Workflow / BPM Software	9	6
Process Collaboration: DARPA Coordinators	3	2
Process Collaboration: RAPPID	4	3
Recommender Systems	9	6
Site Ranking (Google)	9	6
File Sharing Systems (Napster, Mobius, Grokster, …)	9	6
Resource Allocation	9	5-6
Scheduling and Planning	9	5-6

US Army/NASA TRL system [56] (originally designed for space systems), and a more industrially-oriented system [66]. Table 2 compares these two systems. Table 3 reviews the maturity of current applications, while Table 4 reviews the maturity of key enabling technologies.

Table 4: Technology Readiness of Enablers

Enablers	TRL	Parunak
Ubiquitous Sensors	4	3
Ubiquitous Processors	3-4	2-3
Ubiquitous Comms	3-4	2-3
Integrated Social-System Simulation	4	3
Foundations	3-4	2-3
Security	8-9	4-5
Privacy	2	1

This review shows that while there are many fairly mature applications (at least in the civilian sector), the enablers that will be needed to harden these applications for military use and to extend them further are much weaker.

In projecting these technologies into the future, we use the Parunak scale, and treat only those applications and enablers that are not currently at level 6. Table 5 summarizes our expectations about the progress one can expect. These estimates are based on the current state of basic research in the area, and the level of commercial demand. Where both are high (e.g., intelligent transportation systems), one can expect rapid maturation. Where demand exceeds technical capability (e.g., security), we expect early maturation of less capable versions of the technology, while where research is in advance of commercial demand (e.g., integrated simulation), we expect slower progress. The notes in the table reflect our rationale for the projected progress in each case.

Table 5: Technology Readiness Projections (Parunak Scale)					
	2005	2010	2015	2025	2035
Applications					
Intelligent Highway Systems Good market pull and supply pipeline.	2-6	5-6	5-6	6	6
Process Collaboration: DARPA Coordinators Weapons acquisition is slow.	2	4	5	6	6
Process Collaboration: RAPPID Lack of commercial demand in original domain, though concepts will be useful in other domains.	3	4	4	5	6
Enablers					
Ubiquitous Sensors RFID is moving very rapidly.	3	5	6	6	6
Ubiquitous Processors Power requirements make this the most difficult of the three components of ubiquitous computing to realize.	2-3	4	5	6	6
Ubiquitous Comms Will piggyback on community-wide WiFi and equivalents.	2-3	6	6	6	6
Integrated Social-System Simulation This can advance very rapidly given demand.	3	5	6	6	6
Foundations A “self-renewing” area that will uncover new problems as fast as it solves old ones. Anything beyond 3 is no longer “Foundations.”	2-3	3	3	3	3
Security Attacks will increase in sophistication to keep pace with developments. Products will be offered, but they will never be perfect.	4-5	6	6	6	6
Privacy Problem is ill-defined and likely to change shape due to political pressures.	1	3	3	3	3

6 Recommendations

Civilian applications provide low-hanging fruit to motivate military decision-makers, particularly in the areas of collaborative knowledge formation and process management. But from a research perspective, enablers need the most attention. A reasonable strategy would appear to be to sell projects based on military transition of commercially successful applications, while designing each project to address one or more key enablers to permit progress on these issues.

For example, Table 6 suggests how the potential applications outlined in Section 4.2.4 might be pursued with current technology while at the same time exploring needed new development.

Table 6: Recommended Project Strategies

Application	Existing Technology	New Research
Self-Organizing Logistics	RFID for package tracking	Development of statistical mechanics-based models for prediction and control
Army Knowledge Online	Existing web tools (as in the existing US Army site)	Peer-to-peer distribution Smart addressing (Altarum PARTNER)
IPB Process Management	Commercial workflow technology	Exploration of security mechanisms
Viral Psychological Operations	Conventional advertising techniques	Analysis of social network structure of target societies

7 Acknowledgments

This expert assessment has been considerably strengthened by the responses to a survey mailed to nineteen internationally known researchers in this field. We acknowledge the detailed responses provided by the following colleagues:

