Expert Assessment of Stigmergy:

A Report for the Department of National Defence

Contract No.	**W7714-040899/003/SV**
File No.	**011 sv.W7714-040899**
Client Reference No.:	**W7714-4-0899**
Requisition No.	**W7714-040899**

Contact Info.

Tony White
Associate Professor
School of Computer Science
Room 5302 Herzberg Building
Carleton University
1125 Colonel By Drive
Ottawa, Ontario K1S 5B6
(Office) 613-520-2600 x2208
(Cell) 613-612-2708
arpwhite@scs.carleton.ca
http://www.scs.carleton.ca/~arpwhite

Abstract

This report describes the current state of research in the area known as Swarm Intelligence. Swarm Intelligence relies upon stigmergic principles in order to solve complex problems using only simple agents. Swarm Intelligence has been receiving increasing attention over the last 10 years as a result of the acknowledgement of the success of social insect systems in solving complex problems without the need for central control or global information. In swarm-based problem solving, a solution emerges as a result of the collective action of the members of the swarm, often using principles of communication known as stigmergy. The individual behaviours of swarm members do not indicate the nature of the emergent collective behaviour and the solution process is generally very robust to the loss of individual swarm members. This report describes the general principles for swarm-based problem solving, the way in which stigmergy is employed, and presents a number of high level algorithms that have proven utility in solving hard optimization and control problems. Useful tools for the modelling and investigation of swarm-based systems are then briefly described. Applications in the areas of combinatorial optimization, distributed manufacturing, collective robotics, and routing in networks (including mobile ad hoc networks) are then reviewed. Military and security applications are then described, specifically highlighting the groups that have been or continue to be active in swarm research. The final section of the document identifies areas of future research of potential military interest. A substantial bibliography is provided in support of the material provided in the report.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

GLOSSARY

ACO	Ant Colony Optimization
ACS	Ant Colony System
C4ISR	Command, Control, Communications, Computer, Intelligence, Surveillance and Reconnaissance
ISR	Intelligence, Surveillance and Reconnaissance
NASA	National Aeronautical Space Agency
SAM	Surface to air Missile
SEAD	Suppression of Enemy Air Defences
SMDC	Space Missile Defence Command
UAV	Unmanned Autonomous Vehicle

1 EXECUTIVE SUMMARY

This document describes the state of the art in stigmergy. Stigmergy, as originally described by Grassé in 1959, embraces the principle that the environment plays a crucial role in coordinating the activities of agents in a multi-agent system. A stigmergic system is one in which coordination of activity is achieved by individual agents leaving signals in the environment and other agents sensing them and using them to drive their own behaviour. Stigmergic systems solve problems in a bottom-up way – they self-organize – with no central controller or leader. Direct agent-to-agent communication is limited and reduced to local interactions only. Stigmergy is pervasive and is widely observed in social insect systems.

Stigmergy is not new to the military – swarming is an old military technique for harassing an enemy using only local information and decision making. However, only recently have researchers begun to encode stigmergic principles in multi-agent systems and many research projects still rely upon pure simulation environments. Theoretical studies are still lacking, although principles from statistical physics, chaos theory and other disciplines are likely to bear fruit in the next 10 years. A lack of theoretical results is partially to blame for the reluctance to use stigmergic routing algorithms in networks, for example.

The future for stigmergic systems is a bright one as notable successes have been observed in routing, optimization, search, and robot self-assembly, reconfiguration and repair. The latter examples, while immature, are encouraging in that future battlefield robotic systems may be able to rebuild damaged robots, scavenging for "transplants" to maintain their operational state. Growing circuits as exemplified by techniques from Amorphous Computing and self-repairing materials are also exciting areas in which future developments of military interest will occur.

This report identifies a number of stigmergic patterns of potential military value and provides a substantial bibliography related to stigmergic systems and swarm-based systems that utilize it. An excellent overview of stigmergy, swarm intelligence and its value to the military can be found in Dr. Van Parunak's paper entitled, "Making Swarming Happen", which can be found in: *Proceedings of Swarming: Network Enabled C4ISR,* Tysons Corner, VA, ASD C3I, 2003. Certain sections of this paper have been adapted for this report. The report provides several definitions of stigmergy in order to capture the various facets of stigmergic systems, with section 3 providing an overview of stigmergy.

The main body of the report begins in section 4, where principles of swarm-based systems that employ stigmergy are described. A taxonomy for stigmergic systems is introduced in this section. Readers unfamiliar with swarm-based systems should read this section. The section describes 5 patterns for stigmergic systems derived from the behaviour of social insects. When describing the stigmergic patterns suggestions are made as to potential military value. Applications that employ these patterns are briefly described in section 5.

Section 5 is a very long section containing a wide range of examples employing stigmergy. For readers interested only in military applications, section 5 contains several examples that include vehicle assembly and target tracking and acquisition. The technological level of sophistication is included in a brief assessment of the application at the end of the report. Collective robotics and mechatronics are included in this section. The section on collective robotics attempts to describe research on multi-robot systems with no central control and limited inter-robot communication focussing on the Swarm Bots project. The section on mechatronics briefly describes work on self-assembling and self-repairing robots; clearly an area of considerable military interest.

Arguably the most important section of a report of this type is the futures section. As mentioned above, this report has chosen to provide a taxonomy and patterns for stigmergic systems. This is useful because it provides investigators with tools to analyse systems of interest or a toolkit for system composition having target behaviours. This author strongly believes that tools for composition of stigmergic systems will be based upon patterns, much as is the case in current software engineering thinking. The futures section presents a futuristic battlefield scenario and then paints a brief research agenda that includes tools and techniques to support it. The agenda includes theoretical investigation, along with the construction of sophisticated simulators for the evaluation of battlefield scenarios where stigmergic systems are employed.

Section 7 summarizes the report content.

Section 8 provides information on people, companies and research projects in the general area of stigmergy and swarm-based systems. This list is best-effort; the area is rapidly changing. It should provide a rich starting pointing for military researchers wanting to engage in advanced investigation and prototyping.

A separate bibliography is provided containing of almost 300 references. Further references have also been provided electronically.

Finally, this document need not be read cover to cover. Several sections include deep coverage of a particularly important piece of research, which the reader can skip over on first reading. Sections inviting this cursory reading are indicated where appropriate.

2 INTRODUCTION

The Advanced Concept Development group of the Directorate of Defence Analysis (DDA), in partnership with the Directorate of Science and Technology Policy of Defence Research and Development Canada (DRDC) has requested an Expert Assessment of Emerging Technology in the area of Stigmergy, or more generally, Swarm Intelligence. Stigmergy represents an approach to communication in systems wherein simple agents, interacting locally and without recourse to global information, can solve complex problems. Problem solving is considered emergent in that no individual agent has sufficient capabilities to solve the problem alone. In fact, querying individual agents regarding their behaviour may imply little or nothing about the emergent behaviour of the swarm. Swarm-based systems are resilient to the failure of individual agents, and are capable of dealing with rapidly changing environmental conditions; two characteristics which makes them attractive for military environments.

This document provides a description of the state of the art in swarm research with a focus on how such research is relevant to problems of military and security interest.

The report is broken down into several further sections. The next section can be read without reference to the rest of the document. It is intended to provide a rapid introduction overview of swarm intelligence, including the brief description of a number of swarm-based examples. Readers intending to consume the detailed content of the document can reasonably skip over this section. An alternative to reading this section is to read Dr. Van Parunak's paper entitled, "Making Swarming Happen" [234] or the seminal earlier work [156] on engineering swarm-based systems. Section 3 begins by defining stigmergy and swarm intelligence and continues with a description of the essential characteristics of techniques used for swarm-based problem solving. Section 4 describes the principles of Swarm Intelligence. Readers need only take in sections 3 or 4. Section 5 reviews several applications that use swarm intelligence. Section 6 ends the report with a brief discussion of future research in the area of Swarm Intelligence. Section 7 reviews major sources of information on swarm-based problem solving, which includes web sites, influential people and projects.

2.1 OBJECTIVES

The objective of this document is to create a survey of the current state of the art in Swarm Intelligence, specifically highlighting the role of Stigmergy as a problem solving technique. The application of Swarm Intelligence in defence will be indicated, with the state of research being described as it pertains to military and security problems. A research agenda for work related to these areas will be proposed.

2.2 SCOPE

For the objectives to be met the document covers:

- The principles of Stigmergy:
    - Define the characteristics of systems exhibiting Swarm Intelligence.
    - Document several examples of naturally-occurring insect systems that demonstrate these characteristics.

  - Describe mathematical models of stigmergic systems.
- Highlight research of military and security significance in Swarm Intelligence
  - Characterize research according to technology readiness levels.
- Provide a review of emerging trends in Swarm Intelligence research.
  - Propose avenues for future research and development relevant to military and security applications.

2.3 DR. TONY WHITE

Dr. Tony White is an acknowledged expert in the field of Swarm Intelligence. He has published over 60 papers on subjects covering Multi-agent systems, Swarm Intelligence, Network and System Management, Evolutionary Computation and Combinatorial Optimization. He is currently an Associate Professor of Computer Science at Carleton University, Ottawa where he lectures on Swarm Intelligence to graduate students. He has master's degrees in Physics from Cambridge University, England and Computer Science along with a Ph.D. in Electrical Engineering from Carleton University in Ottawa. The focus of his Ph.D. was the use of stigmergic principles to solve control and management problems in communication networks. He has been awarded 7 patents with 3 others pending. Dr. White's current research areas include Swarm Intelligence, Autonomic Computing and the application of biological metaphors to problem solving in Computer Science.

3 AN INTRODUCTION TO SWARMS

This section can be read stand alone. If the reader requires an in depth understanding of stigmergy and swarm intelligence he should access section 4.

During the course of the last 20 years, researchers have discovered the variety of interesting insect and animal behaviours in nature. A flock of birds sweeps across the sky. A group of ants forages for food. A school of fish swims, turns, flees together, etc. [1]. We call this kind of aggregate motion *swarm behaviour*. Recently, biologists and computer scientists have studied how to model biological swarms to understand how such *social animals* interact, achieve goals, and evolve. Furthermore, engineers are increasingly interested in this kind of swarm behaviour since the resulting swarm intelligence can be applied in optimization (e.g. in telecommunication systems) [2], robotics [3, 4], track patterns in transportation systems, and military applications [5].

A high-level view of a swarm suggests that the N agents in the swarm are cooperating to achieve some purposeful behaviour and achieve some goal. This apparent collective intelligence seems to emerge from what are often large groups of relatively simple agents. The agents use simple local rules to govern their actions and via the interactions of the entire group, the swarm achieves its objectives. A type of self-organization emerges from the collection of actions of the group.

Swarm intelligence is the emergent collective intelligence of groups of simple autonomous agents. Here, an autonomous agent is a subsystem that interacts with its environment, which probably consists of other agents, but acts relatively independently from all other agents. The autonomous agent does not follow commands from a leader, or some global plan [6]. For example, for a bird to participate in a flock, it only adjusts its movements to coordinate with the movements of its flock mates, typically its neighbours that are close to it in the flock. A bird in a flock simply tries to stay close to its neighbours, but avoid collisions with them. Each bird does not take commands from any leader bird since there is no lead bird. Any bird can fly in the front, center or back of the swarm. Swarm behaviour helps birds take advantage of several things including protection from predators (especially for birds in the middle of the flock), and searching for food (as each bird is essentially exploiting the eyes of every other bird).

3.1 BIOLOGICAL BASIS AND ARTIFICIAL LIFE

Researchers try to examine how collections of animals, such as flocks, herds and schools, move in a way that appears to be orchestrated. A flock of birds moves like a well-choreographed dance troupe. They veer to the left in unison, and then suddenly they may all dart to the right and swoop down toward the ground. How can they coordinate their actions so well? In 1987, Reynolds created a *boid* model, which is a distributed behavioural model, to simulate on a computer the motion of a flock of birds [7]. Each boid is implemented as an independent actor that navigates according to its own perception of the dynamic environment. A boid must observe the following rules. First, the “avoidance rule" says that a boid must move away from boids that are too close, so as to reduce the chance of in-air collisions. Second, the “copy rule" says a boid must go in the general direction that the flock is moving by averaging the other boids' velocities and directions. Third, the “center rule" says that a boid should minimize exposure to the flock's exterior by moving toward the perceived center of the flock. Flake [6] added a fourth rule, “view," that indicates that a boid should move laterally away from any boid the blocks its view. This boid model seems reasonable if we consider it from another point of view, that of it acting according to attraction and repulsion between neighbours in a flock. The repulsion relationship results in the avoidance of collisions and attraction makes the flock keep shape, i.e., copying movements of neighbours can be seen as a kind of attraction. The center rule plays a role in both attraction and repulsion. The swarm behaviour of the simulated flock is the result of the dense interaction of the relatively simple behaviours of the individual boids. *To summarize, the flock is more than a set of birds; the sum of the actions results in coherent behaviour*.

One of the swarm-based robotic implementations of cooperative transport is inspired by cooperative prey retrieval in social insects. A single ant finds a prey item which it cannot move alone. The ant tells this to its nest mate by direct contact or trail-laying. Then a group of ants collectively carries the large prey back. Although this scenario seems to be well understood in biology, the mechanisms underlying cooperative transport remain unclear. Roboticists have attempted to model this cooperative transport. For instance, Kube and Zhang [2] introduce a simulation model including stagnation recovery with the method of task modeling. The collective behaviour of their system appears to be very similar to that of real ants.

Resnick [8] designed StarLogo (an object-oriented programming language based on Logo), to do a series of micro-world simulations. He successfully illustrated different self-organization and decentralization patterns in the slime mould, artificial ants, traffic jams, termites, turtle and frogs and so on.

Terzopooulos et al. [9] developed artificial fish in a 3D virtual physical world. They emulate the individual fish's appearance, locomotion, and behaviour as an autonomous agent situated in its simulated physical domain. The simulated fish can learn how to control internal muscles to locomotion hydrodynamically. They also emulated the complex group behaviours in a certain physical domain.

Millonas [10] proposed a spatially extended model of swarms in which organisms move probabilistically between local cells in space, but with weights dependent on local morphogenetic substances, or morphogens. The morphogens are in turn affected by the paths of movements of an organism. The evolution of morphogens and the corresponding flow of the organisms constitute the collective behaviour of the group.

Learning and evolution are the basic features of living creatures. In the field of artificial life, a variety of species adaptation genetic algorithms are proposed. Sims [11] describes a lifelike system for the evolution and co-evolution of virtual creatures. These artificial creatures compete in physically simulated 3D environments to seize a common resource.

Only the winners survive and reproduce. Their behaviour is limited to physically plausible actions by realistic dynamics, like gravity, friction and collisions. He structures the genotype by the directed graphs of nodes and connections. These genotypes can determine the neural systems for controlling muscle forces and the morphology of these creatures. They simulate co-evolution by adapting the morphology and behaviour mutually during the evolution process. They found interesting and diverse strategies and counter-strategies emerge during the simulation with populations of competing creatures.

3.2 SWARM ROBOTS

Swarm robotics is currently one of the most important application areas for swarm intelligence. Swarms provide the possibility of enhanced task performance, high reliability (fault tolerance), low unit complexity and decreased cost over traditional robotic systems. They can accomplish some tasks that would be impossible for a single robot to achieve. Swarm robots can be applied to many fields, such as flexible manufacturing systems, spacecraft, inspection/maintenance, construction, agriculture, and medicine [12]. Many different swarm models have been proposed. Beni [4] introduced the concept of cellular robotics systems, which consists of collections of autonomous, non-synchronized, non-intelligent robots cooperating on a finite n-dimensional cellular space under distributed control. Limited communication exists only between adjacent robots. These robots operate autonomously and cooperate with others to accomplish predefined global tasks. Hackwood and Beni [13] propose a model in which the robots are particularly simple but act under the influence of “signpost robots." These signposts can modify the internal state of the swarm units as they pass by. Under the action of the signposts, the entire swarm acts as a unit to carry out complex behaviours. Self-organization is realized via a rather general model whose most restrictive assumption is the cyclic boundary condition. The model requires that sensing swarm “circulate" in a loop during its sensing operation.

The behaviour-based control strategy put forward by Brooks [14] is mature and it has been applied to collections of simple independent robots, usually for simple tasks. Other authors have also considered how a collection of simple robots can be used to solve complex problems. Ueyama et al. [15] propose a scheme whereby complex robots are organized in tree-like hierarchies with communication between robots limited to the structure of the hierarchy.

Mataric [16] describes experiments with a homogeneous population of robots acting under different communication constraints. The robots either act in ignorance of one another, are informed by one another, or intelligently (cooperate) with one another. As inter-robot communication improves, more and more complex behaviours are possible.

Swarm robots are more than just networks of independent agents; they are potentially reconfigurable networks of communicating agents capable of coordinated sensing and interaction with the environment. Considering the variety of possible group designs of mobile robots, Dudek et al. [12] present a swarm-robot taxonomy of the different ways in which such swarm robots can be characterized. It helps to clarify the strengths, constraints and tradeoffs of various designs. The dimensions of the taxonomic axes are swarm size, communication range, topology, bandwidth, swarm reconfigurability, unit processing ability, and composition. For each dimension, there are some key sample points. For instance, swarm size includes the cases of single agent, pairs, finite sets, and infinite numbers. Communication ranges include none, close by neighbours, and “complete" where every agent communicate with every other agent. Swarm composition can be homogeneous or heterogeneous (i.e. with all the same agents or a mix of different agents). We can apply this swarm taxonomy to the above swarm models. For example, Hackwood and Beni's model [13] has multiple agents in its swarm, nearby communication range, broadcast communication topology, free communication bandwidth, dynamic swarm reconfigurability, heterogeneous composition, and its agent processing is Turing machine equivalent [12].

As research on decentralized autonomous robotics systems has developed, several areas have received increasing attention including modeling of swarms, agent planning or decision making and resulting group behaviour, and the evolution of group behaviour. The latter two can be seen as part of the branch of distributed artificial intelligence since several agents coordinate or cooperate to make decisions. There are several optimization methods proposed for the group behaviour. Fukuda et al. [17] introduced a distributed genetic algorithm for distributed planning in a cellular robotics system. They also proposed a concept of self-recognition for the decision making and showed the learning and adaptation strategy [18]. There are also other algorithms proposed.

3.3 EVALUATION OF SWARM INTELLIGENT SYSTEM

Although many studies on swarm intelligence have been presented, there are no general criteria to evaluate a swarm intelligent system's performance. Fukuda et al. [19] try to make an evaluation based on extensibility, which is essentially a robustness property. They proposed measures of fault tolerance and local superiority as indices. They compared two swarm intelligent systems via simulation with respect to these two indices. There is a significant need for more analytical studies.

3.4 STABILITY OF SWARMS

3.4.1 BIOLOGICAL MODELS

In biology, researchers proposed “continuum models" for swarm behaviour based on non-local interactions [20]. The model consists of integro-differential advection-diffusion equations, with convolution terms that describe long range attraction and repulsion. They found that if density dependence in the repulsion term is of a higher order than in the attraction term, then the swarm has a constant interior density with sharp edges as observed in biological examples. They did linear stability analysis for the edges of the swarm.

3.4.2 CHARACTERIZATIONS OF STABILITY

There are several basic principles for swarm intelligence, such as the proximity, quality, response diversity, adaptability, and stability. Stability is a basic property of swarms since if it is not present then it is typically impossible for the swarm to achieve any other objective. Stability characterizes the cohesiveness of the swarm as it moves. How do we mathematically define if swarms are stable? Relative velocity and distance of adjacent members in a group can be applied as criteria. Also, no matter whether it is a biological or mechanical swarm, there must exist some attractant and repellent profiles in the environment so that the group can move so as to seek attractants and avoid repellents. We can analyze the stability of swarm by observing whether swarms stay cohesive and converge to equilibrium points of a combined attractant/repellant profile.

3.4.3 OVERVIEW OF STABILITY ANALYSIS OF SWARMS

Stability of swarms is still an open problem. The current literature indicates that there is limited work done in this area. This is an extremely important consideration when deploying systems. We overview this work next.

Jin et al. [21] proposed the stability analysis of synchronized distributed control of 1-D and 2-D swarm structures. They prove that synchronized swarm structures are stable in the sense of Lyapunov with appropriate weights in the sum of adjacent errors if the vertical disturbances vary sufficiently more slowly than the response time of the servo systems of the agents. The convergence under total asynchronous distributed control is still an open problem. Convergence of simple asynchronous distributed control can be proven in a way similar to the convergence of discrete Hopfield neural network. Beni [22] proposed a sufficient condition for the asynchronous convergence of a linear swarm to a synchronously achievable configuration since a large class of distributed robotic systems self-organizing tasks can be mapped into reconfigurations of patterns in swarms. The model and stability analysis in [21, 22] is, however, quite similar to the model and proof of stability for the load balancing problem in computer networks [23].

References

[1] E. Shaw, The schooling of fishes, Sci. Am., vol. 206, pp. 128-138, 1962.