- Dr. Nicholas Monmarché, Université de Tours [59].
- Dr. David Payton, Hughes Research Laboratories, Malibu, California
- Dr. David Scheidt, Johns Hopkins University Applied Physics Laboratory, Laurel, MD, USA
- Dr. Paul Valckenaers, Dept. of Mechanical Engineering, Katholieke Universiteit Leuven, Leuven, Belgium
- Mr. Danny Weyns, Departement Computerwetenschappen, Katholieke Universiteit Leuven, Leuven, Belgium [99]
- Prof. Franco Zambonelli, Dipartimento di Scienze e Metodi dell'Ingegneria, Università di Modena e Reggio Emilia [110]

References

[1] Action. Action Technologies. 2005. http://www.actiontech.com/index.cfm.

[2] Action Technologies. Enabling Business Process Management. Action Technologies, Inc., Oakland, CA, 2005. http://www.actiontech.com/library/documents/EnablingBPM.pdf.

[3] R. Axtell and J. Epstein. Distributed Computation of Economic Equilibria via Bilateral Exchange. Brookings Institution, Washington, DC, 1997.

[4] A. D. Baker. *Manufacturing Control with a Market Driven Contract Net*. Ph.D. Thesis at Rensselaer Polytechnic Institute, Department of Electrical Engineering, 1991.

[5] T. Balch. Hierarchic social entropy: an information theoretic measure of robot team diversity. *Autonomous Robots*, 2000. http://www.cs.cmu.edu/~trb/papers/divjour.ps.

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

[7] T. Berners-Lee. *Weaving the Web: Origins and Future of the World Wide Web*. Texere Publishing, 1999.

[8] BitTorrent. The Official BitTorrent Home Page. 2005. Web site, http://www.bittorrent.com/index.html.

[9] E. Bonabeau, G. Theraulaz, V. Fourcassié, and J.-L. Deneubourg. The Phase-Ordering Kinetics of Cemetery Organization in Ants. *Physical Review E*, 4:4568-4571, 1998. http://www.santafe.edu/research/publications/workingpapers/98-01-008.ps.

[10] S. Brin and L. Page. The anatomy of a large-scale hypertextual Web search engine. *Computer Networks and ISDN Systems*, 30(1-7):101-117, 1998. http://citeseer.ist.psu.edu/brin98anatomy.html.

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

[12] 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 950-951, 2003. http://www.altarum.net/~vparunak/AAMAS03InfoPhaseChange.pdf.

[13] V. Bush. As We May Think. *The Atlantic Monthly*, vol. 176, pages 101-108, 1945. http://www.theatlantic.com/doc/prem/194507/bush.

[14] 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.

[15] D. Chess. Massively Distributed Systems. 2001. Web site, http://www.research.ibm.com/massive/.

[16] V. A. Cicirello and S. F. Smith. Wasp-like Agents for Distributed Factory Coordination. *Journal of Autonomous Agents and Multi-Agent Systems*, 8(3 (May)):237-266, 2004. http://www.cs.drexel.edu/~cicirello/jaamas-8-3.pdf.

[17] CIDL. Canadian Initiative on Digital Libraries / Initiative Canadienne sur les Bibliothèques Numériques. 2004. Web site, http://www.collectionscanada.ca/cidl/.

[18] S. J. Clark and H. V. D. Parunak. Density-Based Emergent Scheduling System. USA Patent # 5,953,229, Environmental Research Institute of Michigan, 1999. http://patft.uspto.gov/netacgi/nph-Parser?Sect1=PTO1&Sect2=HITOFF&d=PALL&p=1&u=/netahtml/srchnum.htm&r=1&f=G&l=50&s1='5,953,229'.WKU.&OS=PN/5,953,229&RS=PN/5,953,229.

[19] S. H. Clearwater, Editor. *Market-Based Control: A Paradigm for Distributed Resource Allocation*. Singapore, World Scientific, 1996.

[20] T. DeWolf. The Second International Workshop on Environments for Multiagent Systems. 2005. Web site, http://www.cs.kuleuven.ac.be/~distrinet/events/e4mas/.

[21] V. Dimitrov. Swarm-Like Dynamics and their Use in Organizations and Management. 2000. http://www.zulenet.com/VladimirDimitrov/pages/swarmanprint.html.

[22] DLF. Digital Library Federation. 2005. Web site, http://www.diglib.org/.

[23] DLib. D-Lib Magazine. 2005. Web site, http://www.dlib.org.

[24] DOT. ITS Joint Program Office Home. 2005. http://www.its.dot.gov.

[25] J. C. Dyer. Indian Trail Trees. 2004. http://home.att.net/~trailtrees/.

[26] eBay. eBay. 2005. Web site, http://www.ebay.com.

[27] S. Forrest. Forrest Home page. 2005. Web site, http://www.cs.unm.edu/~forrest/.

[28] M. R. Garey and D. S. Johnson. *Computers and Intractability*. San Francisco, CA, W.H. Freeman, 1979.

[29] N. Gilbert, Editor. *Journal of Artificial Societies and Social Simulation*. Surrey, UK, University of Surrey, 1998-.

[30] N. Gilbert. Centre for Research on Simulation in the Social Sciences. 2001. Web Page, http://www.soc.surrey.ac.uk/research/simsoc/cress.html.

[31] Gnutella. Gnutella.com. 2005. Web site, http://www.gnutella.com/.

[32] Google. Google Technology. 2004. Web site, http://www.google.com/technology/.

[33] 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.

[34] J. Gregorio. Stigmergy and the World-Wide Web. 2002. Blog page, http://bitworking.org/news/Stigmergy.

[35] Groove Networks. Groove Virtual Office - Virtual office software for sharing files, projects and data. 2005. Web Site, http://www.groove.net/home/index.cfm.

[36] Y. Guan and W. Zhao. 2005 International Workshop on Security in Distributed Computing Systems (SDCS 2005). 2005. Web site, http://securityworkshop.ece.iastate.edu/.

[37] H. A. Gutowitz. Complexity-Seeking Ants. In *Proceedings of Third European Conference on Artifical Life*, 1993. http://www.santafe.edu/~hag/ants/ants.html.

[38] T. Hammond, T. Hannay, B. Lund, and J. Scott. Social Bookmarking Tools (I): A General Review. *D-Lib Magazine*, 11(4 (April)), 2005. http://www.dlib.org/dlib/april05/hammond/04hammond.html.

[39] 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. http://xxx.lanl.gov/ps/cond-mat/9806097.

[40] F. Heylighen. Collaborative Filtering. 2001. Web page, http://pespmc1.vub.ac.be/COLLFILT.html.

[41] IEEE. IEEE Distributed Systems Online--Security. 2005. Web site, http://dsonline.computer.org/portal/site/dsonline/menuitem.9ed3d9924aeb0dcd82ccc6716bbe36ec/index.jsp?&pName=dso_level1&path=dsonline/topics/security&file=index.xml&xsl=article.xsl&;jsessionid=CGj1WsF77cmpLrdrnLm7558Lpwt8blznMbx075JDjftDV5k57RGf!1502094820.

[42] IEEE. IEEE ITSS Transactions and Society. 2005. http://www.ewh.ieee.org/tc/its/trans.html.

[43] IFLA. Digital Libraries: Resources and Projects. 2003. Web site, http://www.ifla.org/II/diglib.htm.

[44] ITS America. ITS America Home. 2005. http://www.itsa.org.

[45] ITS-STI Canada. ITS - STI Canada - Intelligent Transportation Systems Society of Canada | Systèmes de Transportes Intelligents Société du Canada. 2005. http://www.itscanada.ca/.

[46] R. E. Korf. A Simple Solution to Pursuit Games. In *Proceedings of Eleventh International Workshop on Distributed Artificial Intelligence*, Glen Arbor, MI, pages 183-194, 1992.