[2] E. Bonabeau, M. Dorigo, and G. Theraulaz, Swarm Intelligence: From Natural to Artificial Systems. NY: Oxford Univ. Press, 1999.

[3] R. Arkin, Behaviour-Based Robotics. Cambridge, MA: MIT Press, 1998.

[4] G. Beni and J. Wang, Swarm intelligence in cellular robotics systems, in Proceeding of NATO Advanced Workshop on Robots and Biological System, 1989.

[5] M. Pachter and P. Chandler, Challenges of autonomous control, IEEE Control Systems Magazine, pp. 92-97, April 1998.

[6] G. Flake, The Computational Beauty of Nature. Cambridge, MA: MIT Press, 1999.

[7] C. Reynolds, Flocks, herds, and schools: A distributed behavioural model, Comp. Graph, vol. 21, no. 4, pp. 25-34, 1987.

[8] M. Resnick, Turtles, Termites, and Traffic Jams: Explorations in Massively Parallel Microworlds. Cambridge, MA: MIT Press, 1994.

[9] D. Terzopoulos, X. Tu, and R. Grzeszczuk, Artificial fishes with autonomous locomotion, perception, behaviour, and learning in a simulated physical world, in Artificial Life I, p. 327, MIT Press, 1994.

[10] M. Millonas, Swarms, phase transitions, and collective intelligence," in Artificial Life III, Addison-Wesley, 1994.

[11] K. Sims, Evolving 3d morphology and behaviour by competition, in Artificial Life I, p. 353, MIT Press, 1994.

[12] G. Dudek and et al., A taxonomy for swarm robots, in IEEE/RSJ Int. Conf. on Intelligent Robots and Systems, (Yokohama, Japan), July 1993.

[13] S. Hackwood and S. Beni, Self-organization of sensors for swarm intelligence, in IEEE Int. Conf. on Robotics and Automation, (Nice, France), pp. 819-829, May 1992.

[14] R. Brooks, Intelligence without reason, tech. rep., Artificial Intelligence Memo. No. 1293, 1991.

[15] T. Ueyama, T. Fukuda, and F. Arai, Configuration of communication structure for distributed intelligent robot system," in Proc. IEEE Int. Conf. on Robotics and Automation, pp. 807-812, 1992.

[16] M. Mataric, Minimizing complexity in controlling a mobile robot population, in IEEE Int. Conf. on Robotics and Automation, (Nice, France), May 1992.

[17] T. Fukuda, T. Ueyama, and T. Sugiura, Self-organization and swarm intelligence in the society of robot being, in Proceedings of the 2nd International Symposium on Measurement and Control in Robotics, 1992.

[18] T. Fukuda, G. Iritani, T. Ueyama, and F. Arai, Optimization of group behaviour on cellular robotic system in dynamic environment, in Proceedings of the 1994 IEEE International Conference on Robotics and Automation, pp. 1027-1032, 1994.

[19] T. Fukuda, D. Funato, K. Sekiyam, and F. Arai, Evaluation on extensibility of swarm intelligent system, in Proceedings of the 1998 IEEE International Conference on Robotics and Automation, pp. 3210-3215, 1998.

[20] Mogilner and L. Edelstein-Keshet, A non-local model for a swarm, Journal of Mathematical Biology, vol. 38, pp. 534-570, 1999.

[21] K. Jin, P. Liang, and G. Beni, Stability of synchronized distributed control of discrete swarm structures," in IEEE International Conference on Robotics and Automation, pp. 1033-1038, 1994.

[22] G. Beni and P. Liang, Pattern reconfiguration in swarms convergence of a distributed asynchronous and bounded iterative algorithm," IEEE Trans. on Robotics and Automation, vol. 12, June 1996.

[23] K. Passino and K. Burgess, Stability Analysis of Discrete Event Systems. New York: John Wiley and Sons Pub., 1998.

4 PRINCIPLES OF SWARM INTELLIGENCE

This section provides a detailed introduction to swarm intelligence. Readers need not access section 3 to understand the content of this section.

4.1 OVERVIEW

The objective of this engagement is to provide a comprehensive assessment of the state of the art in Swarm Intelligence; specifically the role of stigmergy in distributed problem solving. In order to do this, working definitions have to be provided along with the essential properties of systems that are swarm-capable; i.e. problem solving is an emergent property of a system of simple agents. Several models of stigmergic systems are provided in this section; applications using the various models (singly or in combination) are described in a later section.

4.2 DEFINITIONS

The following definition for stigmergy has been proposed:

> *Grassé coined the term stigmergy (previous work directs and triggers new building actions) to describe a mechanism of decentralized pathway of information flow in social insects. In general, all kinds of multi-agent groups require coordination for their effort and it seems that stigmergy is a very powerful means to coordinate activity over great spans of time and space in a wide variety of systems. In a situation in which many individuals contribute to a collective effort, such as building a nest, stimuli provided by the emerging structure itself can provide a rich source of information for the working insects. The current article provides a detailed review of this stigmergic paradigm in the building behaviour of paper wasps to show how stigmergy influenced the understanding of mechanisms and evolution of a particular biological system. The most important feature to understand is how local stimuli are organized in space and time to ensure the emergence of a coherent adaptive structure and to explain how workers could act independently yet respond to stimuli provided through the common medium of the environment of the colony. [Istvan Karsai]*

A similar, but distinct, definition is:

> *Stigmergy is a class of mechanisms that mediate animal-animal interactions. Its introduction in 1959 by Pierre-Paul Grassé made it possible to explain what had been until then considered paradoxical observations. In an insect society individuals work as if they were alone while their collective activities appear to be coordinated. In this article we describe the history of stigmergy in the context of social insects and discuss the general properties of two*

distinct stigmergic mechanisms: quantitative stigmergy and qualitative stigmergy. [Theraulaz and Bonabeau]

In both definitions, the principle of stigmergy implies the interaction of simple agents through a common medium with no central control. This principle implies that querying individual agents tells one little or nothing about the emergent properties of the system. Consequently, simulation is often used to understand the emergent dynamics of stigmergic systems. Stigmergic systems are typically stochastic in nature; individual actions being chosen probabilistically from a limited behavioural repertoire. Actions performed by individual agents change the nature of the environment; for example a volatile chemical called a pheromone is deposited. This chemical signal is sensed by other agents and results in modified probabilistic choice of future actions.

The advantages of such a system are clear. Being a system in which multiple actions of agents are required for a solution to emerge, the activity of an individual agent is not as important. That is, stigmergic systems are resilient to the *failure of individual agents* and, more importantly still react extremely well to *dynamically changing environments*.

Optimal use of resources is often a significant consideration in designing algorithms. Another stigmergic system -- the raid army ant model – efficiently and effectively forages for food using pheromone-based signalling. In a raid army ant system, agents develop a *foraging front* that covers a wide path, leading to extremely effective food finding. This model has been simulated using NetLogo (for example) and the results agree extremely well with experimental observation. This model is described in some detail in Section 4.9.1.2.

This model has military value in that it could potentially be exploited as a series of mechanisms for searching for land mines, a problem that, tragically, is all too common in parts of the world.

A third stigmergic model of military interest is that of flocking or aggregation. Here, large numbers of simple agents can be made to move through a space filled with obstacles (and potentially threats) without recourse to central control. The environmental signals here are the position and velocities of the agents themselves. The utility of this model is that tanks could potentially be made to move across a terrain taking into account only tanks that are close by. A similar use of the model might be the self-organization of a squadron of flying drones.

Clearly, there are many examples of stigmergic systems that might be of use in a military environment and the examples described above are provided as a demonstration of an understanding of the area.

4.3 SWARM SYSTEMS

Considerable interest has been shown in Swarm Intelligence in the popular literature (e.g. Scientific American) and that interest is demonstrated both in industry and in research activity. As examples, using Google as a search engine with “swarm intelligence” as a search query, over 300,000 pages are returned; using Citeseer to search over 716,000 documents (academic papers), 172 are returned using the same query.

Interest in swarm systems reflect the belief that biologically-inspired problem solving – learning from and exploiting biological metaphors – holds considerable promise in terms of creating large, scalable, fault resistant agent systems.

The areas in which the applications of swarm principles have been applied are very diverse, they include: optimization, network management, collective robotics, supply chain management, manufacturing and military applications.

4.3.1 EMERGENT PROBLEM SOLVING

Emergent problem solving is a characteristic of swarm systems. Emergent problem solving is a class of problem solving where the behaviour of individual agents is not goal directed; i.e. by looking at the behaviour of single agents little or no information on the problem being solved can be inferred.

4.3.2 SWARM PROBLEM SOLVING

Swarm problem solving is a bottom-up approach to controlling and optimizing distributed systems. It is a mindset rather than a technology that is inspired by the behaviour of social insects that has evolved over millions of years.

The Scientific American article by Bonabeau and Theraulaz [151] is an excellent (and digestible) overview of swarm-based problem solving. The article discusses a number of social insect systems and practical problems that can be solved using algorithms derived from them. Peterson [152] suggests that swarms calculate faster and organize better.

Swarm systems are characterized by simple agents interacting through the environment using signals that are spatially (and temporally) distributed. By simple we mean that the agents possess limited cognition and memory; sometimes no memory at all. Furthermore, the behaviour of individual agents is characterized by a small number of rules. In this document we consider the complexity (or simplicity) of an agent to be a function of the number of rules that are required to explain its behaviour.

4.3.3 RELEVANCE TO MILITARY APPLICATIONS

Why is this important from a military perspective?

First, traditional military systems have been designed to be top-down, centralized control systems. They often assign fixed roles to entities within systems thereby allowing for system failure when a critical role becomes unavailable. Social insect systems, using response threshold mechanisms, exhibit no such characteristic. They exhibit flexible role assignment based upon perceived threats and stimuli. We have employed response threshold mechanisms in simulated robotic soccer, where the roles of defender, midfielder and attacker are dynamically assigned and can change during the game. Furthermore, the players can tire or become injured, as is the case in a real game. The results have been encouraging and require further investigation. While soccer is a game, it shares obvious characteristics with military war games, where a threat must be countered using an optimal distribution of available resources.

Knight [153] talks about the robot swarms for mine sweeping and search and rescue. Each agent in the swarm uses algorithms inspired by social insects. There are several examples of military applications that will be discussed later in the document.

4.3.4 ADVANTAGES AND DISADVANTAGES

There are several advantages:

A. Agents are not goal directed; they react rather than plan extensively.

B. Agents are simple, with minimal behaviour and memory.

C. Control is decentralized; there is no global information in the system.

D. Failure of individual agents is tolerated; emergent behaviour is robust with respect to individual failure.

E. Agents can react to dynamically changing environments.

F. Direct agent interaction is not required.

The table below (due to Eric Bonabeau) provides an alternative description of the advantages of swarm systems.

Flexible:	the colony can respond to internal perturbations and external challenges
Robust:	tasks are completed even if some individuals fail
Scalable:	from a few individuals to millions
Decentralized:	there is no central control(ler) in the colony
Self-organized:	paths to solutions are emergent rather than predefined

Table 1: Advantages of Swarm Systems

There are certain disadvantages:

A. Collective behaviour cannot be inferred from individual agent behaviour. This implies that observing single agents will not necessarily allow swarm-defeating behaviour to be chosen. (This can be viewed as an advantage too from an aggressive point of view).

B. Individual behaviour looks like noise as action choice is stochastic.

C. Designing swarm-based systems is hard. There are almost no analytical mechanisms for design.

D. Parameters that define the swarm system can have a dramatic effect on the emergence (or not) of collective behaviour.

Behaviour:	Difficult to predict collective behaviour from individual rules.
Knowledge:	Interrogate one of the participants, it won't tell you anything about the function of the group.
Sensitivity:	Small changes in rules lead to different group-level behaviour.
Actions:	Individual behaviour looks like noise: how do you detect threats?

Table 2: Disadvantages of Swarm Systems

4.4 MECHANISMS FOR UNDERSTANDING SWARM

The previous section indicated that there are several issues to address in order to design successful swarm systems. Essentially, three questions need to be answered:

1. How do we define individual behaviour and interactions to produce desired emergent patterns?

2. How do we shape emergence?

3. How do we fight swarms – organizations that operate on swarm principles?

Question 1 may often be answered through a combination of simulation and design using evolutionary computing. A detailed discussion of agent-based simulation and evolutionary computing is out of the scope of this report. However, agent-based simulation is a rapidly maturing area.

The idea with agent-based simulation is to associate simple rules with individual agents and run the simulation for a period of time until the emergent dynamics (if any) are manifest. An assessment of the emergent dynamics (driven by a human observer or through automation) can be used to guide a learning process (often drawn from evolutionary computation) which refines the rules used in the simulation. This iterative process of simulation followed by agent behaviour refinement is common in the literature; e.g. the development of swarm robot behaviours in the Swarm-Bots project.

Question 2 may be answered through simulation. Here, the idea is to understand how the swarm can be controlled through the parameters that characterize the system. For example, if a particular signal dissipates at a given rate, what should that rate be and how sensitive is the collective behaviour to it? Automated approaches to parameter space evaluation are possible [154], [155].

Question 3 is a difficult question to answer, but arguably the most important. Given that observing individual agent behaviour does not provide must insight into the collective behaviour of the swarm, it would seem to be an open question. However, section 4.9.5 provides some insight into the possibilities in the context of a particular stigmergic pattern.

4.5 HOW SELF-ORGANIZATION WORKS

Self-organization in swarm systems occurs through several means, not all of which have to be present in a system for effective problem solving to occur. It should be noted that agent memory is not an important aspect of a swarm system; the effects below are the principal components for self organization.

4.5.1 POSITIVE FEEDBACK

When an agent performs an action in the environment, the value of that action needs to be reflected in some change in the environment. For example, in ant foraging behavior, an ant successfully finding food returns to the nest dropping pheromone with an intensity that is proportional to the quality of the food source. A second example of positive feedback comes from nest building. An ant deposits a ball of mud; other ants seeing this deposit the ball of mud that they are carrying on top of it. As a result of this reinforcement, a wall is built. In the first example stigmergy is present explicitly, with an independent signal (the pheromone) providing the feedback. In the second example, stigmergy is present through the actual work being done – the wall being built. This second form of stigmergy is sematectonic stigmergy. The figure on the next page shows that the mud pile forms a *stimulus* to ant carrying a mud ball, which the ant responds to by dutifully adding its mud ball to the top of the pile. A third example is the clustering behaviour of ants; preferring to add something to a pre-existing pile, with the pile size making the addition all the more likely.

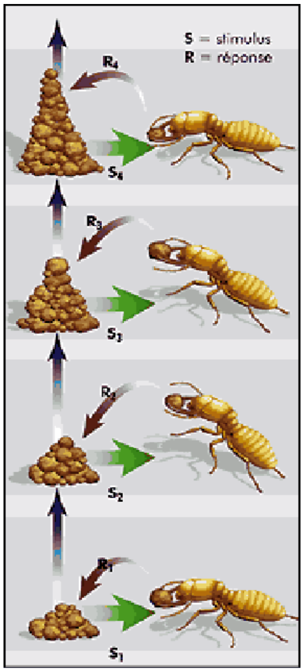

Figure 1: Sematectonic stigmergy

Positive feedback in a self-organized system drives agents in the system to reinforce actions that provide most gain to the collective. Positive feedback in stigmergic systems are often said to form an *autocatalytic process*.

4.5.2 NEGATIVE FEEDBACK

While positive feedback attracts more and more agents to participate in a problem solving process -- reinforcing the actions of other agents by making them more likely -- this can cause premature convergence to a suboptimal solution if negative feedback is not provided.

Negative feedback is used for stabilization and is designed to ensure that one decision, or a small number of poor decisions, will not bias the entire problem solving process. The highly volatile nature of pheromones provides this in ant systems. Pheromone volatility ensures that signals must be constantly reinforced in order to persist in the environment. Think of pheromone volatility – or negative feedback generally -- as "forgetting".

4.5.3 AGENT DIVERSITY

It is important that the behavior of agents exhibit diversity. This means that different decisions can be made for a given environment. Usually, when faced with several competing actions, an action value will be associated with each action and a stochastic choice will be made. Agent diversity can be achieved in other ways. As an example, imagine 3 distinct choices, with action values of 1, 2, and 3 respectively; then action 3 will be the most likely choice with probability 3/6 (=1/2) assuming a uniform distribution.

4.5.4 AMPLIFICATION OF FLUCTUATIONS

In most swarm systems there is stochastic behavior. For example, ants make choices as to where to forage for food, the decisions being made based upon pheromone levels. If we imagine 3 distinct choices, with levels of 1, 2, and 3 respectively; then direction 3 will be the most likely choice with probability 3/6 (=1/2) assuming a uniform distribution. However, in an absence of pheromone, a decision will be made and the value of this action will be amplified by other ants having their action choices biased by the pheromone laid down by other ants. It has also been shown that periodically ignoring signals in the environment can be beneficial. In this case an action is chosen randomly. Often crucial, this allows discovery of new solutions to occur.

Often the nature of the fluctuations in a swarm system is chaotic; that is, even though the emergent dynamics are predictable in some macro sense, the trajectory of the system and the micro structure cannot be predicted. For example, consider the clustering behavior of ants. While it can be demonstrated that given enough time ants will cluster all objects of a given type in a single pile, the location of the pile and the actual positions of individual objects cannot be predicted. Successive "runs" of a simulation of the system will yield significantly different structures; however, the objects will be sorted into clusters.

4.5.5 MULTIPLE INTERACTIONS

Another key attribute of swarm systems is that they rely on multiple interactions; i.e. many agents taking the same action, in order for problem solving behaviour to emerge. The interactions with the environment cause change, with the changes being reflected in the environment. Agent memory is not a significant factor in problem solving; the spatio-temporal patterns in the environment are. Signals from one individual have to be sensed by others for these multiple interactions to have value. The degree with which agents can sense other agents changes in the environment determines the value of multiple agent actions as those changes affect the decisions being made by the sensing agent.

4.5.6 CREATING SWARMING SYSTEMS

A swarm-based system can be generated using the following principles:

1. Agents are independent, they are autonomous. They are not simply functions as in the case of a conventional object oriented system.
2. Agents should be small, with simple behaviours. They should be situated and capable of dealing with noise. In fact, noise is a desirable characteristic.
3. Decentralized – do not rely on global information. This makes things a lot more reliable.
4. Agents should be behaviourally diverse – typically stochastic.
5. Allow information to leak out of the system; i.e. introduce disorder at some rate.
6. Agents must share information – locally is preferable.
7. Planning and execution occur concurrently – the system is reactive.

The principles outlined above come from Parunak [156]. More recently, the importance of gradient creation and maintenance has been stressed and that digital pheromones can be made to react in the environment, thereby creating new signals of use to other swarm agents [157].

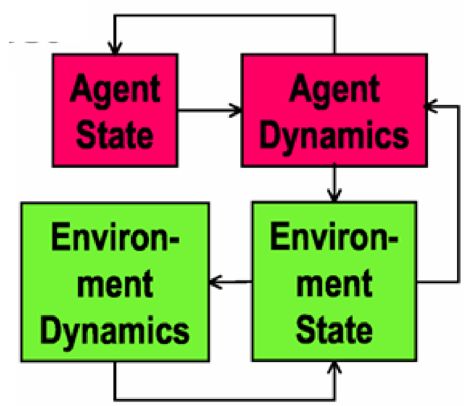

Figure 2: Agent-Environment Interaction

The above figure summarizes the interactions between agent and environment. Agent state along with environment state drives agent dynamics; i.e. agent action selection. Agent action selection changes environment state through the creation or modification of signals. Environment state is used as input to environment dynamics. The dynamics of the environment causes changes to occur in environment state. What is important in the above figure is that agent state is hidden – only the agent has access to it. Environment state is visible to the agent but has to be stored by the agent if it is to be reused at some later point in time when the agent has (presumably) moved to a different location.

4.6 HOW CAN WE MEASURE AND CONTROL SWARMING?

This section adapted from Parunak's presentation at the Conference on Swarming and C4ISR, Tyson's Corner, VA, 3rd June, 2003.

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

4.6.1 MEASUREMENT

Altarum have defined swarming as "useful self-organization of multiple entities through local interactions." The terms in this definition offer a useful template for measuring the performance of a swarm. The criteria of "multiple entities" and "local interactions" identify independent variables that characterize the kind of swarm being considered, while the notion of "useful self organization" leads to several dependent variables. Because of the nonlinearities involved in both individual agent behaviour and the interactions among agents, the values of the dependent variables can change discontinuously as the independent variables are adjusted, and qualification of a swarm requires careful study of such "phase shifts." An example of such a study is [223].

Multiple entities.—Sometimes mechanisms that work satisfactorily for small numbers of entities do not scale well as the population increases. In other cases, there may be a critical minimum population below which the swarm will not function. In evaluating swarms, it is crucial to study how the performance varies with population.

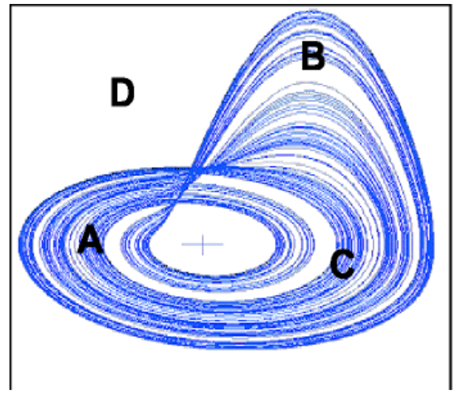

Figure 3: The Roessler Attractor

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

Measures of Usefulness.—The measures used to assess the usefulness of a swarm are drawn directly from the measurements in the problem domain. For example, in a target tracking problem the percentage of targets detected would be an important measure.

Measures of Self-Organization.— Some of the benefits of swarming are difficult to measure directly, but are directly correlated with the degree to which a swarm can organize itself. For example, directly assessing a swarm's robustness to unexpected perturbations would require a very large suite of experiments, but our confidence in this robustness can be strengthened if we can measure its self-organizing capabilities. Altarum has found a variety of measures derived from statistical physics to be useful indicators of self-organization, including measures of entropy over the messages exchanged by agents, their spatial distribution, or the behavioural options open to them at any moment [196]. Frequently, local measures of these quantities permit us to deduce the global state of the swarm, a crucial capability for managing a distributed system [223]. It has recently been suggested that a Lebesgue measure of the portion of the swarm's space of behaviours that is dominated by the Pareto frontier might also be a useful measure of self-organization [207].

4.6.2 CONTROL

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

Envelope Shaping.—While the details of a swarm's behaviour may vary from one run to the next, those variations often are constrained to an envelope that depends on the configuration of the swarm. An illustration of this distinction can be seen in the Roessler attractor from chaos theory (Figure 3). This figure is a plot in three-dimensional phase space of a set of differential equations in their chaotic regime. The line that twists through this figure indicates the trajectory of this system, a trajectory that is so intertwined that arbitrarily small differences in initial conditions can lead to widely varying outcomes. For instance, if the system starts at location "A," it is in principle impossible to predict whether at a specified future time it will be at location B or location C. However, in spite of its detailed unpredictability, the system is confined to a highly structured envelope, and it is impossible for it to visit the point D.

To shape a swarm's envelope, it is exercised in simulation, and human overseers evaluate its performance, rewarding appropriate behaviour and punishing inappropriate behaviour. Evolutionary or particle swarm methods then adjust the behaviours of individual swarm members so that desirable behaviour increases and undesirable behaviour decreases [193], [228]. The process adjusts the envelope of the system's behaviour so that undesirable regions are avoided. Incidentally, these techniques enable swarms to be trained rather than designed, an approach that reduces the need for specialized software skills on the part of the warfighter. Evolution can also be used to explore the behavioural space of a swarm in much greater detail than exhaustive simulation would permit, by selectively altering later simulation runs based on the results of earlier ones [197].

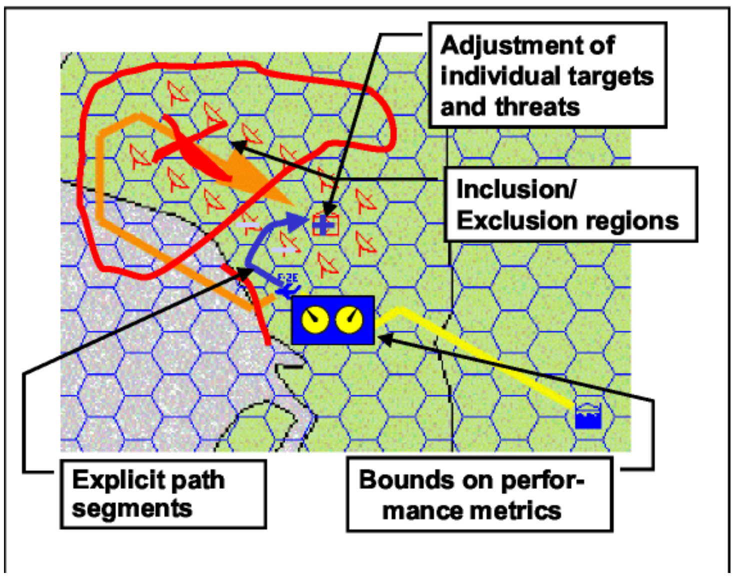

Figure 4: Annotating Performance

Managing by Exception.—Once a swarm has been launched, human overseers can observe its emerging behaviour and intervene on an exception basis. For example, a swarm with kill capability can autonomously detect a target and configure itself for attack, then apply for human permission to execute. Digital pheromones are especially amenable to human direction. Graphic marks on a map can be translated directly into pheromone deposits that modify the emergent behaviour of the swarm in real-time (Figure 4). A path being formed by the system can be blocked or a whole region excluded; the priority of individual targets and threats can be adjusted; segments of paths can be explicitly designated; and bounds can be placed on performance metrics. The important point is that human intervention is on an exception basis. Routine operation proceeds without detailed human control, freeing human warfighters to concentrate on more strategic concerns and calling their attention to situations where their judgment is required.

4.7 TAXONOMY FOR STIGMERGY

Carleton University
Canada's Capital University

Taxonomy of Stigmergy

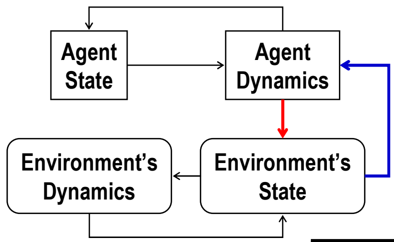

Agents sense and manipulate environmental state variables
•To coordinate with one another
•To solve their individual problems
Sometimes these categories are the same; sometimes different.

Insect Examples	Marker-Based **Artificial signs for coordination**	Sematectonic **Domain elements only**
Quantitative **Scalar quantities**	Gradient following in a single pheromone field	Ant cemetery clustering
Qualitative **Symbolic distinctions**	Decisions based on combinations of pheromones	Wasp nest construction

Figure 5: Taxonmy for Stigmergy

The figure above provides a taxonomy for stigmergy. The taxonomy is due to Parunak [234]. All examples described in this report can be described using this taxonomy. As the figure indicates, there are two dimensions to stigmergy. The first is shown horizontally and refers to the difference between a signal simply pointing in a certain direction (driving a particular action decision) and actually contributing to the solution of the problem. The second dimension, shown vertically, describes the complexity of signal content. Scalar quantities are simple; e.g. the concentration of a particular pheromone. However, more complex signals can also be represented; e.g. the configuration of a set of blocks in a structure.

4.8 TOOLS FOR INVESTIGATING SWARM SYSTEMS

As mentioned in a previous section, predicting the emergent behaviour of swarm systems based upon the behaviour of individual agents is generally not analytically tractable. Consequently, agent-based simulation is used to investigate the properties of these systems. This section briefly describes two tools useful for such investigations.

4.8.1 NETLOGO

NetLogo is a simple agent simulation environment based upon StarLogo, an environment by Resnick and described in his book entitled, “Turtles, Termites and Traffic Jams”. Users program using agents and patches (the environment). In NetLogo, the environment has active properties and is ideal in its support of stigmergy as agents can easily modify or sense information of the local patch or patches within some neighbourhood. Unlike conventional programming languages, the programmer does not have control over agent execution and cannot assume uninterrupted execution of agent behaviour. A fairly sophisticated user interface is provided and new interface components can be introduced using a drag-and-drop mechanism. Interaction with model variables is easily achieved through form-based interfaces. The user codes in NetLogo’s own language, which is simple and type-free (i.e. dynamically bound).

The environment, written in Java, is freely available from http://ccl.northwestern.edu/netlogo/. The environment comes with a large number of models that include several from biology, the social sciences, computer science and mathematics. Several community models are also available, which include economics, evolutionary biochemistry and games.

An example of a NetLogo interface is shown below.

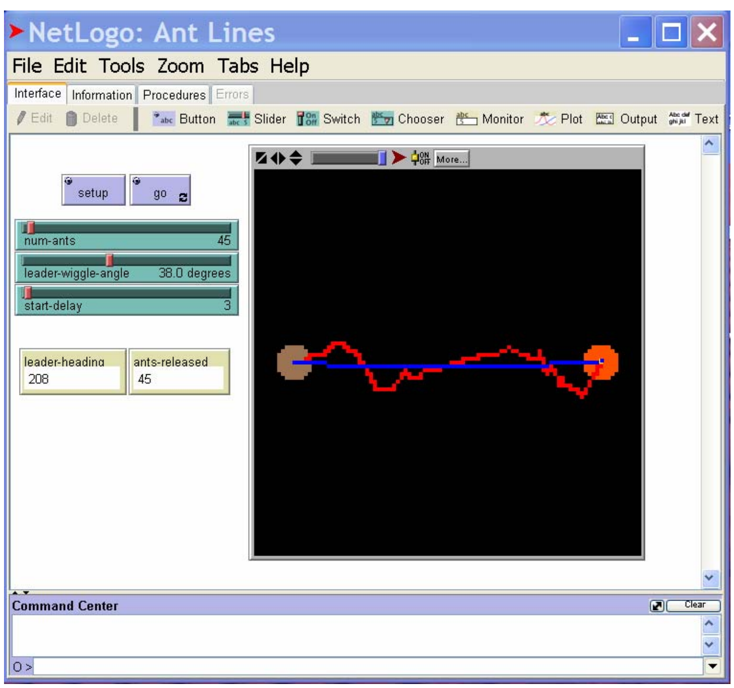

Figure 6: Example NetLogo interface

4.8.2 REPAST

Repast is a more sophisticated Java-based simulation environment that forces the developer to provide Java classes in order to create an application. From the Repast web site, http://repast.sourceforge.net/:

> The Recursive Porous Agent Simulation Toolkit (Repast) is one of several agent modeling toolkits that are available. Repast borrows many concepts from the Swarm agent-based modeling toolkit [1]. Repast is differentiated from Swarm since Repast has multiple pure implementations in several languages and built-in adaptive features such as genetic algorithms and regression. For reviews of Swarm, Repast, and other agent-modeling toolkits, see the survey by Serenko and Detlor, the survey by Gilbert and Bankes, and the toolkit review by Tobias and Hofmann [2] [3] [4]. In particular, Tobias and Hofmann performed a review of sixteen agent modeling toolkits and found that "we can conclude with great certainty that according to the available information, Repast is at the moment the most suitable simulation framework for the applied modeling of social interventions based on theories and data" [4].

Of particular interest is the built-in support for genetic algorithms (which can be used to evolve controllers for robot swarms, for example) and sophisticated modelling neighbourhoods. Repast is widely used for social simulation and models in crowd dynamics, economics and policy making among others have been constructed. The tutorial link http://complexityworkshop.com/cw/tutorial/RePast/index.html provides most of the information required to create simple simulations. An example user interface for the SugarScape model due to Axtell and Eppstein that provides fairly sophisticated instrumentation and data gathering capabilities is shown in Figure 7.

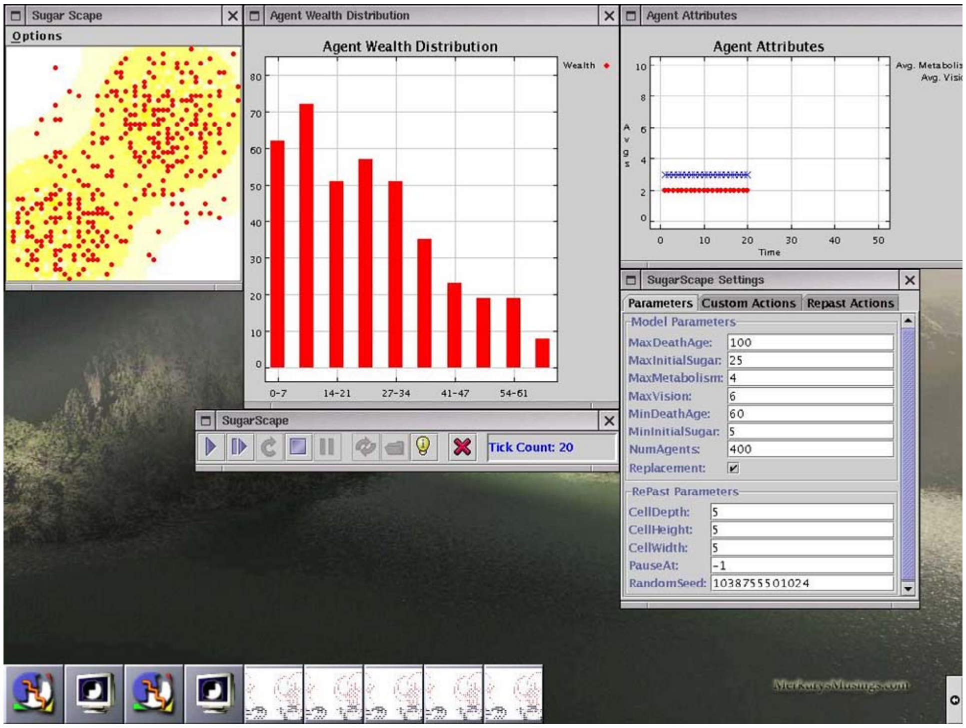

Figure 7: Example Repast Interface

4.9 MODELS OF STIGMERGIC SYSTEMS

This section provides details of several stigmergic systems that have been examined in a research setting and exploited in various industrial applications. Applications of these models are described in the section on applications.

4.9.1 FORAGING

4.9.1.1 ANT FORAGING

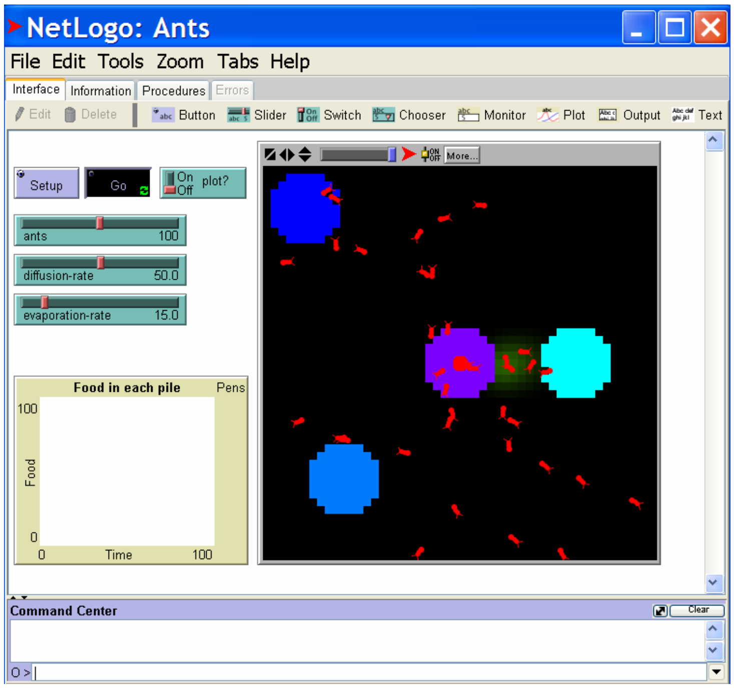

Figure 8: Start of Foraging

Figure 8 shows a nest (centre of the display) with 3 potential food sources. The picture shows ants leaving the nest and performing a random walk in the plane. In this model, ants lay down pheromone trails as they return to the nest, which they do when they have discovered a food source. The pheromone trail both diffuses and evaporates in this model. Evaporation ensures that the pheromone trails to depleted or exhausted food sources will eventually disappear and ants will not visit these sites. Diffusion ensures that ants wandering in the plane will eventually pick up a scent and can use gradient following in order to follow the trail to the food source. Pheromone trails in the foraging figures are shown in green to white, where white represents a very strong trail.

Figure 9 shows a well-established foraging pattern. Food source 1 is almost depleted, while a trail is beginning to form from source 2.

Figure 10 shows ants with well-established trails to food sources 2 and 3, with food source 2 being depleted at a faster rate. Note that foraging still occurs elsewhere in the plane; i.e. not all ants are employed in bringing back food.

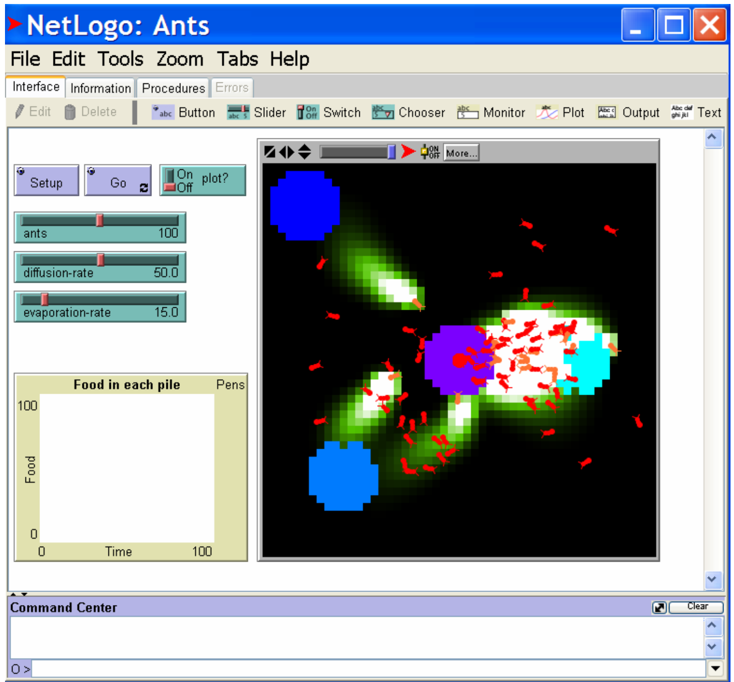

Figure 9: Food source 1 almost depleted

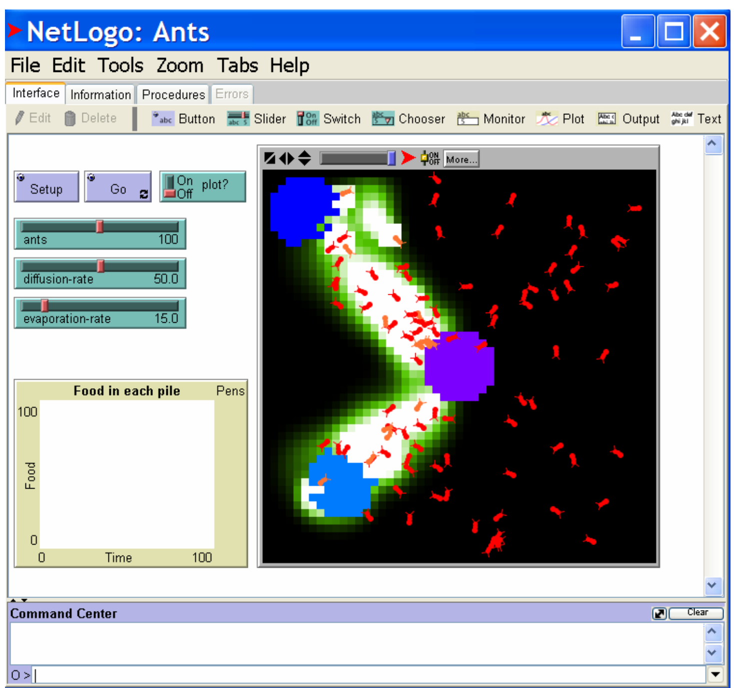

Figure 10: Food sources 2 and 3 being exploited

Once the colony finishes collecting the closest food, the chemical trail to that food naturally disappears, freeing up ants to help collect the other food sources. The more distant food sources require a larger "critical number" of ants to form a stable trail.

The ant colony generally exploits the food source in order, starting with the food closest to the nest, and finishing with the food most distant from the nest. It is more difficult for the ants to form a stable trail to the more distant food, since the chemical trail has more time to evaporate and diffuse before being reinforced. Variations on this characteristic behaviour are possible if the amount of pheromone dropped reflects the quality of the food source.

Trail laying clearly demonstrates a recruitment process. Once a food source has been found, other ants quickly follow the trail to the source and, in turn, enhance the trail. This is an example of an autocatalytic process.

While foraging in this example is represented by food, it could equally well be represented by quality of information.

The model shown in the above figures is included with the NetLogo models library. This marker-based stigmergy model can be used for target acquisition and tracking. This is further described in section 5.8.1.

4.9.1.2 Raid Army Ant Foraging

Raid army ant foraging is considerably different from simple ant foraging. In a raid army ant system ants lay pheromone trails both to and from the ants' nest, with the outward concentration (1 unit) being somewhat smaller than the inbound concentration (10 units). Raid army ants make two decisions. The first is whether to move or not. This is determined by p_m, as shown in the equation below.

$$p_m = \frac{1}{2}\left[1 + \tanh\left(\frac{\ell_l + \ell_r}{100} - 1\right)\right]$$

Here, λ_l and λ_r represent the concentrations of pheromone to the left and right of the ant respectively. Having chosen to move, the direction of movement is decided based upon the equation:

$$p_l = \frac{(5 + \ell_l)^2}{(5 + \ell_l)^2 + (5 + \ell_r)^2}$$

Here, p_l represents the probability of moving to the left. The probability of moving to the right is given by 1- p_l. The constants 5 and 2 are often more generally represented by k and n respectively. A wide range of raid structures can be generated by varying n and k; however, the raid front is a remarkably stable structure across a wide range of values.