[47] T. J. Lambert, III, M. A. Epelman, and R. L. Smith. A Fictitious Play Approach to Large-Scale Optimization. *Operations Research*, (forthcoming), 2005.

[48] C. Langton, R. Burkhart, and G. Ropella. The Swarm Simulation System. 1997. http://www.swarm.org.

[49] 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. http://polaris.ing.unimo.it/didattica/curriculum/marco/Web-Co-Fields/stuff/Swarm.pdf.

[50] K. Lerman and A. Galstyan. A General Methodology for Mathematical Analysis of Multi-Agent Systems. ISI-TR-529, USC Information Sciences Institute, Marina del Rey, CA, 2001. http://www.isi.edu/%7Elerman/papers/isitr529.pdf.

[51] B. Leuf and W. Cunningham. Wiki: Welcome Visitors. 2002. Web site, http://wiki.org/.

[52] LOCKSS. LOCKSS Program Home. 2005. Web site, http://lockss.stanford.edu/.

[53] M. Mamei and F. Zambonelli. Spreading Pheromones in Everyday Environments via RFID Technologies. In *Proceedings of 2nd IEEE Symposium on Swarm Intelligence*, IEEE, 2005. http://polaris.ing.unimo.it/Zambonelli/PDF/Swarm05.pdf.

[54] 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. http://polaris.ing.unimo.it/Zambonelli/PDF/Teamwork.pdf.

[55] 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. http://polaris.ing.unimo.it/Zambonelli/PDF/isads.pdf.

[56] J. C. Mankins. Technology Readiness Levels. NASA, 1995. http://advtech.jsc.nasa.gov/downloads/TRLs.pdf.

[57] L. Martineau. *The Rocks Begin to Speak*. Las Vegas, NV, KC Publications, 1973.

[58] F. Michel. *Formalisme, méthodologie et outils pour la modélisation et la simulation de systèmes multi-agents*. Doctorat Thesis at Université des Sciences et Techniques du Languedoc, Department of Informatique, 2004. http://www.lirmm.fr/~fmichel/these/index.html.

[59] N. Monmarché. HaNT Web Site: Nicholas Monmarché. 2005. Web site, http://www.hant.li.univ-tours.fr/webhant/index.php?pageid=16.

[60] N. Monmarché. Personal Communication. 2005.

[61] M. Newman. Mark Newman. 2005. Web site, http://www-personal.umich.edu/~mejn/.

[62] M. E. J. Newman. The structure and function of complex networks. *SIAM Review*, 45:167-256, 2003. http://xxx.lanl.gov/pdf/cond-mat/0303516.

[63] NSF. Digital Libraries Initiative Phase 2. 2003. Web site, http://www.dli2.nsf.gov/.

[64] L. Page, S. Brin, R. Motwani, and T. Winograd. The PageRank Citation Ranking: Bringing Order to the Web. Stanford Digital Library Technologies Project, Palo Alto, CA, 1998. http://citeseer.ist.psu.edu/page98pagerank.html.

[65] H. V. D. Parunak. Manufacturing Experience with the Contract Net. In M. N. Huhns, Editor, *Distributed Artificial Intelligence*, pages 285-310. Pitman, London, 1987. http://www.altarum.net/~vparunak/YAMS87.pdf.

[66] H. V. D. Parunak. Agents in Overalls: Experiences and Issues in the Development and Deployment of Industrial Agent-Based Systems. *International Journal of Cooperative Information Systems*, 9(3):209-227, 2000. www.altarum.net/~vparunak/PAAM99Keynote.pdf.

[67] H. V. D. Parunak. Making Swarming Happen. In *Proceedings of Swarming and Network-Enabled C4ISR*, Tysons Corner, VA, ASD C3I, 2003. http://www.altarum.net/~vparunak/MSH03.pdf.

[68] H. V. D. Parunak. Challenging Old Assumptions in Global Information Management. In *Proceedings of ONR Conference on Collaborative Decision-Support Systems*, Quantico, VA, CADRC, California Polytechnic Institute, 2004. http://www.altarum.net/~vparunak/ONRCDS04.pdf.

[69] H. V. D. Parunak, A. D. Baker, and S. J. Clark. The AARIA Agent Architecture: From Manufacturing Requirements to Agent-Based System Design. *Integrated Computer-Aided Engineering*, 8(1):45-58, 2001. http://www.altarum.net/~vparunak/icaa98.pdf.

[70] H. V. D. Parunak and L. Barto. Agent-based Models and Manufacturing Processes. In *Proceedings of Embracing Complexity: The 1999 Colloquium on the Application of Complex Adaptive Systems to Business*, Boston, MA, pages 109-118, Ernst & Young Center for Business Innovation, 1999.

[71] H. V. D. Parunak and S. Brueckner. Entropy and Self-Organization in Multi-Agent Systems. In *Proceedings of The Fifth International Conference on Autonomous Agents (Agents 2001)*, Montreal, Canada, pages 124-130, ACM, 2001. www.altarum.net/~vparunak/agents01ent.pdf.

[72] H. V. D. Parunak and S. Brueckner. Swarming Coordination of Multiple UAV's for Collaborative Sensing. In *Proceedings of Second AIAA "Unmanned Unlimited" Systems, Technologies, and Operations Conference*, San Diego, CA, AIAA, 2003. http://www.altarum.net/~vparunak/AIAA03.pdf.

[73] 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. www.altarum.net/~vparunak/AAMAS02Effort.pdf.

[74] H. V. D. Parunak, S. Brueckner, and R. Savit. Universality in Multi-Agent Systems. In *Proceedings of Third International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2004)*, New York, NY, pages 930-937, IEEE, 2004. http://www.altarum.net/~vparunak/AAMAS04Universality.pdf.

[75] H. V. D. Parunak and S. A. Brueckner. Engineering Swarming Systems. In F. Bergenti, M.-P. Gleizes, and F. Zambonelli, Editors, *Methodologies and Software Engineering for Agent Systems*, pages 341-376. Kluwer, 2004. http://www.altarum.net/~vparunak/MSEAS03.pdf.

[76] H. V. D. Parunak, S. A. Brueckner, and J. Odell. Swarming Pattern Detection in Sensor and Robot Networks. In *Proceedings of 10th International Conference on Robotics and Remote Systems for Hazardous Environments*, Gainsville, FL, American Nuclear Society (ANS), 2004. http://www.altarum.net/~vparunak/ANS04.pdf.

[77] H. V. D. Parunak, S. A. Brueckner, J. A. Sauter, and R. Matthews. Global Convergence of Local Agent Behaviors. In *Proceedings of Fourth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS05)*, Utrecht, The Netherlands, pages (forthcoming), 2005. http://www.altarum.net/~vparunak/AAMAS05Converge.pdf.

[78] H. V. D. Parunak, A. C. Ward, M. Fleischer, and J. A. Sauter. The RAPPID Project: Symbiosis between Industrial Requirements and MAS Research. *Autonomous Agents and Multi-Agent Systems*, 2:2 (June):111-140, 1999. http://www.altarum.net/~vparunak/RAPPID99JAAMAS.pdf.

[79] H. V. D. Parunak, A. C. Ward, and J. A. Sauter. A Systematic Market Approach to Distributed Constraint Problems. In *Proceedings of International Conference on Multi-Agent Systems (ICMAS'98)*, Paris, France, pages 455-456, American Association for Artificial Intelligence, 1998. http://www.altarum.net/cec/rappid/icmas98.pdf.

[80] H. V. D. Parunak, A. C. Ward, and J. A. Sauter. The MarCon Algorithm: A Systematic Market Approach to Distributed Constraint Problems. *AI-EDAM: Artificial Intelligence for Engineering Design, Analysis and Manufacturing*, 13(3):217-234, 1999. http://www.altarum.net/cec/rappid/aiedam98.pdf.