This model reproduces the models of army ant foraging developed by Deneubourg et al. (1989. The blind leading the blind: modeling chemically mediated army ant raid patterns. J.

Insect Behav., 2, 719-725) and the extension of this model analyzed by Sole et al in 2000 (Pattern formation and optimization in army ant raids. Artificial Life, 6(3), 219-226).

The characteristic raid front is shown in Figure 11. In this figure we see that the ants are capable of creating a wide front while foraging. This is particularly effective at clearing a path through a region and is quite apparent that the ants are working as teams. Looking closely at the figure we see that beyond the raid front there is also structure in the trails that lead back to the nest. These trails have value too in that they represent regions of the space which have been searched; i.e. their contents are known. In a military scenario these trails have value in that they represent "safe" or known threats.

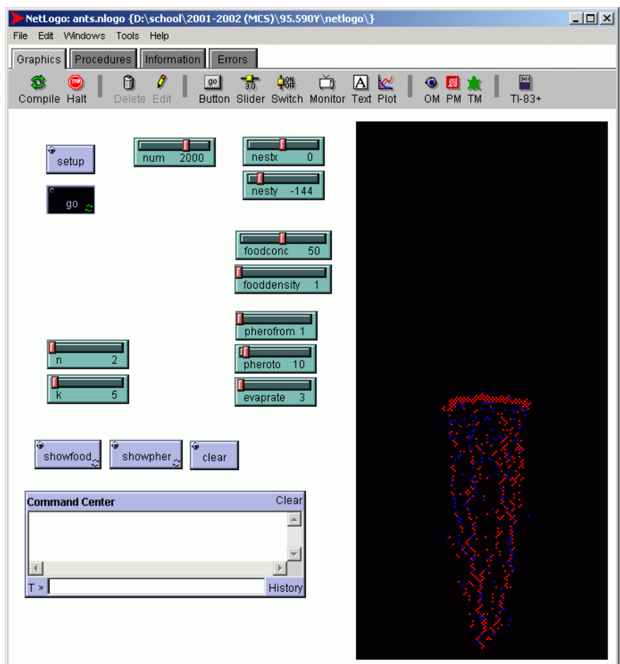

Figure 11: Raid Army Ant Foraging

It has been hypothesized that raid army networks represent optimal distribution networks; however, this remains a conjecture.

The model shown in the above figure was created at Carleton University. However, a more sophisticated model written by Tim Brown as part of his Ph.D. research can be found at: http://www.infiniteworld.org/ant/model/. The background information and research goals on this site are interesting in that they discuss several issues of military importance; such as allocation and distribution of individuals to achieve a particular goal and how teams can be dynamically formed. His third goal, reproduced here, is particularly relevant:

> *A well designed computer model of army ant swarm behaviour which incorporates real-world measures of efficiency provides a powerful tool for*

exploring key questions in multi-agent system design and collective intelligence. In particular, one can perform precise sensitivity testing to examine how specific parameters influence the ant swarm's ability to solve their collective goal of efficiently exploring the environment. The relative importance of communication rate, network size (number of ants), task fidelity and task specialization will be examined in detail.

4.9.2 DIVISION OF LABOUR AND TASK ALLOCATION

Stigmergy is used extensively in determining how many agents are required to undertake a particular task and what part individual agents play in it. Division of labour and task allocation algorithms use both marker-based and sematectonic forms of stigmergy.

The algorithm Bonabeau [2] suggests a model of task specialization based upon a model insect division of labour; it is designed to model behavioural castes, also referred to as behavioural roles From an initially homogenous set of individuals, the result of the algorithm is to end up with a heterogeneous set of individuals, each member of which is specialized to a specific task.

In order to model this problem, each individual has a certain threshold for working on a task, as well as a *stimulus* for doing that task. The stimulus is the stigmergic signal in this system. The threshold lowers when they engage in that task (or learn it) and rises when they're not doing that task (forgetting it). Depending on the threshold value, the individual can have a greater or lessened probability of responding to the exact same level of stimulus.

The idea behind the algorithm is that individuals with more experience, and which are thus better equipped to handle a specific task, are more inclined to undertake that task than individuals who have less experience with that task.

The probability of an individual *i* undertaking a task *j* is expressed as:

$$T_{\theta i,j}(s_j) = \frac{s_{i,j}^2}{s_{i,j}^2 + \alpha\theta_{i,j}^2 + \beta d_{i,j}^2}$$

Where $\theta_{i,j}$ is the self-reinforcing threshold for individual *i*, task *j*, and $d_{i,j}$ is the distance from individual *i* to where task *j* is performed. α and β are tuning coefficients, which are often set to 1. Whenever individual *i* is performing task *j*, the self-reinforcing equation is:

$$\theta_{i,j} \leftarrow \theta_{i,j} - \xi\Delta t$$

Whenever individual *i* is not performing task *j*, the self-reinforcing equation is:

$$\theta_{i,j} \leftarrow \theta_{i,j} + \varphi\Delta t$$

The value of $\theta_{i,j}$ is restricted to between 0 and a maximum value, typically 1.

As the individual performs one task more than others, this causes the threshold for that task to drop, while the thresholds for other tasks increase. Since the probability function is based on the threshold, a lower threshold means a greater tendency to perform that task, reinforcing the selection of that task, thereby reinforcing the behaviour.

The use of a distance in the equation for T allows a higher probability to those individuals that are closer to the task performance location. Using the above model, individuals can specialize in particular tasks over time. Systems employing these algorithms are also capable of responding to the failure of specialized agents as other agents will take over once a stimulus gets high enough.

This model has been successfully applied by Cicirello [183] to a dynamic, distributed factory scheduling scenario where jobs have to be scheduled on particular machines. White [184] has applied the same principles to robotic soccer where soccerbot roles are dynamically assigned rather than being static. This later usage of the algorithm is particularly pertinent to the military in that it raises the possibility that unmanned autonomously vehicles could be assigned roles dynamically as the battlefield scenario unfolds.

4.9.3 SORTING AND CLUSTERING

Sorting and clustering in ants is achieved with simple sematectonic stigmergy. Essentially, ants wander in a plane being able to perceive the local density of classes of object. Their behaviour is quite simple; they either pick up objects with a given probability based upon their perception of object density in the region if they are not carrying anything or if carrying something, they drop if based upon a perception of density. Mathematically this can be stated as:

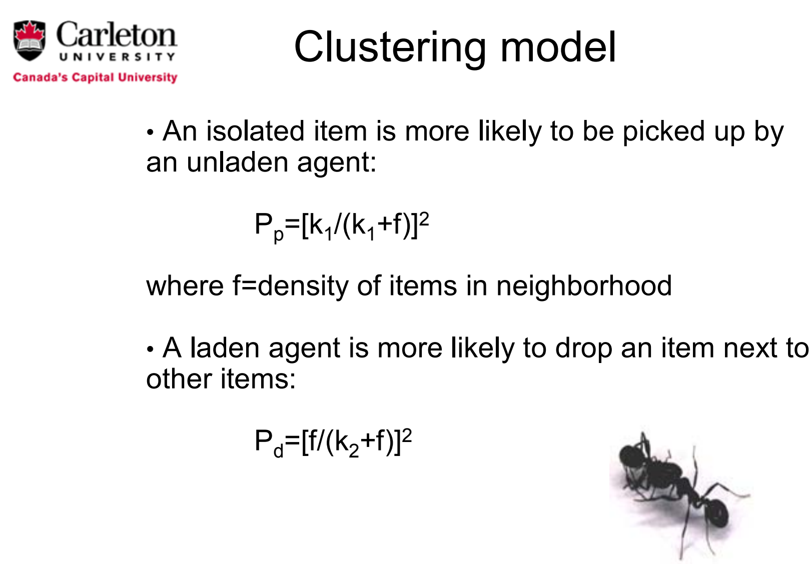

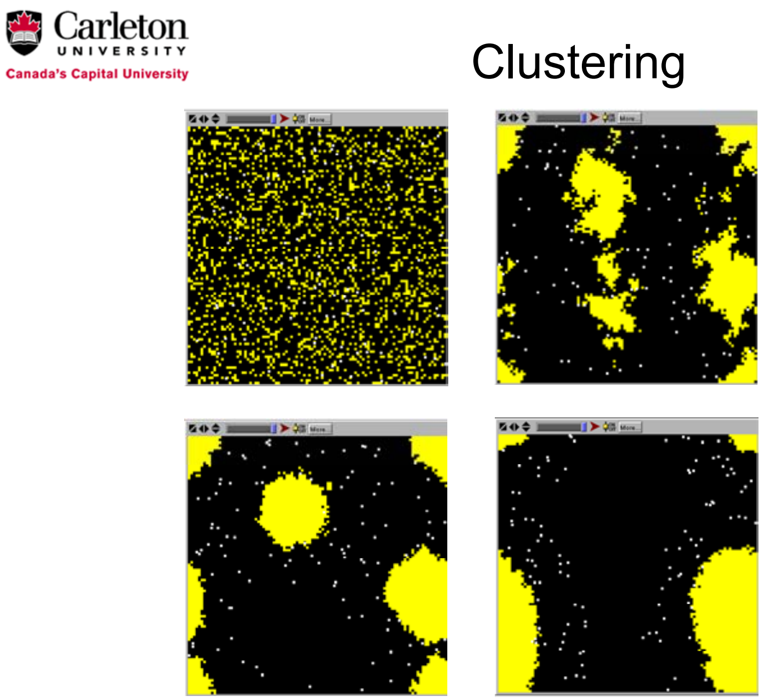

Figure 12: Clustering using Termites

The figure above shows the time evolution of the mathematical model shown on the previous page. The reader should note that the 2 dimensional grid is toroidal which implies that a single pile has emerged in the bottom-right snapshot.

An extension of the model to multiple classes of object can be described as: The same principle can be applied to sort items of several types (i=1,...,n), f is replaced by f_i, the fraction of type i items in the agent's neighborhood:

$$P_p(i)=[k_1/(k_1+f_i)]^2$$

$$P_d(i)=[\, f_i \,/(k_2+f_i)]^2$$

The value of this model from a military perspective is two-fold. First, the model can be used literally to accumulate items in a single location that does not have to be communicated to any of the participating agents. This is an advantage from a security perspective.

Secondly, in concept space, this algorithm can be used to determine useful relations between pieces of information. A number of applications using this approach have been reported.

Finally, robot swarms have been programmed using the above algorithms to perform sorting and clustering.

4.9.4 NEST BUILDING

Building structures using distributed, stigmergic algorithms is a hard problem. In human-controlled structure building algorithms, one individual controls the process of construction and plans are drawn up prior to construction. Construction is a mainly sequential process, although certain phases allow for some parallelism.

Taking wasps as an example, hive construction is a distributed process. The process uses sematectonic stigmergy. Wasps recognize patterns in the structure that is being built and augment it with new components. In essence, the wasp has a small number of rules of the form, “if I see a 3 dimensional pattern of cells then I should add a new cell at a particular point”.

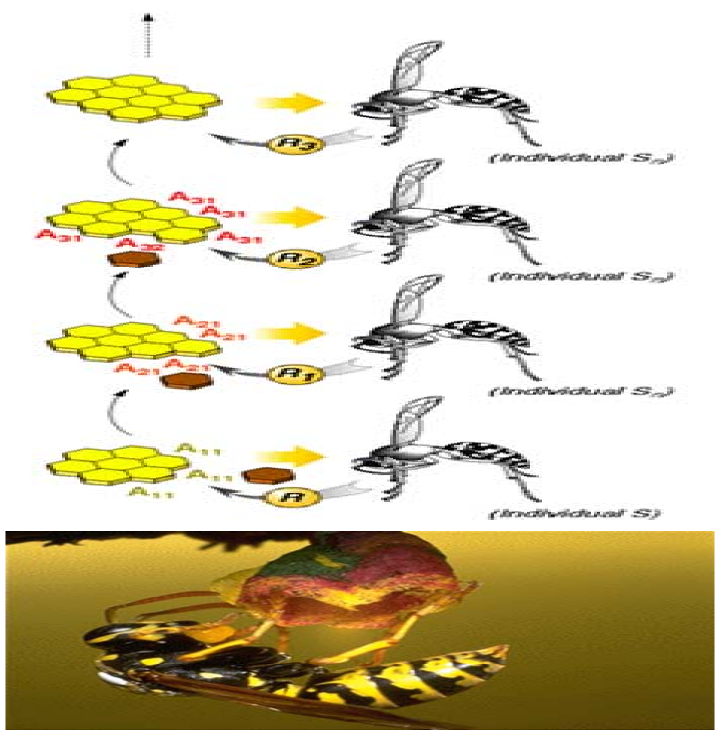

Figure 13: Nest building

The figure above demonstrates the stigmergic mechanism of nest building. A pattern is perceived by an individual wasp and it adds a new cell in the appropriate place, thereby changing the configuration of cells. Another wasp then sees the changed configuration, recognizes the new pattern and adds another cell. This process continues until a space filling structure has been created or no further additions of cells are possible; i.e. no pattern-action rules match.

Building model

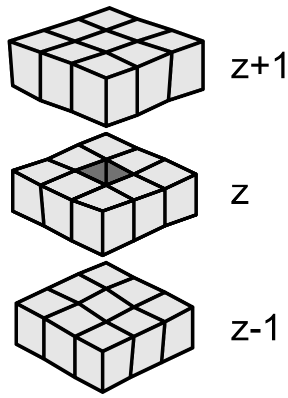

- Agents move randomly on a 3D grid of sites.
- An agent deposits a brick every time it finds a stimulating configuration.
- Rule table contains all such configurations. A rule table defines an algorithm.
- Rule space is very large.

Figure 14: Model for building hive

The figure above highlights the essential characteristics of the process. The wasp “sees” in 3 dimensions, being able to sense a total of 26 cells. A cell either is present or not. This pattern of 26 ones or zeros may match a rule that says create a new cell in position 15. Building is asynchronous, with no central control. As the figure above indicates, the possible rule space is extremely large – genetic algorithms have been used to search for viable rule sets.

This stigmergic system can be used as a model for the construction of structures using relatively simple agents. It does not rely on steps being pre-ordered and each agent is capable of completing the entire structure. Therefore, individual agents may fail but the structure can still be completed.

NASA has used these principles to demonstrate how space stations of the future could be constructed. It would seem to be the case that military structures could be constructed in a similar way.

Engineered emergent patterns

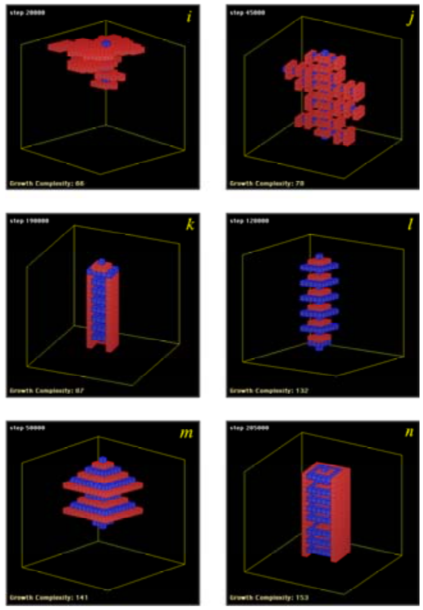

Figure 15: Emergent Structures

The figure above shows a number of example structures generated using the stigmergic principles above. The bottom 2 structures are candidate examples for space structures.

4.9.5 FLOCKING

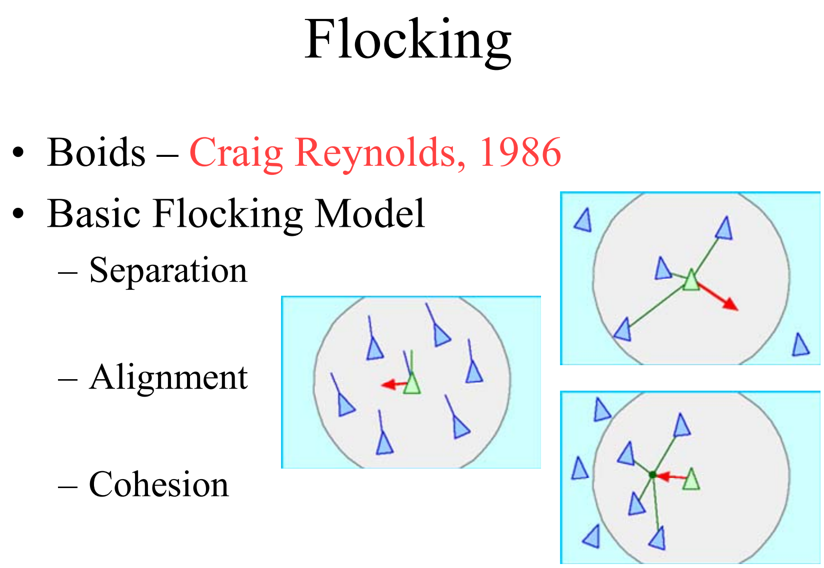

Figure 16: Principles of Flocking

The figure above demonstrates the essential principles of flocking as described by Reynolds [7]. Emergent group control of a collection of birds – Reynolds called them boids -- can be achieved by consideration of 3 independent effects. The first – separation – ensures that birds remain a discrete distance away from one another. The goal here is to avoid collisions. Stigmergy in this system is represented by the birds themselves and their relative positions and velocities. The second effect is that of alignment – the birds try and move with the same average velocity. Finally, the birds are cohesive in that they attempt to move towards the average position of the local group. This last point – locality – should be stressed here. The birds only look at a small number of birds nearby.

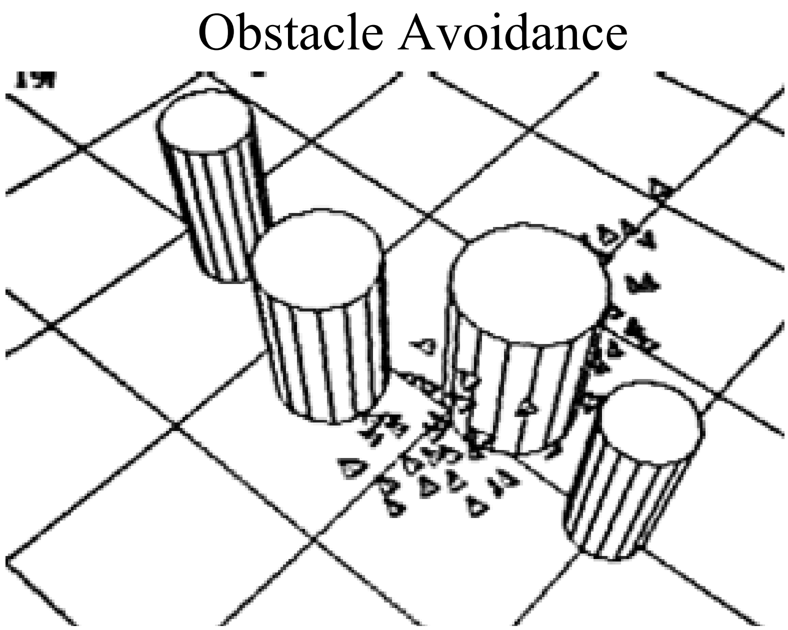

Figure 17: Elements of Flocking

The figure above demonstrates the effectiveness of the above effects. Even in the presence of complex obstacles coherent flight is observed. There is no communication required between the boids in order to observe this emergent behaviour.

Clearly, this system has applications in the area of coordination of groups of unmanned autonomous vehicles. While the above applies to vehicles moving in 3 dimensions, similar algorithms have been developed for 2 dimensions. NASA has been active in this area and proposes to use algorithms of this type for deep space exploration using multiple, small spacecraft.

Couzin [182] has recently discovered that in a heterogeneous collection of boids a small number of leaders can cause the collective to move in a specific direction. This further supports the view that flocking algorithms can be used to control groups of unmanned autonomous vehicles.

Finally, the question of how to infiltrate and disturb swarm systems was raised earlier in the report. An interesting extension to the flocking model available in the models library provided with the Netlogo distribution allows a user to set a level of renegade behaviour. Renegades are boids that *appear* to adhere to the rules of behaviour but sometimes do not. It can be shown that with appropriate levels of renegade behaviour flocking can be disrupted.

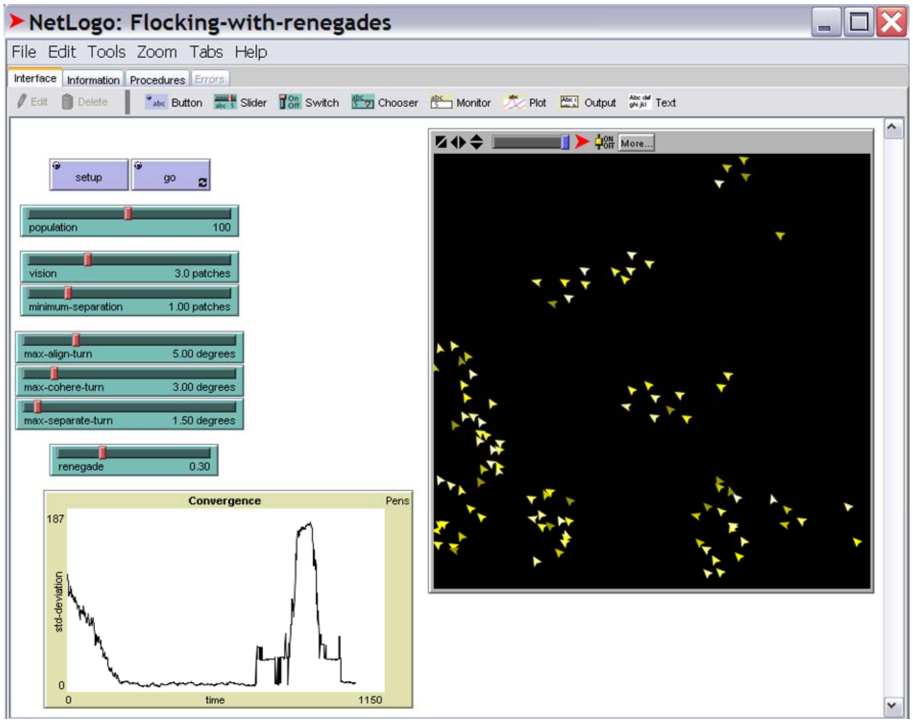

Figure 18: Flocking with some renegades

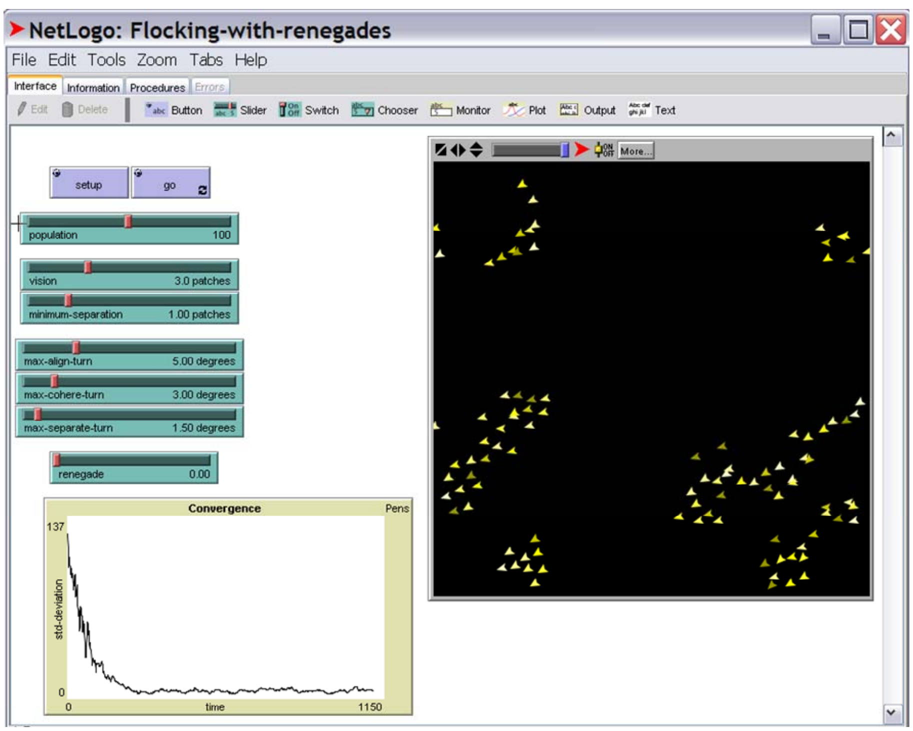

Figure 19: Flocking only

Comparing the convergence graphs in Figure 18 and Figure 19 shows that in pure flocking, the boids converge to motion in a given direction whereas in a flock with renegades convergence occurs but then periodic catastrophic changes in direction occur. The period is not shown in the first figure but can be reproduced.

This is an encouraging result in that it appears to imply that behaviourally similar agents can be introduced into a swarm to disrupt its emergent behaviour. While this provides anecdotal evidence, a comprehensive study should be undertaken to evaluate the “cuckoo effect”.

4.9.6 SUMMARY

The 5 patterns (or models) described in the previous sections represent examples of stigmergic systems that use either marker-based or sematectonic stigmergy. They are *not* a comprehensive set of examples, such a description would require an extended analysis going far beyond the scope of this report. However, the table below provides several other examples of stigmergic patterns observed in nature.

Table 3: Stigmergic Patterns in Nature

Swarm Behaviour	**Entities**
Pattern Generation	Bacteria, Slime mold
Path Formation	Ants
Nest Sorting	Ants
Cooperative Transport	Ants
Food Source Selection	Ants, Bees
Thermoregulation	Bees
Task Allocation	Wasps
Hive Construction	Bees, Wasps, Hornets, Termites
Synchronization	Fire Flies
Feeding Aggregation	Bark Beetles
Web Construction	Spiders
Schooling	Fish

Flocking	Birds
Prey Surrounding	Wolves

5 APPLICATIONS OF SWARM INTELLIGENCE

5.1 ANT COLONY OPTIMIZATION

Ant algorithms (also known as Ant Colony Optimization) are a class of metaheuristic search algorithms that have been successfully applied to solving NP hard problems [159]. Ant algorithms are biologically inspired from the behaviour of colonies of real ants, and in particular how they forage for food. One of the main ideas behind this approach is that the ants can communicate with one another through indirect means (stigmergy) by making modifications to the concentration of highly volatile chemicals called pheromones in their immediate environment.

The Traveling Salesman Problem (TSP) is an NP complete problem addressed by the optimization community having been the target of considerable research [164]. The TSP is recognized as an easily understood, hard optimization problem of finding the shortest circuit of a set of cities starting from one city, visiting each other city exactly once, and returning to the start city again. The TSP is often used to test new, promising optimization heuristics. Formally, the TSP is the problem of finding the shortest Hamiltonian circuit of a set of nodes. There are two classes of TSP problem: symmetric TSP, and asymmetric TSP (ATSP). The difference between the two classes is that with symmetric TSP the distance between two cities is the same regardless of the direction you travel; with ATSP this is not necessarily the case.

Ant Colony Optimization has been successfully applied to both classes of TSP with good, and often excellent, results. The ACO algorithm skeleton for TSP is as follows [164]:

```
procedure ACO algorithm for TSPs
  Set parameters, initialize pheromone trails
  while (termination condition not met) do
    ConstructSolutions
    ApplyLocalSearch % optional
    UpdateTrails
  end
end ACO algorithm for TSPs
```

5.1.1 ANT SYSTEM (AS)

Ant System was the earliest implementation of Ant Colony Optimization metaheuristic. The implementation is built on top of the ACO algorithm skeleton shown above. A brief description of the algorithm follows. For a comprehensive description of the algorithm, see [158], [159], [160] or [164].

5.1.1.1 ALGORITHM

Expanding upon the algorithm above, an ACO consists of two main sections: *initialization* and a *main loop*. The main loop runs for a user-defined number of iterations. These are described below:

Initialization

- Any initial parameters are loaded.
- Each of the roads is set with an initial pheromone value.
- Each ant is individually placed on a random city.

Main loop begins

Construct Solution

- Each ant constructs a tour by successively applying the probabilistic choice function and randomly selecting a city it has not yet visited until each city has been visited exactly once.

$$p_{ij}^{k}(t) = \frac{[\tau_{ij}(t)]^{\alpha} \cdot [\eta_{ij}]^{\beta}}{\sum_{l \in N_i^k} [\tau_{il}(t)]^{\alpha} \cdot [\eta_{il}]^{\beta}}$$

- The probabilistic function, $p_{ij}^{k}(t)$, is designed to favour the selection of a road that has a high pheromone value, τ, and high visibility value, η, which is given by: $1/d_{ij}$, where d_{ij} is the distance to the city. The pheromone scaling factor, α, and visibility scaling factor, β, are parameters used to tune the relative importance of pheromone and road length in selecting the next city.

Apply Local Search

- Not used in Ant System, but is used in several variations of the TSP problem where 2-opt or 3-opt local optimizers [164] are used.

Best Tour check

- For each ant, calculate the length of the ant's tour and compare to the best tour's length. If there is an improvement, update it.

Update Trails

- Evaporate a fixed proportion of the pheromone on each road.
- For each ant perform the "ant-cycle" pheromone update.
- Reinforce the best tour with a set number of "elitist ants" performing the "ant-cycle" pheromone update.

In the original investigation of Ant System algorithms, there were three versions of Ant System that differed in how and when they laid pheromone. They are:

- “Ant-density” updates the pheromone on a road traveled with a fixed amount after every step.
- “Ant-quantity” updates the pheromone on a road traveled with an amount proportional to the inverse of the length of the road after every step.
- “Ant-cycle” first completed the tour and then updates each road used with an amount proportional to the inverse of the total length of the tour.

Of the three approaches “Ant-cycle” was found to produce the best results and subsequently receives the most attention. It will be used for the remainder of this paper.

Main Loop Ends

Output

- The best tour found is returned as the output of the problem.

5.1.1.2 DISCUSSION

Ant System in general has been identified as having several good properties related to directed exploration of the problem space without getting trapped in local minima [158]. The current state of the art is described in [159]. The initial form of AS did not make use of elitist ants and did not direct the search as well as it might.

The addition of elitist ants was found to improve ant capabilities for finding better tours in fewer iterations of the algorithm, by highlighting the best tour. However, by using elitist ants to reinforce the best tour the problem now takes advantage of global data with the additional problem of deciding on how many elitist ants to use. If too many elitist ants are used the algorithm can easily become trapped in local minima [158], [160]. This represents the dilemma of exploitation versus exploration that is present in most optimization algorithms. While the ant foraging behaviour on which the Ant System is based has no central control or global information on which to draw, the use of global best information in the Elitest form of the Ant System represents a significant departure from the purely distributed nature of ant-based foraging. Use of global information presents a significant barrier to *fully* distributed implementations of Ant System algorithms in a live network, for example. This observation motivated the development of a fully distributed algorithm – the Ant System Local Best Tour (AS-LBT) [165].

There have been a number of improvements to the original Ant System algorithm. They have focused on two main areas of improvement [164]. First, they more strongly exploit the globally best solution found. Second, they make use of a fast local search algorithm like 2-opt, 3-opt, or the Lin-Kernighan heuristic to improve the solutions found by the ants.

The algorithm improvements to Ant System have produced some of the highest quality solutions when applied to the TSP and other NP complete (or NP hard) problems [158], [159]. Applications to vehicle routing problems, quadratic assignment problems, job shop scheduling, graph colouring and several other areas have been documented in the literature. Design of ant-based algorithms in these application areas requires the designer to develop a heuristic for the visibility function, η. Clearly, for the TSP this is simply $1/d_{ij}$, the distance between the i^{th} and j^{th} cities.

5.1.1.3 Potential Applications of Military significance

Ant search has been used in an industrial setting by the Icosystem Corporation. They have applied sophisticated variants of the algorithm to perform schedule optimization for a large US airline. It would seem that similar algorithms could be used for logistical optimizations in military organizations.

5.1.1.4 ACO Tools

Several implementations of ACO metaheuristics for various problems can be found at http://iridia.ulb.ac.be/~mdorigo/ACO/aco-code/public-software.html. Dr. White also has several implementations in Java or C.

5.2 ROUTING

Readers only interested in an outline of the marker-based stigmergy approaches to routing (and not the details of ad hoc or sensor network approaches) should read sections 5.2, 5.2.2.1, 5.2.2.2 and 5.2.12 only.

Given the increasing importance of sensor networks to the military, routing has provided a fertile research area for stigmergic solutions. In the solution examined in this section, stigmergy is marker based, similar in concept to the pheromone-based foraging of ants. A large number of research papers are described here owing to this author's belief as to the importance of sensor networks in future military conflicts.

Routing has been a significant area of research for swarm intelligence. Starting with Schonderwoerd in 1997, and Di Caro in 1998, the exploitation of the foraging behaviour of ants has been shown to significantly improve the quality of routing in networks. Most recently, research into ad hoc network routing has been active; with Di Caro (AntHocNet) having provided the most compelling research.

Ad hoc networks consist of autonomous self-organized nodes. Nodes use a wireless medium for communication, thus two nodes can communicate directly if and only if they are within each other's transmission radius. Examples are sensor networks (attached to a monitoring station), rooftop networks (for wireless Internet access), and conference and rescue scenarios for ad hoc networks, possibly mobile. In a routing task, a message is sent from a source to a destination node in a given network. Two nodes normally communicate via other nodes in a multi-hop fashion. Swarm intelligence follows the behaviour of cooperative ants in order to solve hard static and dynamic optimization problems. Ants leave pheromone trails at nodes or edges which increases the likelihood of other ants to follow these trails. Routing paths are then found dynamically on the fly, using this so called notion of stigmergy. In this article we survey existing literature on swarm intelligence based solutions for routing in ad hoc networks. We identified 13 different methods, covering non-position and position based approaches, flooding and path based search methods. Some of the articles consider related problems such as multicasting or data centric routing. All of the articles were published after 2001. The ideas coming from existing swarm intelligence based routing in communication networks are incorporated into the wireless domain, with some new techniques which are typical for the wireless domain (such as flooding, use of position, monitoring traffic at neighbouring nodes) being incorporated. We observed that the experimental data provided by these articles is insufficient to make a firm conclusion about scenarios which show the advantages of the proposed swarm intelligence based methods with respect to other existing methods.

5.2.1 INTRODUCTION

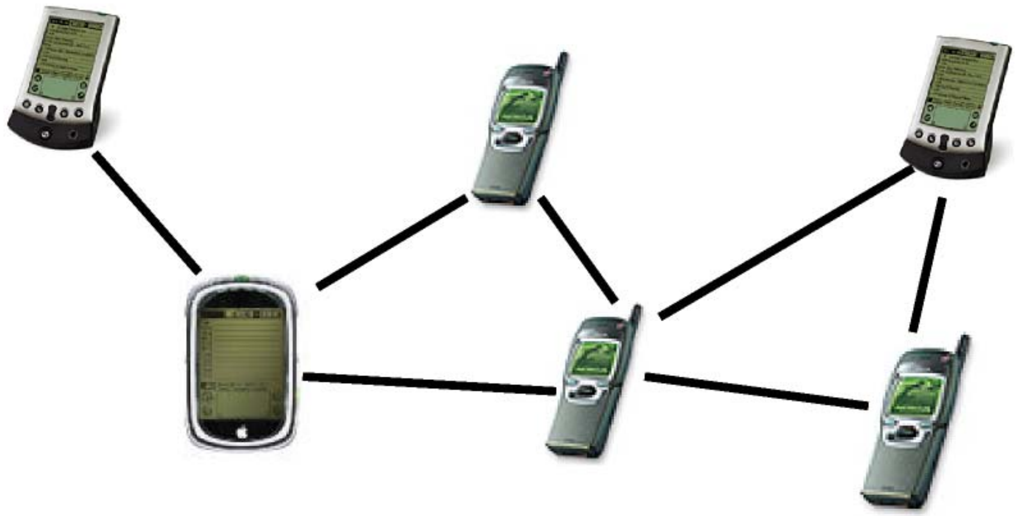

Figure 20: Self-organized ad hoc wireless network

In ad hoc wireless networks, nodes are self-organized and use wireless links for communication between themselves. Ad hoc networks are dynamically created. Examples are conference, battlefield, rescue scenarios, sensor networks placed in an area to monitor the environment, mesh networks for wireless Internet access etc. Nodes in ad hoc networks can be mobile in many scenarios, or mostly static in other scenarios, as in sensor networks. Nodes may decide to go to sleep mode to preserve energy, and wake up later to rejoin the network. Routing solutions must address the nature of the network, and aim at minimizing control traffic, to preserve both bandwidth and energy at nodes. Ant colony based algorithms use a number of control traffic, or existing traffic, sets of information to create best routes. It is a challenging task to discover good routes with controlled traffic, so that overall the swarm intelligence approach outperforms existing routing protocols for ad hoc networks.

Swarm intelligence is a set of methods to solve hard static and dynamic optimization problems using cooperative agents, usually called ants. Ant inspired routing algorithms were developed and tested by British Telecomm and NTT for both fixed and cellular networks with superior results [BH, DD, BHGGKT, SHB, WP]. AntNet, a particular such algorithm, was tested in routing for communication networks [DD]. The algorithm performed better than OSPF, asynchronous distributed Bellman-Ford with dynamic metrics, shortest path with a dynamic cost metric, the Q-R algorithm and predictive Q-R algorithm [BH, DD, BHGGKT, SHB, WP].

This section will review the literature on swarm intelligence based solutions for routing in ad hoc networks. After an extensive search on http://citeseer.nj.nec.com/cs and www.google.com, 13 different relevant articles (two of the articles were published twice, so the total count is 15) were found. They are all very recent, published in 2001 or later, and they propose some swarm intelligence based routing methods for ad hoc wireless networks. Their list is given in the references section. The goal of the article is to summarize existing solutions, classify them according to assumptions and approaches taken, compare them, report on experimental findings from the article, and to draw some conclusions. It was observed that cross referencing between these articles is poor, which is not surprising since many of them appeared simultaneously and all of them were published within the last two years. Two articles were published in 2001, two were published in 2002, and nine out of these 13 articles were published in 2003. There were some independent discoveries of the same ideas, which was also not surprising. It was observed, however, that a number of summaries of other works given in these articles was incorrect, and that many articles do not clearly state which ideas come from the existing research, and which ideas are new. The approach taken in this article is to first present existing swarm intelligence based methods for routing in communication networks, and existing routing schemes for ad hoc networks (in both cases, we only presented methods that were actually used in the surveyed articles), and then referred to them when ad hoc network scenarios are considered, so that additions and differences between them are underlined.

This section is organized as follows. Section 5.2.2 describes swarm intelligence based routing schemes for communication networks. Section 5.2.3 presents routing schemes for ad hoc networks, which do not use swarm intelligence, and which are adapted in the surveyed articles by adding ants for enhanced performance. Section 5.2.4 summarizes path based routing schemes with swarm intelligence, which are close to the schemes used in communication networks. Section 5.2.5 describes routing schemes which use a wireless medium to flood the ants; therefore each initial ant multiplies into a number of ants in the process, which is a non-traditional understanding of what an ant is. Section 5.2.6 presents solutions which assume that nodes have position information, that is, they know their geographic coordinates. Two related routing problems, multicasting, and data centric routing in sensor networks, are discussed in sections 5.2.7 and 5.2.8.

5.2.2 SWARM INTELLIGENCE FOR ROUTING IN COMMUNICATION NETWORKS

5.2.2.1 General principles

We will first describe general principles in all swarm intelligence based solutions. They are used in all of the described solutions, each with particular details starting from this general

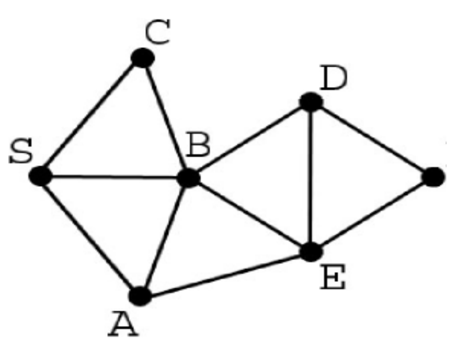

Figure 22: Network

	A	B	C	D	E	F
A	0.9	0.1	0.1	0.4	0.5	0.5
B	0.1	0.8	0.2	0.6	0.4	0.4
C	0.0	0.1	0.7	0.0	0.1	0.1

Figure 21: Routing table for node S

approach. The ants navigate their designated selection of paths while depositing a certain amount of substance called pheromone on the ground, thereby establishing a trail. The idea behind this technique is that the more ants follow a particular trail, the more attractive is that trail for being followed by other ants. They therefore dynamically find a path on the fly, using the explained notion of stigmergy to communicate indirectly amongst themselves. In the case of routing, separate pheromone amounts are considered for each possible destination (that is, on each link pheromone trails are placed in a sequence, one trail for each possible destination). An ant chooses a trail depending on the amount of pheromone deposited on the ground. Each ant compares the amounts of trails (for the selected destination) on each link toward the neighbouring nodes. The larger the concentration of pheromone in a particular trail, the greater the probability of the trail being selected by an ant. The ant then reinforces the selected trail with its own pheromone. The concentration of the pheromone on these links evaporates with time at a certain rate. It is important that the decay rate of pheromone be well tuned to the problem at hand. If pheromone decays too quickly then good solutions will lose their appeal before they can be exploited. If the pheromone decays too slowly, then bad solutions will remain in the system as viable options.