[81] R. Pastor-Satorras and A. Vespignani. Epidemic Spreading in Scale-Free Networks. *Physical Review Letters*, 86:3200-3203, 2001.

[82] D. Payton, M. Daily, and K. Martin. Dynamic collaborator discovery in information intensive environments. *ACM Computing Surveys*, 31(2 (June)), 1999.

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

[84] 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.

[85] J. Robb. Stigmergic Learning and Global Guerillas. 2004. Web page, http://globalguerrillas.typepad.com/globalguerrillas/2004/07/stigmergic_syst.html.

[86] R. Savit, S. A. Brueckner, H. V. D. Parunak, and J. Sauter. Phase Structure of Resource Allocation Games. *Physics Letters A*, 311:359-364, 2002. http://arxiv.org/pdf/nlin.AO/0302053.

[87] R. Savit, S. A. Brueckner, H. V. D. Parunak, and J. Sauter. General Structure of a Class of Resource Allocation Games. *Physica A*, 345:676-704, 2005. www.altarum.net/~vparunak/RAGpaper.pdf.

[88] D. Scheidt. Personal Communication. 2005.

[89] N. Streitz and P. Nixon. Special Issue on The Disappearing Computer. *Communications of the ACM*, vol. 48, pages 32-71, 2005.

[90] Transport Canada. Intelligent Transportation Systems. 2005. http://www.its-sti.gc.ca/en/menu.htm.

[91] U.S. Army. Army Knowledge Online. 2005. https://www.us.army.mil/suite/login/welcome.html#.

[92] University of Chicago. RePast: An Agent Based Modelling Toolkit for Java. 2003. Web Page, http://repast.sourceforge.net/.

[93] P. Valckenaers, H. V. Brussel, K. Hadeli, O. Bochmann, B. S. Germain, and C. Zamfirescu. On the design of emergent systems: an investigation of integration and interoperability issues. *Engineering Applications of Artificial Intelligence*, 16:377-393, 2003.

[94] H. Varian. *Intermediate Microeconomics*. Fifth ed. New York, NY, W. W. Norton, 1999.

[95] T. Wagner. Coordinators. 2004. Web site, http://www.darpa.mil/ipto/Programs/coordinators/index.htm.

[96] F. Wang. Self-Organizing Communities Formed by Middle Agents. In *Proceedings of First International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2002)*, Bologna, Italy, pages 1333-1339, 2002.

[97] P. Weinstein, H. V. D. Parunak, P. Chiusano, and S. Brueckner. Agents Swarming in Semantic Spaces to Corroborate Hypotheses. In *Proceedings of AAMAS 2004*, New York, NY, pages 1488-1489, 2004. http://www.altarum.net/~vparunak/AAMAS04AntCAFE.pdf.

[98] 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.

[99] D. Weyns. Homepage of Danny Weyns. 2005. Web site, http://www.cs.kuleuven.ac.be/~danny/home.html.

[100] D. Weyns, H. V. D. Parunak, F. Michel, T. Holvoet, and J. Ferber. Multiagent Systems, State-of-the-Art and Research Challenges. In *Proceedings of Workshop on Environments for Multi-Agent Systems (E4MAS 2004)*, New York, NY, Springer, 2004. http://www.cs.kuleuven.ac.be/~danny/e4mas_survey_2004.pdf.

[101] Wikipedia. RFID. 2005. Web page, http://en.wikipedia.org/wiki/RFID.

[102] Wikipedia. Wikipedia, The Free Encyclopedia. 2005. Web site, http://en.wikipedia.org/wiki/Main_Page.

[103] U. Wilensky. NetLogo. 1999. Web site, http://ccl.northwestern.edu/netlogo.

[104] R. F. Wilson. The Six Simple Principles of Viral Marketing. *Web Marketing Today*, 2000. http://www.wilsonweb.com/wmt5/viral-principles.htm.