Each node in the network has a routing table which helps it determine where to send the next packet or ant. These routing tables have the neighbours of the node as rows, and all of the other nodes in the network as columns. In Figure 22, we see an example of a network, and in Figure 21 we see the routing table for node *S* in this network.

An ant or message going from node *S* to node *F,* for example, would consider the cells in column *F* to determine the next hop. Ants and messages can determine the next hop in a variety of ways. The next hop can be determined uniformly; which means that any one of the neighbours has an equally likely probability of being chosen. It can be chosen probabilistically, that is, the values in the routing table in column *F* are taken as the likelihoods of being chosen. Taking the highest value in the column of *F* could be another way of choosing the next hop. It could also be chosen randomly, which means choosing uniformly if there is no pheromone present, and taking the highest value if there is. There is also an exploratory way of choosing the next hop, which means taking a route with a value of 0 if one exists.

There are a few swarm intelligence (ant-based) routing algorithms developed for wired networks, and the most well known of which are AntNet [DD] and Ant-Based Control (ABC) [SHB]. The fundamental principle behind both AntNet and ABC is similar – they use ants as exploration agents. These ants are used for traversing the network node to node and updating routing metrics. A routing table is built based on the probability distribution functions derived from the trip times of the routes discovered by the ants. The approaches used in AntNet and ABC are, however, dissimilar – in AntNet, there are *forward* and *backward* ants, whereas in ABC, there is only one kind of ant. Another difference between AntNet and ABC is in the routing front. In ABC, the probabilities of the routing tables are updated as the ants visit the nodes, and are based on the life of the ant at the time of the visit; while in AntNet, the probabilities are only updated when the backward ant visits a node.

5.2.2.2 ANT-BASED CONTROL (ABC) ROUTING

Schoonderwoerd, Holland, and Bruten [SHB] proposed the Ant-Based Control (ABC) scheme for routing in telephone networks. In the ABC routing scheme [SHB], there exist two kinds of routing tasks: exploratory ants which make probabilistic decisions, and actual calls which made deterministic decisions (that is, choosing the link with the most pheromone in the column corresponding to the destination). Exploratory ants are used for source updates. Each source node S issues a number of exploratory ants. Each of these ants goes toward a randomly selected destination D (the ant is deleted when it reaches D). The routing table at each node contains neighbours as rows and all possible destinations as columns, and each entry corresponds to the amount of pheromone on the link towards a particular neighbour for a particular destination. These amounts are normalized in each column (the sum is one), so that they can be used as probabilities for selecting the best link. At each current node C, the entry in the routing table at C corresponding to the source node S is updated. Exploratory ants make the next node choice by generating a random number and using it to select a link based on their probabilities in the routing table. The amount of pheromone left on a trail depends on how well the ant performs. Aging is used to measure performance. In each hop, the delay depends on the amount of spare capacity of the node, and is added to the age. Both ants and calls travel on the same queue. Calls make a deterministic choice of a link with the highest probability, but do not leave any pheromone. The pseudo code of the ABC algorithm is presented below. RT[S][X][Y] is the probability of going from node *S* to node *Y* via node *X.* Referring back to Figure 21, for example, the value of RT[S][A][C] = 0.

Each ant chooses source *S* and destination *D* at random; *C*=*S; T*=0

```
While C ≠ D do {
    Choose next node B using probabilities from RT[C][B][D]:
    Delay = c . exp(–d*sparecapacity(B));
    T ← T + Delay;
    Delta = a/T + b
    // Update the routing table, assuming symmetry
    RT[B][C][S] ← (RT[B][C][S] + Delta)/(1 + Delta)
    RT[B][X][S] ← RT[B][X][S]/(Delta + 1) for X ≠ C
    C=B
}
```

The variables a, b, c and d are parameters with empirically determined values. There is an exploration threshold, g, as well. The threshold g, if crossed determines the next hop uniformly instead of consulting the routing table. This g value is used to ensure that not only one path is used. It is there to make sure that other routes are tried from time to time.

Guerin proposed an all column update enhancement to the ABC scheme. While moving forward, the ABC algorithm only updates routing tables corresponding to source S. Guerin [G] proposed updating the routing tables for all other nodes visited in the route. For example, let the route be: SABCD. In ABC, the routing tables for S are updated at nodes A, B, C and D as an ant moves toward D. The all column update scheme [G] adds updating routing tables for A at B, C and D, routing tables for B at C and D, and routing table for C at D.

5.2.2.3 AntNet and other schemes

In the AntNet scheme [DD], each node periodically sends a forward ant packet to a random destination. The forward ant records its path as well as the time needed to arrive at each intermediate node. The timing information recorded by the forward ant, which is forwarded with the same priority as data traffic, is returned from the destination to the source by means of a high priority backward ant. Each intermediate node updates its routing tables with the information from the backward ant. Routing tables contain per destination next hop biases. This way, faster routes are used with greater likelihood.

Subramaniam, Druschel, and Chen [SDC] described a method which has characteristics of both the AntNet and ABC schemes, and applied it to packet switching networks. Routing tables are probabilistic and are updated as in ABC [SHB]. They [SDC] introduce *uniform ants* that uniformly randomly choose the next node to visit (all neighbours have the same probability of being selected). Ants accumulate cost as they progress through the network. Their method is called Ants-Routing. Only backward exploration is used to update routing tables.

White [W, WP] suggested another routing algorithm for circuit switched networks. The approach is based on three kinds of ants. The first class collects information, the second class allocates network resources based on the collected information and the third class makes allocated resources free after usage.

5.2.2.4 Routing in ad hoc networks without swarm intelligence

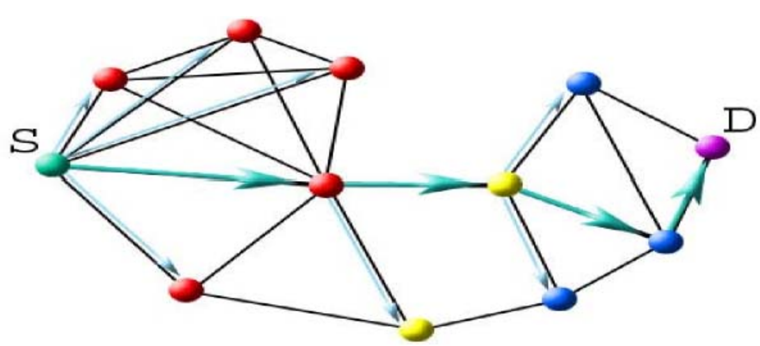

Figure 23: Route discovery from S to D

Routing methods in literature are divided into two groups based on the assumptions made on the availability of position information. There exist non-position and position based approaches. In position based approaches, it is assumed that each node knows its geographic coordinates, the coordinates of its all neighbours, and is somehow informed about the position of the destination. Location based systems have recently been making rapid technological and software advances, and there are cheap solutions with tiny hardware already available. Non-position based solutions assume no knowledge of position information.

5.2.2.5 Non-position based routing

In AODV [PR], the source node floods a route discovery message throughout the network. Each node receiving the message for the first time retransmits it, and ignores further copies of the same message. This method is known as blind flooding. The destination node replies back to the source upon receiving the first copy of the discovery message using the memorized hops of the route. The source node then sends the full message using the recorded path. The method may easily provide multipaths for quality of service, and each node may introduce forwarding delays which may depend on the energy left at the node, or is imposed by a queuing delay. Local route maintenance methods are developed for mobile ad hoc networks. The expanding ring search is also considered to reduce the overhead coming from blind flooding. An adaptive distance vector (ADV) routing algorithm for mobile, ad hoc networks is proposed in [BK], where the amount of proactive activity increases with increasing mobility.

The zone routing protocol [HPS] applies a combination of proactive and reactive routing. Proactive routing is applied for nodes within the same zone, while reactive on-demand routing (such as AODV) is applied if the source and destination are not in the same zone. Within the zone, routes can be proactively maintained using one of several options. One option is to broadcast local topological change within the zone so that shortest paths can be computed. The other option is to periodically exchange routing tables between neighbours, so that each node can refresh its route selection using new information from its neighbours.

5.2.2.6 Position based routing

Finn [F] proposed a position based localized greedy routing method. Each node is assumed to know the position of itself, its neighbours, and the destination. The source node, or node currently holding the message, adopts the greedy principle: choose the successor node that is closest to the destination. The greedy method fails when none of the neighbouring nodes are closer to the destination than the current node. Finn [F] also proposed a recovery scheme from failure: searching all *n*-hop neighbours (nodes at a distance of at most *n* hops from the current node) by limited flooding until a node closer to the destination than C is found, where *n* is a network dependent parameter. The algorithm has nontrivial details and does not guarantee delivery.

5.2.3 PATH BASED ANT ROUTING FOR AD HOC NETWORKS

Our literature review will begin with swarm intelligence based routing methods which do not use the geographic positions of nodes, and which follow the well known traditional definition of an ant, as a single entity that travels through the network, creating a path, possibly travels back to its source, and eventually disappears. There are three protocols described in this category, by Matsuo and Mori [MM] in 2001, Islam, Thulasiraman and Thulasiram [ITT] in April 2003, and by Roth and Wicker [RW] in June 2003. The following section will cover an alternative notion of an ant as an entity that can multiply itself.

5.2.3.1 ACCELERATED ANTS ROUTING

Matsuo and Mori [MM] apparently described the first ant based routing scheme for ad hoc networks, called *accelerated ants routing* in 2001. It appears that it is a straightforward adaptation of a well known scheme for communication networks, with two additions which themselves do not appear to be novel. They followed the Ants-Routing method [SDC] and added a ‘no return’ rule which does not allow ants to select the neighbour where the message came from. They also added an ‘N step backward exploration rule’. This is identical to the all column update scheme proposed by Guerin [G]. In [MM], it is applied when an ant moves backward (and consequently routing entries toward the destination are updated). Performance evaluation showed that the new ants routing algorithm achieves good acceleration for routing table’s convergence with respect to the Ants-Routing method, even if network topology was dynamically changed.

The accelerated ants routing scheme [MM] uses both probabilistic and uniform ants. Uniform ants are important in ad hoc networks because of link instabilities. When a link on a favourite route is broken, uniform ants may quickly establish an alternative route. The whole algorithm is illustrated in the following figures.

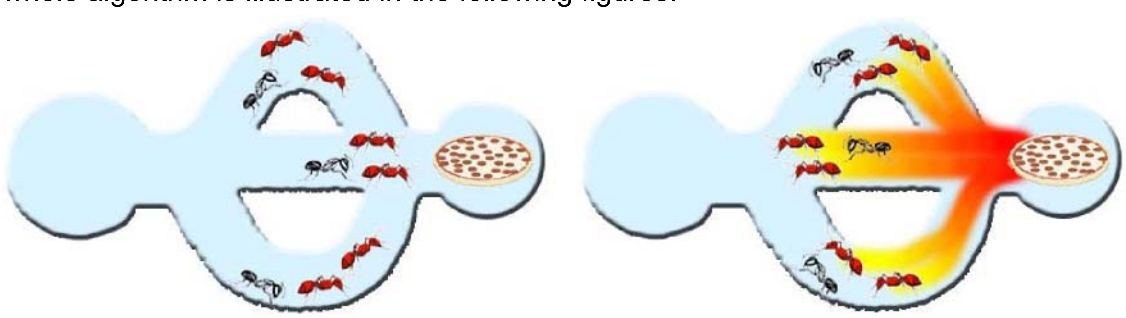

Figure 24: (a) Searching for destination (b) Pheromone leads to destination

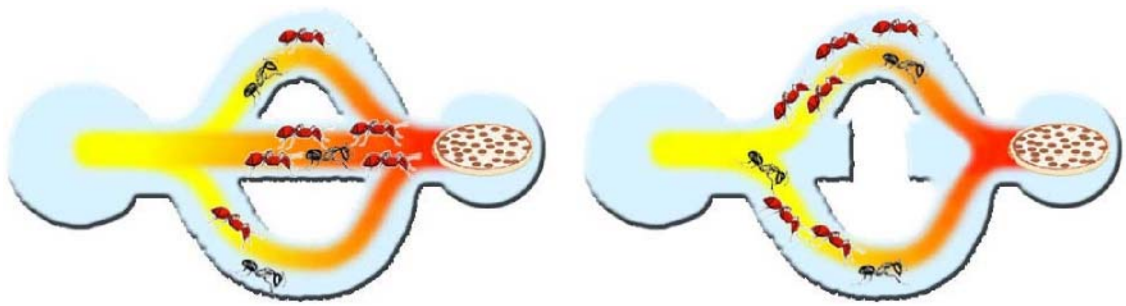

Figure 25: (a) shortest path is most reinforced (b) link is lost

Figure 24a illustrates both the probabilistic (red) and the uniform (black) ants choosing the paths uniformly since there is no pheromone present in the network. Figure 24b shows the returning ants marking the path with pheromone. The path in the middle is the shortest, and therefore has the highest concentration of pheromone. This is why most of the probabilistic ants in Figure 25a follow this trail. Figure 25b shows that the ants will adapt if a path disappears. The top path is shorter than the bottom one; therefore, the probabilistic ants have a higher chance of choosing it.

5.2.3.2 SOURCE UPDATE ROUTING

Islam, Thulasiraman and Thulasiram [ITT] recently proposed an ant colony optimization (ACO) algorithm, called *source update,* for all-pair routing in ad hoc networks. ‘All pair’ routing means that routing tables are created at each node, for all source-destination pairs, in the form of a matrix with neighbours as rows and destinations as columns, so that the table assists in any randomly chosen source-destination pair. The algorithm is claimed to be scalable, but apparently this is with respect to the number of processors on a parallel computer, not the number of nodes in an ad hoc network. The authors also claim that it is an on-demand routing algorithm for ad hoc networks; this is true if ants are launched just before data traffic. They, [ITT], develop a mechanism to detect cycles, and parallelize this algorithm on a distributed memory machine using MPI.

In the source update technique [ITT], each ant memorizes the whole path to its destination and uses it to return back to the source. While the ant is searching for the destination, the routing table updates are performed to form a trail that leads back to the source. During the backward move, updates are made with respect to the selected destination D (with D as the starting point in the route, thus erasing the accumulated weight first), which then in fact serves as the source of the new message, therefore the procedure for the backward move is algorithmically identical to the one used in the forward move. Backward routing is needed so that S finally places some pheromone in its routing table for D. The amount of pheromone placed at each selected edge is not constant in [ITT]. It depends on the weight, which can be a function of distance, transmission time (delay), congestion, interaction time or other metrics ([ITT] used the delay as weight). Note that the amount of new pheromone left on a traveled link is inversely proportional to the cumulative weight from S to the current node, so that longer paths are less enforced. The amount of pheromone in other entries is decreased by a certain fixed percentage. The authors do not normalize the total pheromone count (that is, the sum is not equal to 1), which is done in some traditional approaches such as [SHB]. Comparing several different ants going toward the same destination, longer created paths obviously evaporation more and accumulate less pheromone, and shorter path therefore have a higher chance of be selected.

To memorize the path, ants in [ITT] use a stack data structure containing all of the nodes along the path from S (these nodes are called *stack* nodes). The same stack is used in [ITT] for loop detection and avoidance. This is achieved by ignoring neighbours which are already in the stack when deciding the next hop. Therefore, a loop is never created. If a node has no neighbour which is not already in the stack (such a node becomes a *visited* node), the search backtracks to the previous node. The authors do not discuss the possible reappearance of such visited nodes in the stack later on, which could lead to infinite loops. However, this can be avoided by keeping such nodes in a separate list of visited nodes, so that it does not reappear on the route (and loop creation is avoided). The algorithm, therefore, is a simple *depth-first search* scheme, which the authors [ITT] do not note.

Exploratory ants [ITT] apply the following semi-deterministic scheme when deciding the next node to continue the depth first search with. If there is any link toward *unseen* neighbours (unseen neighbours are nodes which are neither stack nodes nor visited nodes) that also has not yet been tried by any other ants, it is selected (if there are a few such links, one at random is selected). The reason is that the quality of all path candidates needs to be tested. This is important for ad hoc networks, since a newly created edge may provide good quality path. If there is no such unseen node, the ant searches for the next hop by considering the pheromone concentration. It selects the neighbour whose pheromone trail in the column corresponding to destination D is the largest.

The experimental results in [ITT] concentrate only on the parallel implementation for the algorithm, and discuss issues like parallel speed up, scalability with respect to number of processors used, and time versus number of ants. The only comparison is with a basic technique without source update, which is a technique where ants make random decisions at each node, without leaving any pheromone behind. There is no discussion on the impact of various parameters. Since ad hoc networks are self-organized networks where each node makes independent decisions (generally following the pre-agreed protocol), parallel implementations (aiming at speedup optimization), where one processor simulates the work of several nodes from the ad hoc network, do not provide the needed insight into the performance of a particular routing protocol. The insight provided by the authors [ITT] is only on the quality of their parallelization.

Figure 27a and Figure 27b illustrate the source update routing algorithm presented by [ITT]. Figure 27 shows how ants prefer unvisited nodes in their path to the destination. They pick the node with the highest concentration of pheromone if no unvisited nodes exist in their path. The arrows in both figures depict the forward movement of the ants, and the pheromone trails depict the backward movement. In Figure 27, the brown ant was last to move, and it found a path that is shorter than that of its predecessors.

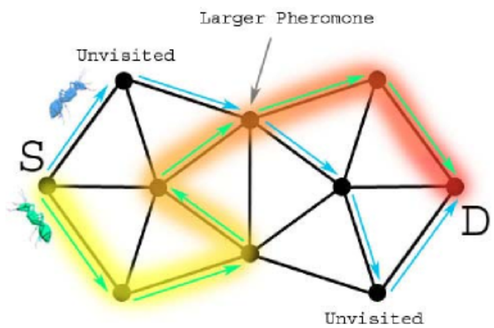

Figure 27a: Second ant begins routing

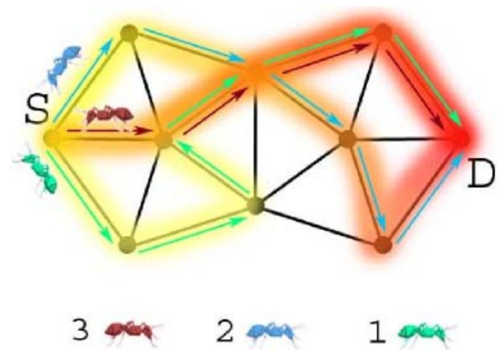

27b: Third ant returns to source

5.2.3.3 RANDOM WALK BASED ROUTE DISCOVERY

Roth and Wicker [RW] presented the scheme called 'Termite' which expands on the ABC algorithm [SHB], but does away with the idea that only specialized packets may update routing tables. In the Termite protocol [RW], data traffic follows the largest pheromone trails, if any exist on any link. If there are no pheromone trails on any link, a route request is performed by a certain number of ants. Each ant performs a random walk over the network. In the random walk, ants and packets uniformly randomly choose their next hop, except for the link they arrived on. During the random walk, pheromone trails with respect to the source are left. If an ant cannot be forwarded, it is dropped. Any number of ant packets may be sent for each route request; the exact number of which may be tuned for a particular environment. An ant is not looking for an explicit route to the destination. Rather it is searching for the beginning of a pheromone trail to the destination. The route will be strengthened by future communications. Once an ant reaches a node containing pheromone to the requested destination, a route reply packet is returned to the requestor. The message is created such that the source of the packet appears to be the requested destination and the destination of the packet is the requestor. The reply packet extends pheromone for the requested destination back to the requestor without any need to change the way in which pheromone is recorded at each node. The reply packet is routed normally through the network probabilistically following a pheromone trail to the requestor. Intermediate nodes on the return path automatically discover the requested node. Hello packets are used to search for neighbours when a node has become isolated. Proactive seed packets are used to actively spread a node's pheromone throughout the network. Seeds make a random walk through the network and serve to advertise a node's existence. They can be useful for reducing the necessary number of explicit route request transactions. All routing decisions in Termite are random. A time to live field is used to prevent propagation of bad routes. The size of the pheromone table may be reduced by implementing a clustering scheme.

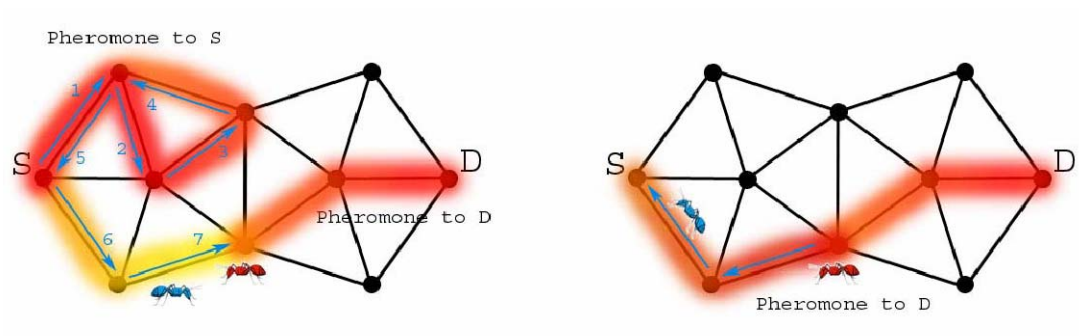

Figure 28: (a) Blue ant searches for trail (b) Blue ant returns to source

Termite can take advantage of the wireless broadcast medium, since it is possible for each node to promiscuously listen to all transmissions. Routing information can be gained from listening to all traffic, rather than only to specifically addressed traffic. New nodes can quickly be detected when their transmissions are overheard. Also, a great deal of information about the network can be gained from the destinations that neighbours are forwarding to. While promiscuity can boost the performance of Termite, it also creates some problems. The same packet overheard a few times shall not be processed more than once, to avoid misleading pheromone gradients. In order to prevent the double counting of packets, a message identification field is included in Termite packets. Another problem is that energy consumption increases when traffic at neighbouring nodes is monitored. Finally, Termite assumes bidirectional links. This article therefore presented a number of novel ideas for ant based routing. However, the experimental data only presented the performance of the Termite protocol, without comparing it with any other routing scheme.

The Termite scheme [RW] differs from the source routing [ITT] by applying pheromone trails or random walks instead of a stack based depth first search. Therefore it allows loops. It differs from the accelerated ants routing [MM] by applying random walk ants rather than uniform or probabilistic ones. Random walk ants differ from uniform ants since they follow pheromone trails, if any. Termite [RW] also does not apply all column updates. Finally, the Termite scheme applies monitoring traffic at neighbouring nodes, which is not present in [MM] and [ITT].

Figure 28 illustrates the Random walk based route discovery algorithm. The red ant in (a) has left a pheromone trail from its current location to destination *D.* The blue ant makes a random walk (labelled by the numbered blue arrows) along the network until it reaches the pheromone trail left by the red ant to the destination. As it searches for a trail to the destination, it leaves a trail which leads back to the source. It then turns around, and lays a second pheromone trail (which leads to the destination) from this node back to the source, as seen in (b). This forms a trail that leads from source to destination.

5.2.4 FLOODING BASED ANT ROUTING

Nearly half, (that is, six out of 13) of the published articles that we surveyed fall into this category. Two such methods are proposed in 2002, by Marwaha, Tham, and Srinivasan [MTS1, MTS2], and by Gunes, Sorges and Bouazizi [GSB]. This later method was improved by Gunes, Kahmer and Bouzazizi [KBB] in June 2003. Baras and Mehta [BM] added a method in March 2003. Eugster [E] derived some formulas for probabilistic guarantees of protocols [GSB] and [MTS2]. Finally, in May 2003, Rajagopalan, Jaikaeo and Shen [RJS] applied flooding in the context of their zonal routing scheme.

5.2.4.1 ANTAODV REACTIVE ROUTING

Marwaha, Tham and Srinivasan [MTS1, MTS2] studied a hybrid approach using both AODV and reactive Ant based exploration. Their technique is called AntAODV. Routing tables in AntAODV are common to both ants and AODV. If the sender node (or node currently holding the message) S has a fresh route toward the destination, it uses it to forward the packet. The authors claim that this is different from AODV which starts route discovery first, but there are modifications of AODV in literature that use fresh routes in the same way. Otherwise (no fresh route available) it will have to keep the data packets in its send buffer until an ant arrives and provides it with a route to that destination. Each ant follows a blind flooding approach and therefore multiplies into several ants. If an ant reaches a node with a fresh route, it stops the advance and converts into a backward ant to report the route to S. Note that again, a similar provision already exists in AODV variations. Ants take a 'no return' rule, meaning that they never return to the node they came from. Overall, it appears that the only difference between AODV and its variants, and AntAODV, is that routing tables are larger, listing all neighbors with their trail amounts for each destination instead of simple routing tables used in AODV, listing only the best choice. This allows a random selection of the next hop, based on pheromone trails. The definition of a fresh route is similar in the two schemes. In the experimental section, comparing a new scheme against AODV (without the mentioned variations), the authors added a proactive component to AntAODV. If no ant visited a node within a certain visit period, the node would generate a new ant and transmit it to one of its neighbours selected randomly. This article does not discuss pheromone trails (that is, what they mean by ‘fresh’ routes) and therefore does not sufficiently underline how the ant based approach really works compared to already existing equivalent AODV variants.

5.2.4.2 ARA REACTIVE ROUTING

Gunes, Sorges and Bouazizi [GSB] presented a detailed routing scheme, called ARA, for MANETs, including route discovery and maintenance mechanisms. Route discovery is achieved by flooding forward ants to the destination while establishing reverse links to the source. Their approach uses ants only for building routes initially and hence is a completely reactive algorithm. A similar mechanism is employed in other reactive routing algorithms such as AODV. Routes are maintained primarily by data packets as they flow through the network. In the case of a route failure, an attempt is made to send the packet over an alternate link. Otherwise, it is returned to the previous hop for similar processing. A new route discovery sequence is launched if the packet is eventually returned to the source. The scheme also uses a notion of reinforcement of currently used routes. A forward ant establishes a pheromone track back to the source, while a backward ant establishes a pheromone track to the destination. ARA prevents loops by memorizing traffic at nodes. If a node receives a duplicate packet, it will send the packet back to the previous node. The previous node deactivates the link to this node, so that the packet cannot be sent in that direction any longer. This loop prevention mechanism is problematic, since further backtracking, if needed, is not resolved, and is based on traffic memorization. Regular data packets are used to maintain the path. In case of link failure, the pheromone trail is set to 0, and the node will send the packet on the second best link. If that link also fails, the node informs the source node about the failure, which then initiates a new route discovery process. Their algorithm is implemented in the ns-2 simulator and compared with AODV. The algorithm, however, is inherently not scalable. The protocol is similar to the AntAODV [MTS1, MTS2] but gives more specific ant behaviour by discussing pheromone use and updates. It also additionally memorizes past traffic and applies pheromone table values instead of ‘fresh’ link indicators.

5.2.5 PROBABILISTIC GUARANTEES FOR ANT-BASED ROUTING IN AD HOC NETWORKS

Eugster [E] considers the probabilistic behaviour of routing (ant-based and gossip-based), multicast, and data replication schemes. His analysis is centered around flooding based methods presented in [GSB] and [MTS2]. He tries to bridge the gap between the different views of reliability-centered distributed systems and communication-centered networking communities. Rather than imposing a rigid deterministic system model on dynamic ad hoc networks in an attempt to obtain "exactly once" reliability guarantees for distributed computations taking place among nodes, the author proposes to embrace the nondeterministic nature of these settings, and work with probabilities and hence notions of partial success. Although the paper builds on existing literature and is more like a survey, it brings out an interesting issue. The paper is based on formal notations used in traditional distributed systems. The properties of the common terms such as unicast, multicast and replication are defined by formal distributed system terms, which is hard for a general audience to understand. He adds many formulas by referring to the original paper, without explaining where and how they come from. The article appears technically sound, but also appears to be mainly of theoretical interest for readers, not offering much for potential designers of ad hoc networks.

5.2.6 ENHANCED ARA PROTOCOL: PRIORITIZED QUEUE, BACKWARD FLOODING AND TAPPING

Gunes, Kahmer, and Bouzazizi [GKB] presented some extensions and improvements to their previous article [GSB]. Probabilistic routing is used instead of selecting the path with the maximal pheromone trail. Pheromone values decrease continually rather than in discrete intervals. Ant packets use a prioritized queue rather than handling them as ordinary data packets. Backward ants use the same type of flooding as forward ants instead of returning on the constructed path. For several packets on the same connection, only one forward ant is created. Finally, similarly as in [RW], MAC-Tap extracts information from packets from the neighbourhood. Experimental data shows improvements, however the need to flood the network is a big disadvantage in mobile ad hoc networks. A flooding technique with less overhead is desirable.

5.2.7 PERA: PROACTIVE, STACK AND AODV BASED ROUTING PROTOCOL

Baras and Mehta [BM] described two ant-based routing schemes for ad hoc networks. One scheme only uses one-to-one or unicast communications where a message sent by one node is only processed at one neighbouring node, while the other utilizes the inherent broadcast one-to-all nature of wireless networks to multicast control and signalling packets (ants), where a message sent by one node is received by all its neighbours. Both algorithms are compared with the well known ad hoc reactive routing scheme, AODV [PR].

The first algorithm in [BM] is similar to the swarm intelligence algorithm described in [DD, SDC]. It uses regular forward, uniform forward and backward ants. Regular forward ants make probabilistic decisions based on pheromone trails, while uniform forward ants use the same probability of selecting each neighbour. Forward ants use the same queue as data packets. When a forward ant is received at a node, and that node is already in the stack of the ant, the forward ant has gone into a loop and is destroyed. Backward ants use the stack which memorized the path to return to the source, using high priority queues. Only backward ants leave pheromones on the trails. Newly created edges are assigned a small amount of pheromone, while broken edges are followed by the redistribution of pheromone to other nodes with normalization.

The second algorithm [BM] is called PERA (Probabilistic Emergent Routing Algorithm). The algorithm applies a route discovery scheme used in AODV to proactively establish routes by the ants. This is a very similar type of route discovery, used reactively in AODV, the difference being that metrics other than hop count may be used. If hop count is used, forward and backward ants travel on high priority queues. If delay is used as metric, they use data queues, so that routes with less congestion are preferred. Multi-path routes are established. Each initial forward ant (only regular forward ants are used) creates multiple forward ants. Only backward ants change the probabilities in the routing tables (pheromone trails are placed using a different reinforcement model than in other articles). Data packets can be routed probabilistically, or deterministically (using the neighbour with the highest probability for the next hop). The simulation was performed on the ns-2 with 20 nodes, and PERA was compared with AODV. The authors observe that end-to-end delay for swarm based routing is low compared to AODV, but the goodput (ratio of data to control packets at each node) is worse (lower) than in AODV. The later conclusion is due to heavy proactive overheads in situations with heavy topological changes. We also note that AODV is used with the hop count as a metric which is unfair when delay is used for comparison (AODV schemes with other metrics are already proposed in the literature).

5.2.8 ANSI: ZONE, FLOODING AND PROACTIVE/REACTIVE ROUTING

Rajagopalan, Jaikeo, and Shen [RJS] described the ANSI (Ad hoc Networking with Swarm Intelligence) protocol. Route discovery and maintenance in ANSI is a combination of proactive and reactive activities. Proactive ants are broadcast periodically to maintain routes in a local area. Whenever other routes are required, a forward reactive ant is broadcast. The outline of the process of ANSI routing is as follows:

- Every node periodically broadcasts *proactive ants* which reach a number of nodes in its local area. Each ant is allocated a certain maximum energy, which is reduced by the energy needed to transmit to a given node. The zone of each node is equal to the transmission radius used in the broadcast. Each receiving neighbour decides to retransmit with a certain fixed probability.
- When a route to a destination *D* is required, but not known at source *S, S* broadcasts a *forward reactive ant* to discover a route to *D.* The number of hops that ant can travel is limited.
- When *D* receives the forward reactive ant from *S,* it source-routes a *backward reactive ant* to the source *S.* The backward reactive ant updates the routing table of all the nodes in the path from *S* to *D.*
- When a route fails at an intermediate node *X,* ANSI buffers the packets which could not be routed and initiates a route discovery to find *D.* Additionally, *X* sends a route error message back to the source node *S.*

The simulation is performed using Qualnet with up to 30 nodes, and comparison is made with AODV. ANSI consistently performed better than AODV with respect to delay characteristics, but the packet delivery rate in ANSI needs to be improved. The scalability of ANSI remains to be investigated. If the zone size remains limited, and hop count for reactive ants becomes unlimited, the performance is expected to be close to that of AODV. If zone size is increased, a comparison with ZRP becomes more appropriate.

Hybrid routing protocols like ZRP [HPS], ADV [BK], and AntAODV [MTS1, MTS2] have leveraged the power of proactive routing with the flexibility and scalability of purely reactive routing. ZRP has a fixed zone radius, while ANSI has a flexible implicit zone radius, which can adapt itself to changing network requirements. This adaptive model resonates with the approach used in ADV [BK], where the amount of proactive activity increases with increasing mobility. Furthermore, the timeout period (equivalent to the beacon timeout in ZRP [HPS]) in ANSI can also be adaptive to reflect the routing needs as the mobility and route errors in a network increase.

5.2.9 ANT AND POSITION BASED ROUTING IN LARGE SCALE AD HOC NETWORKS

5.2.9.1 PROACTIVE, ZONE GROUPING, LOGICAL LINK BASED ROUTING

Heissenbüttel and Braun [HB] described a proactive position and ant based routing algorithm for large and possibly mobile ad hoc networks. The plane is divided into geographical areas (e.g. squares) with all nodes within the same area belonging to the same logical router (LR). All the nodes within a LR share and use the same routing tables. Every logical router has its own set of logical links (LLs). A set of LRs is considered as a communication endpoint for the LLs. For that purpose, a LR groups the other LRs into zones depending on their position relative to it (as shown in Fig. 12). More LRs are grouped together as they are located farther away. It is not a pure hierarchical approach since these zones look different for different LRs. LLs are now established from a specific LR to all its zones. The routing table at each LR has a row for every outgoing LL and a column for every zone. Therefore this is a table with zones as both rows and columns. For a given row zone entry, the table gives probabilities to select column zone entries as the next logical hops, if the destination is located in a row zone. The link costs of incoming LLs are stored in another table. This information will be used to determine the quality of the followed path by the ants.

Ants and data packets are both marked in the header fields with source and destination coordinates. Further, they keep track of the followed path by storing the coordinates of each intermediate relaying node. The followed path can be approximated by a sequence of straight lines. Data packets and ants are routed basically in the same way. The LR determines in which zone the destination coordinates are located and then selects an outgoing LL for that zone with the probability given in the routing table. Multipath routing and load balancing are therefore achieved with this approach. Forward ants are launched periodically from every LR to a random destination. After reaching the destination, the ant becomes a backward ant, and returns to the source node over the recorded path. Pheromone trails are left both ways (whose amount depends on path costs), which evaporate over time. The reason for using a different LL from the zone LL itself when routing is that perhaps there is an obstacle on the direct line, thus greedy routing along exact directions may fail. Ants are supposed to go around such obstacles, and their path is then decomposed into several straight line segments. Each such straight line segment represents a path between two zones, which can be achieved using any existing position based routing scheme (examples are the greedy scheme and greedyface-greedy).

The authors did not present any experimental data on the performance of the proposed scheme, which appears very interesting and appealing. Division into zones requires network pre-processing, and for large networks with n nodes, the number of zones is $O(\log^2 n)$. For n nodes, there are therefore $O(n \log^2 n)$ searches for table entries, and each of them needs a number of ants before the best neighbouring zone is selected. If a constant number of ants is used to test most of the candidate zones, there are $O(n \log^4 n)$ ants generated. In a network where topologies change frequently, the overhead of doing proactive routing may far overweight the benefits of doing so.

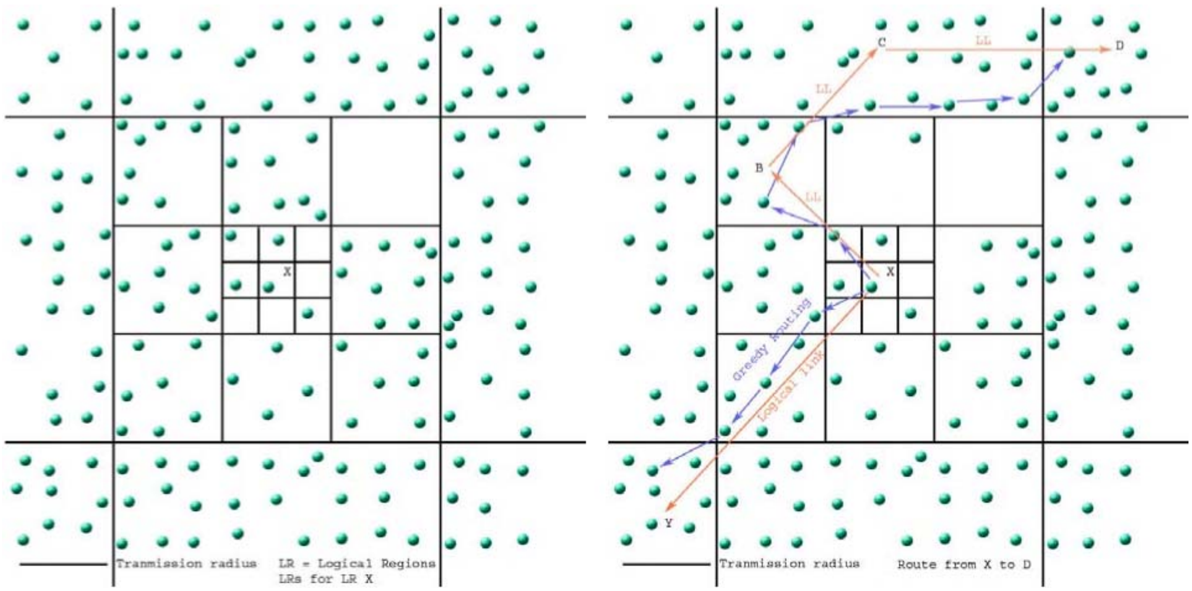

Figure 29: (a) Logical regions as seen from X (b) Logical links for X

Figure 29 demonstrates the main steps in the proactive, zone grouping algorithm presented by Heissenbüttel and Braun [HB]. The transmission radius of the nodes in the network is seen on the bottom left of the figure. Assume that all nodes that are within the transmission radius of each other can communicate directly. These links are not drawn in order to simplify the diagram. The large scale ad hoc network is divided into logical regions, as seen in Figure 29a. The partitioning of the networks is only depicted for logical region *X,* however. The routing table of LR *X* contains the next hops toward all of the logical regions. Only a few of these logical links are drawn in Figure 29b for the purposes of clarity. The actual, physical routing between nodes is done using a greedy algorithm. There exists a direct logical link from LR *X* to LR *Y,* sine the greedy algorithm between them works. On the other hand, three logical links are necessary to reach LR *D.* As seen in figure Figure 29b, each logical link requires its own greedy algorithm. Therefore, the messages may be routed via other logical regions.

5.2.9.2 ANT-BASED LOCATION UPDATES

Camara and Loureiro [CL] proposed the GPSAL protocol which employs ants only to collect and disseminate information about node's locations in an ad hoc network. The destination for an ant could be the node with the oldest information in the routing table. Routing tables contain information about previous and current locations and timestamps of each node, and whether each node is fixed or mobile. When a host receives an ant it compares the routing table present in the ant packet with its routing table and updates the entries that have older information. The protocol, therefore, does not make use of the 'auto-catalytic' effect for finding shortest paths. Furthermore, a shortest path algorithm is applied to determine the best possible route to a destination. Therefore, the protocol assumes that a node knows a lot about the links currently present in the network, and a lot about the positions of other nodes, which certainly will not be true for large scale ad hoc networks. However, once location information is available, localized routing algorithms can be applied, such as greedy [F] or greedy-face-greedy. The algorithm is compared with a position and flooding based algorithm, and decreasing routing overhead is reported. However, the algorithm selected for comparison has significant and unnecessary communication overhead.

5.2.10 MULTICASTING IN AD HOC NETWORKS

Shen and Jaikaeo [SJ] described a swarm intelligence based multicast routing algorithm for ad hoc networks. In the multicasting problem, a source node sends the same message to several destination nodes. The sender and its recipients create a multicast group. There could be several multicasting groups running in the same network. In their algorithm each source starts its session by using shortest paths to each recipient (group member), which is obtained by flooding the message to the whole network, with each group member responding using a reverse broadcast tree (forwarding nodes are decided in this step). Ants are then used to look for paths with a smaller overall cost, that is, to create a multicast core. The cost of multicasting will be reduced if the number of forwarding nodes is reduced. This is achieved by using common paths to several members as much as possible, before splitting into individual or subgroup paths. In addition, each member which is not in the core periodically deploys a small packet that behaves like an ant to opportunistically explore different paths to the core. This exploration mechanism enables the protocol to discover new forwarding nodes that yield lower forwarding costs (the cost represents any suitable metric, such as number of retransmissions, total energy for retransmitting, load balancing, security level etc.). When a better path is discovered, a backward ant (using the memorized path) returns to its origin and leaves a sufficient amount of pheromone to change the route. To avoid cross cutting the 'roads', forwarding nodes keep the highest ID of the nodes that use it to connect to the core, and only the link to a higher ID forwarding node is allowed. Adaptation to ad hoc network dynamics is achieved by cancelling appropriate information whenever a link is broken, and using the best current pheromone trails to continue the multicast. Exploratory ants or periodic core announce messages will restore the connectivity if pheromone trails do not lead toward all group members. The experiments [SJ] are performed on the Qualnet simulator with 50 nodes, and the ant-based protocol is compared with a similar multicasting scheme that does not use ants, and with a simple flooding scheme. The new method performed better, however there exist other multicasting schemes (such as the one that constructs the core based tree first) which are not taken for comparison.