[105] M. J. Wooldridge and N. R. Jennings. Pitfalls of Agent-Oriented Development. In *Proceedings of 2nd Int. Conf. on Autonomous Agents (Agents-98)*, Minneapolis, MN, pages 385-391, 1998. http://citeseer.nj.nec.com/wooldridge98pitfalls.html.

[106] K. Wunderlich, D. Kaufman, and R. L. Smith. Link Travel Time Prediction for Decentralized Route Guidance Architectures. *IEEE Transactions on Intelligent Transportation Systems*, 1(1 (March)):4-14, 2000. http://www-personal.engin.umich.edu/~rlsmith/Predict%20for%20Decent.pdf.

[107] C.-Z. Xu and X. Zhou. The First International Workshop on Security in Systems and Networks (SSN2005). 2005. Web site, http://www.cs.uccs.edu/~SNS/sns2005.html.

[108] F. Ygge and H. Akkermans. Power Load Management as a Computational Market. In *Proceedings of Second International Conference on Multi-Agent Systems (ICMAS-96)*, Kyoto, Japan, pages 393-400, AAAI, 1996. http://www.enersearch.se/knowledgebase/publications/conference-journals/Icmas96/ICMAS96.html.

[109] F. Ygge and H. Akkermans. Decentralized Markets versus Central Control: A Comparative Study. *Journal of Artificial Intelligence Research*, 11:301-333, 1999.

[110] F. Zambonelli. Home page for Franco Zambonelli. 2005. Web site, http://www.dismi.unimo.it/Members/fzambonelli.

[111] F. Zambonelli, M. P. Gleizes, M. Mamei, and R. Tolksdorf. Spray Computers: Explorations in Self-Organization. *Journal of Pervasive and Mobile Computing*, 1(1 (March)), 2005. http://polaris.ing.unimo.it/Zambonelli/PDF/ACC.pdf.

[112] F. Zambonelli and V. Parunak. Towards a Paradigm Change in Computer Science and Software Engineering: a Synthesis. *The Knowledge Engineering Review*, (forthcoming), 2004. http://polaris.ing.unimo.it/Zambonelli/PDF/KER.pdf.

DOCUMENT CONTROL DATA

(Security classification of the title, body of abstract and indexing annotation must be entered when the overall document is classified)

1. ORIGINATOR (The name and address of the organization preparing the document, Organizations for whom the document was prepared, e.g. Centre sponsoring a contractor's report, or tasking agency, are entered in section 8.)

Publishing: DRDC
Performing: Altarum Institute, 3520 Green Court, Suite 300, Ann Arbor, MI 48105−1579
Monitoring:
Contracting: DRDC

2. SECURITY CLASSIFICATION (Overall security classification of the document including special warning terms if applicable.)

UNCLASSIFIED

3. TITLE (The complete document title as indicated on the title page. Its classification is indicated by the appropriate abbreviation (S, C, R, or U) in parenthesis at the end of the title)

Expert Assessment of Human−Human Stigmergy (U)
Évaluation experte de la stigmergie entre humains

4. AUTHORS (First name, middle initial and last name. If military, show rank, e.g. Maj. John E. Doe.)

H. Van Dyke Parunak

5. DATE OF PUBLICATION (Month and year of publication of document.)

May 2005

6a NO. OF PAGES (Total containing information, including Annexes, Appendices, etc.)

34

6b. NO. OF REFS (Total cited in document.)

7. DESCRIPTIVE NOTES (The category of the document, e.g. technical report, technical note or memorandum. If appropriate, enter the type of report, e.g. interim, progress, summary, annual or final. Give the inclusive dates when a specific reporting period is covered.)

Contract Report

8. SPONSORING ACTIVITY (The names of the department project office or laboratory sponsoring the research and development − include address.)

Sponsoring: DDA−ACD
Tasking:

9a. PROJECT OR GRANT NO. (If appropriate, the applicable research and development project or grant under which the document was written. Please specify whether project or grant.)

Contract W7714−040899/002/SV

9b. CONTRACT NO. (If appropriate, the applicable number under which the document was written.)

10a. ORIGINATOR'S DOCUMENT NUMBER (The official document number by which the document is identified by the originating activity. This number must be unique to this document)

DRDC CR 2005−003

10b. OTHER DOCUMENT NO(s). (Any other numbers under which may be assigned this document either by the originator or by the sponsor.)

11. DOCUMENT AVAILABILIY (Any limitations on the dissemination of the document, other than those imposed by security classification.)

Unlimited distribution

12. DOCUMENT ANNOUNCEMENT (Any limitation to the bibliographic announcement of this document. This will normally correspond to the Document Availability (11), However, when further distribution (beyond the audience specified in (11) is possible, a wider announcement audience may be selected.))

Unlimited announcement

DOCUMENT CONTROL DATA

(Security classification of the title, body of abstract and indexing annotation must be entered when the overall document is classified)

13. ABSTRACT (A brief and factual summary of the document. It may also appear elsewhere in the body of the document itself. It is highly desirable that the abstract of classified documents be unclassified. Each paragraph of the abstract shall begin with an indication of the security classification of the information in the paragraph (unless the document itself is unclassified) represented as (S), (C), (R), or (U). It is not necessary to include here abstracts in both official languages unless the text is bilingual.)

(U) Human−Human Stigmergy is pervasive. A wide range of pre−computer social systems fit the pattern of stigmergic coordination, and have provided a rich set of metaphors on which a diverse set of computer−enabled systems for enabling human stigmergy have been constructed. It would be more difficult to show a functioning human institution that is not stigmergic, than it is to find examples of human stigmergy.

The reason that human−human stigmergy is so common can be understood from the growing body of experience in constructing large−scale distributed computing systems with resource−constrained elements. It has become clear that central control of such systems is not feasible, since resource−constrained components cannot cope with the large−scale, distributed aspects of such systems. The central insight of stigmergy is that coordination can be achieved by resource−constrained agents interacting locally in an environment. Two fundamental principles govern the success of this strategy: 1. No matter how large the environment grows, because agents interact only locally, their limited processing capabilities are not overwhelmed; and 2. Through the dynamics of self−organization, local interactions can yield a coherent system−level outcome that provides the required control. The essence of stigmergy is the coordination of bounded agents embedded in a (potentially unbounded) environment, whose state they both sense (to guide their actions) and modify (as a result of their actions). The ability of stigmergic systems to scale without overwhelming individual agents requires that the environment offer a topology in which agents are situated locally, and within which their actions and senses have a limited horizon. Previous studies of stigmergic coordination, inspired by insects and applied to software and robotic systems, have identified four classes of stigmergic interaction between the agent and the environment [67]. Stigmergy can be either marker−based or sematectonic, and (independently) either quantitative or qualitative. All four of these can be identified in human−human stigmergy. Against this theoretical background, we classify a wide range of human stigmergic mechanisms, both pre−computational and computational. This survey shows that applications in the civilian world are both common and fairly mature. However, extending their effectiveness and transitioning them to military applications will require advances in several categories of enabling technology, including ubiquitous computing hardware and foundations, integrated social and distributed systems simulation technology, and technology for security and privacy. As a way ahead for the Canadian Defense Organization, we recommend leveraging the wide range of commercial systems in order to deliver short−term benefits to the military, but casting these transition efforts in a framework in which each supports the exploration of one or more critical enabling technologies.

14. KEYWORDS, DESCRIPTORS or IDENTIFIERS (Technically meaningful terms or short phrases that characterize a document and could be helpful in cataloguing the document. They should be selected so that no security classification is required. Identifiers, such as equipment model designation, trade name, military project code name, geographic location may also be included. If possible keywords should be selected from a published thesaurus, e.g. Thesaurus of Engineering and Scientific Terms (TEST) and that thesaurus identified. If it is not possible to select indexing terms which are Unclassified, the classification of each should be indicated as with the title.)

(U) stigmergy, stimergic systems