5.2.11 DATA CENTRIC ROUTING IN SENSOR NETWORKS

Singh, Das, Gosavi, and Pujar [SDGP1, SDGP2] proposed an ant colony based algorithm for data centric routing in sensor networks. This problem involves establishing paths from multiple sources in a sensor network to a single destination, where data is aggregated at intermediate stages in the paths for optimal dissemination. The optimal path amounts to a minimum Steiner tree in the sensor network. The minimum Steiner tree problem is a classic NP-complete problem that has numerous applications. It is a problem of extracting a sub-tree from a given graph with certain properties. The algorithm makes use of two kinds of ants, forward ants that travel from the sources to the destination, exploring new paths and gathering information, and backward ants that travel back to the sources from the destination to update the information in each sensor node as they move. A Steiner tree is obtained when the paths traced by forward ants merge into each other or reach the destination. This Steiner tree defines the paths along which data is to be transmitted from the sources to the destination. Because the forward ants move from the sources to the destination, they can also carry packets of data. In the proposed algorithm [SDGP1, SDGP2], each sensor node *i* contains two vectors, the pheromone trails *ph,* and the node potential *pot,* with one entry per each of its neighbours. This node potential is a measure of the proximity of the node to the Steiner tree. The pheromone trails are all initialized to a sufficiently high value to make the algorithm exploratory, and the initial node potentials are based on heuristic estimates. Each sensor node also maintains a variable *tag,* which is initialized to zero, and contains information about how many ants have visited the node.

The total number of forward ants is equal to the number of source sensors, and each ant begins its path from a source sensor. Each such forward ant *m* maintains the tabu list *T* of nodes already visited, as well as a variable *pCost* that indicates the partial cost contributed by the ant's path to the Steiner tree. The list T is initialized to the source sensor where the ant is located, while *pCost* is set to zero. The probability of an ant moving from the current node *i* to its neighbour *j* is proportional to pheromone trail *ph,* and inversely proportional to potential *pot.* In order to prevent the formation of cycles, nodes in *T* that are already visited are excluded. The next location for ant *m* is chosen based on this probability, the new location *j* is pushed into *T,* and *tag* is examined. If *tag* is zero, indicating that location *j* is previously unvisited, the cost of the path *i,* is added to *pCost.* A non-zero value indicates that another ant has already visited the node, and therefore the cost of the path *i* is already incorporated in another ant's *pCost.* Under these circumstances, the forward ant *m* has already merged into an already existing path. It simply follows the previous ant's path to the destination node. The destination node, *d* contains a variable *cost,* the total cost of the Steiner tree path from the sources to *d.* When a forward ant enters the destination node, *d* it increments *cost* by an amount *pCost.* In the present version of the online algorithm, it is assumed that the total number of source nodes is known by the destination at the beginning of the computation. When all forward ants have arrived at the destination, backward ants are generated at the destination. There is a one-to-one correspondence between the forward and the backward ants, and a backward ant, also indexed as *m* acquires the list T of the corresponding forward ant *m.*

Each time a backward ant moves, it pops *T* to obtain the next destination. The backward ants carry a copy of the destination variable *cost.* This information is used to update the pheromones. Updating the tables of node potentials is somewhat more complex. A node's potential is considered low if it is either close to the destination, or brings a forward ant closer to the rest of the Steiner tree. In order to detect the cost of a node to *d,* each backward ant *m* maintains a variable *pCost* similar to a forward moving one, initially zero at the destination d, that gets incremented by an amount equal to *dis(i,j)* whenever a backward ant moves from *j* to *i.* When a backward ant is in any node, *pCost* is the cost of the path joining the node to *d.* In order to compute the cost of joining a node to another route, i.e. only a branch of the Steiner tree, another variable *rCost* is used by backward ants that are updated in the same manner. However, *rCost* is reset to zero each time a backward ant detects a split in a path leading to more than one branch of the Steiner tree. A split, leading to another branch is detected by examining the *tag* variable of a node *i.* If the previous node of the backward ant was *j,* then node *i* is a separate branch if *tag(i)< tag(j).* A backward ant *m* leaving node *j* decrements the *tag(j)* variable. Backward ants travel back to the sources in *S* and reset these tag variables to zero for future ants. The updating rule for the potential is a linear combination of *rCost* and *pCost.* This updating is carried out only if the node potential gets lowered.

The experimental data showed that the ant based algorithm performed significantly better than the address-centric one, where shortest paths are used from each source sensor to the destination.

5.2.12 SUMMARY

The dynamic and wireless nature of ad hoc networks has led to some modifications and new ideas in ant based routing schemes. The frequent edge creation and breakage has added the portion of exploratory ants that behave at random or with uniform probability, so that new paths are quickly discovered and reinforced, or new edges incorporated quickly into the path. Most articles exploit the one-to-all nature of message transmission, which gave the opportunity to multiply an ant and flood it throughout the network instead of simply following a path as in other considered communication networks. It also allowed nodes to overhear transmissions from neighbouring nodes and use them to update their pheromone tables.

While new opportunities for ad hoc networks are exploited in the proposed solutions, their experimental evaluation apparently was not done properly. Most authors only compare their methods with other weaker ant based methods, or with the standard version of the AODV protocol, without considering existing AODV improvements that might prove competitive. Also, position based schemes and routing schemes were not compared with the best existing position based methods. Therefore future articles are expected to provide a realistic evaluation of ant based routing in ad hoc networks, with emphasis on the primary question, whether the communication overhead imposed by using ants is worthwhile for obtaining gains in paths, especially in dynamic scenarios.

The need for improved accuracy of simulations exists also in ant based routing for communication networks. The routing problem becomes more challenging if constraints are added, for example to achieve quality of service. Flow control and admission control in routing are also important to incorporate. The existing reported simulation results that are encouraging are not done in real networks with real equipment. The primary concerns for routing are about convergence to a steady state, adaptation to changing environments, and oscillation [W, WP]. One of the interesting challenges for ant based routing is in their applications for routing and searching in Internet networks.

Further ant based methods can be expected soon; especially for position based routing. The recovery scheme proposed by Finn [F] is based on flooding up to *n* hops, hoping that a node closer to the destination than the current node will be found. This introduces a lot of flooding but still does not guarantee delivery. We believe that is worthwhile to consider the application of ants in search for such a node. A certain number of ants can be sent, each with a certain limited distance from the current node. The distances traveled by the ants could be set incrementally so that, if a closer node is not found by a certain time, new ants with a longer search range are sent. This is a preliminary idea, and obviously extensive simulation and modification is needed to get an acceptable version.

References

- References on routing in ad hoc networks

[BM] J. S. Baras and H. Mehta, A Probabilistic Emergent Routing Algorithm for Mobile Ad Hoc Networks, Proc. WiOpt, Sophia-Antipolis, France, March 2003.

[CL] D. Camara, A. Loureiro, GPS/Ant-like routing algorithm in ad hoc networks, Telecommunication Systems, 18, 1-3, 2001, 85-100.

[E] P. Eugster, Probabilistic guarantees and algorithms for ad hoc networks, manuscript, 2003.

[GKB] M. Günes, M. Kähmer, and I. Bouazizi, Ant-routing-algorithm (ARA) for mobile multi-hop ad-hoc networks - new features and results, Proceedings of the 2nd Mediterranean Workshop on Ad-Hoc Net- works (Med-Hoc-Net'2003), Mahdia, Tunisia, 25-27, June 2003.

[GSB] M. Gunes, U. Sorges, and I. Bouazizi, ARA - the ant colony based routing algorithm for manets, Proc. ICPP Workshop on Ad Hoc Networks IWAHN, 2002, 79-85.

[HB] M. Heissenbüttel and T. Braun, Ants-based routing in large scale mobile ad-hoc Networks, Kommunikation in Verteilten Systemen KiVS03, Leipzig, Germany, February 25-28, 2003.

[ITT] M.T. Islam, P. Thulasiraman, R.K. Thulasiram, A parallel ant colony optimization algorithm for all-pair routing in MANETs, Proc. IEEE Int. Parallel and Distributed Processing Symposium IPDPS, Nice, France, April 2003.

[MM] H. Matsuo, K. Mori, Accelerated ants routing in dynamic networks, 2nd Int. Conf. Software Engineering, Artificial Intelligence, Networking & Parallel Distributed Computing, Nagoya, Japan, 2001.

[MTS1] S. Marwaha, C. K. Tham, D. Srinivasan, A novel routing protocol using mobile agents and reactive route discovery for ad hoc wireless networks, Proc. IEEE ICON, 2002.

[MTS2] S. Marwaha, C. K. Tham, D. Srinivasan, Mobile agent based routing protocol for mobile ad hoc networks, Proc. IEEE GLOBECOM, 2002.

[RJS] S. Rajagopalan, C. Jaikaeo, and C.C. Shen, Unicast Routing for Mobile Ad hoc Networks with Swarm Intelligence, Technical Report #2003-07, University of Delaware, May 2003.

[RW] M. Roth and S. Wicker, Termite: Emergent ad-hoc networking, Proceedings of the 2nd Mediterranean Workshop on Ad-Hoc Net- works (Med-Hoc-Net'2003), Mahdia, Tunisia, 25-27, June 2003.

[SDGP1] G. Singh, S. Das, S. Gosavi, S. Pujar, "Ant Colony Algorithms for Steiner Trees: An Application to Routing in Sensor Networks", *Recent Developments in Biologically Inspired Computing,* Eds. L. N. de Castro, F. J. von Zuben, 2003, under preparation.

[SDGP2] S. Das, G. Singh, S. Gosavi, S. Pujar, "Ant Colony Algorithms for Data-Centric Routing in Sensor Networks", *Proceedings, Joint Conference on Information Sciences,* Durham, North Carolina, 2003.

[SJ] C.C. Shen, and C. Jaikaeo, Ad hoc Multicast Routing Algorithm with Swarm Intelligence, Technical Report #2003-08, DEGAS Group, Dept. of CIS, University of Delaware, 2003.

- ***References on swarm intelligence for routing in communication networks***

[BHGSKT] E. Bonabeau, F. Henaux, S. Guerin, D. Snyers, P. Kuntz, G. Theraulaz, Routing in telecommunication networks with 'smart' ant-like agents, Proc. 2nd Int. Workshop on Intelligent Agents for Telecommunication Applications, Paris, France, July 1998.

[DD] G. DiCaro, M. Dorigo, AntNet: Distributed stigmergetic control for communication networks, J. of Artificial intelligence Research, 9, 1998, 317-365.

[G] S. Guerin, Optimisation multi-agents en environment dynamique: Application au routage dans les reseaux de telecommunications, DEA Dissertation, University of Rennes I, France, 1997.

[SDC] D. Subramaniam, P. Druschel, J. Chen, Ants and reinforcement learning: A case study in routing in dynamic networks, Proc. IEEE MILCOM, Atlantic City, NJ, 1997.

[SHB] R. Schoonderwoerd, O.E. Holland, J.L. Bruten, Ant-like agents for load balancing in telecommunication networks, Proc. First ACM Int. Conf. on Autonomous Agents, Marina del Ray, CA, USA, 1997, 209-216.

[W] T. White. Swarm intelligence and problem solving in telecommunications, *Canadian Artificial Intelligence Magazine,* Spring 1997.

[WP] T. White and B. Pagurek, Towards multi-swarm problem solving in networks, Proc. Third Int. Conf. Multi-Agent Systems ICMAS, July 1998, 333-340.

- ***References on routing in ad hoc networks***

[BK] R.V. Boppana, S.P. Konduru, An adaptive distance vector routing algorithm for mobile, ad hoc networks, IEEE INFOCOM, Anchorage, Alaska, 2001.

[F] G.G. Finn, Routing and Addressing Problems in Large Metropolitan-Scale Internetworks, Research Report ISU/RR-87-180, Inst. for Scientific Information, Mar. 1987.

[HPS] Z. J. Haas, M. R. Pearlman, and P. Samar. The zone routing protocol (ZRP) for ad hoc networks, July 2002. IETF Internet Draft, draft-ietf-manetzone-zrp-04.txt.

[PR] C. Perkins, E. M. Royer, Ad hoc on demand distance vector routing, Proc. IEEE Workshop on Mobile Computing Systems and Applications, WSMA, Feb. 1999.

5.3 DISTRIBUTED MANUFACTURING OR MAINTENANCE

Agent-based approaches to manufacturing scheduling and control are attractive because they offer increased robustness against the unpredictability of factory operations.

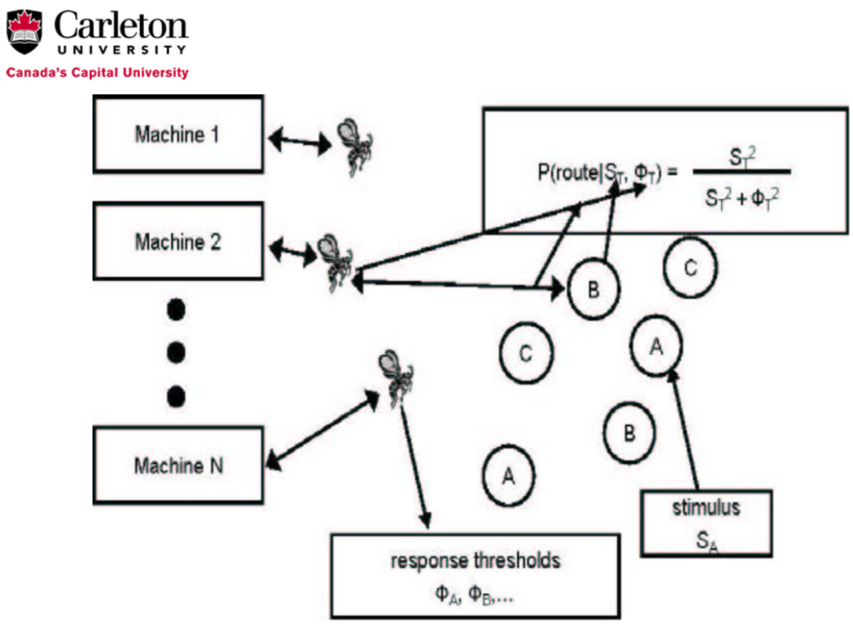

Figure 30: Wasps for Distributed Manufacturing

If the figure above, wasps use response threshold mechanisms to choose which tasks should be routed to which machines using the principles briefly described in section 4.9.2. Cicirello extends the basic algorithm to resolve conflicts that arise when two wasps respond to the same stimulus (job). In this scenario a calculation of "dominance occurs" -- similar to the self-organized hierarchies of real wasps -- as shown in the figure on the next page.

Using these algorithms has been shown to significantly improve scheduling. The implications of military maintenance scheduling are clear.

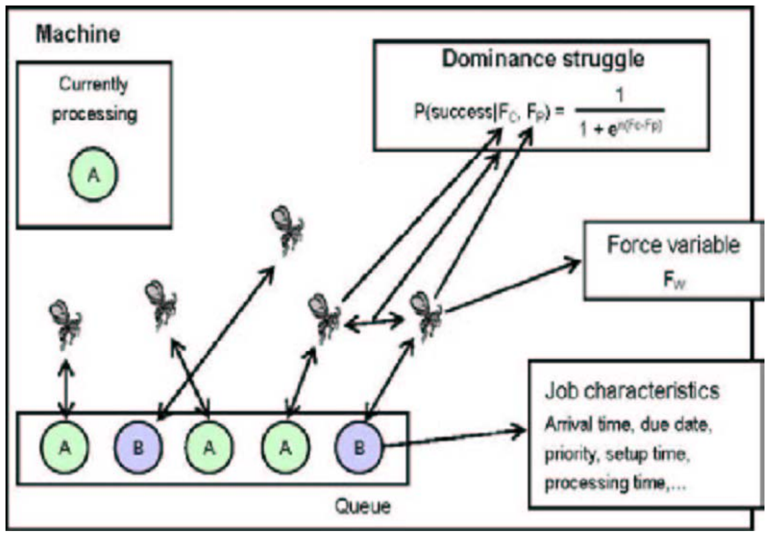

Figure 31: Resolving Conflicts

5.4 DYNAMIC ORGANIZATIONAL STRUCTURE

The algorithms of section 4.9.2 can be used for role assignment without central control. A role here is the ability to perform a set of task with a specified level of expertise. This is a problem of considerable importance to the military which typically relies upon centralized, top down decision making. In the military, soldiers and equipment can perform a variety of roles; the question is given a group of soldiers and equipment what roles should be assigned to them in order to meet the threat at hand?

Consider the game of soccer. Simplifying the game somewhat, teams have four basic roles: goalkeeper, defender, midfield and forward. Two sides compete to score the most goals. Threats are determined based upon the position of the ball and the number of opponent players within a given distance from the goal. A soccer side also has a strategy, which cannot often be inferred until the match is underway. Players get injured, or become less effective as they tire.

The previous description corresponds quite closely to what goes on in a battlefield scenario. Threats are perceived based upon intelligence – presumably gathered by sensor networks and unmanned surveillance drones in the future. Strategy only becomes apparent as the battlefield engagement unfolds. Soldiers die and equipment fails. Soldiers and equipment deployed on the battlefield have associated roles and the goal of distributed command and control is to optimize the response to any perceived threat.

Using division of labour and task allocation based upon response thresholds allows a stigmergic system to overcome real time failures and facilitates the emergence of specialization over time. Furthermore, if a specialized agent should fail, other agents will eventually take over the role as a threat will continue to escalate until it can no longer be ignored.

5.5 COLLECTIVE ROBOTICS

5.5.1 INTRODUCTION

Collective, or swarm-based, robotics is a relatively new field. One of the earliest researchers in the field was Kube (see section 8.1, number 10) who demonstrated that simple robots with no inter-robot communication could collectively push heavy objects and cluster objects in a manner similar to ants. His robots were homogeneous.

Martinoli (see section 8.1, number 9) is an active researcher in the field. Martinoli has undertaken considerable work in the areas of distributed exploration and collaboration. His PhD [180] provides a very good introduction to the problems of creating swarms of robots that exhibit complex distributed collective problem solving strategies. More recently, March 2005, the Swarm Bots project lead by Marco Dorigo (see section 8.1, number 3) completed its 3.5 year investigation into the creation of teams of small robots using stigmergy.

5.5.2 AUTONOMOUS NANOTECHNOLOGY SWARMS

NASA's autonomous nanotechnology swarms (ANTS) creates communities of intelligent teams of agents where redundancy is built in. The ANTS architecture uses a biologically inspired approach, with ants as primary inspiration. It is the most sophisticated of all of the stigmergic systems currently in design. Swarms of up to 1000 nodes will be deployed on deep space missions to study asteroids, with sub-swarms of 100 nodes being independently tasked with given mission parameters. Several classes of swarm unit have been defined with measurement (imaging, for example), communication and leadership characteristics. A generic worker class has also been designed. The ANTS project timeline extends beyond 2030 when the first missions are envisaged. However, several important engineering concepts have already been developed (See http://ants.gsfc.nasa.gov/netactivities.html for details)

In the ANTS system, the basic physical structure is a tetrahedron that flexes, changing shape causing a tumbling motion thereby allowing movement over a surface. Tetrahedral structures are used at all levels of the ANTS design, the designers arguing that this structure is one of the most stable naturally-occurring structures. The ANTS system consists of small, spatially distributed units of autonomous, redundant components. These components exhibit high plasticity and are organized as hierarchical (multilevel, dense heterarchy) and inspired by the success of social insect colonies. The ANTS system uses hybrid reasoning – symbolic and neural network systems – for achieving high levels of autonomous decision making.

5.5.3 SWARM BOTS

The main scientific objective of the recently completed *Swarm-bots* project (see http://www.swarm-bots.org) was to study a novel approach to the design and implementation of self-organising and self-assembling artifacts. This novel approach used as theoretical roots recent studies in swarm intelligence, that is, studies of the self-organizing and self-assembling capabilities shown by social insects and other animal societies employing stigmergic principles extensively.

The main tangible objective of the project was the demonstration of the approach by means of the construction of at least one of such artifact. A *swarm-bot* was constructed. That is, an artifact composed of a number of simpler, insect-like, robots (*s-bots*), built out of relatively cheap components, capable of self-assembling and self-organizing to adapt to its environment.

Three distinct components were developed: *s-bots* (hardware), simulation (software), and swarm-intelligence-based control mechanisms (software). A set of hardware *s-bots* that can self assemble into a shape-changing *swarm-bot* were developed that were capable of accomplishing a small number of tasks. Tasks completed were dynamic shape formation and shape changing and navigation on rough terrain. In both cases, teaming is crucial as a single *s-bot* cannot accomplish the task and the cooperative effort performed by the *s-bots* aggregated in a *swarm-bot* is necessary.

5.5.3.1 PROJECT RESULTS

An s-bot was developed, an example of which is shown in the figure below:

A Bot

Figure 32: Example of an s-bot

As can be seen in the figure, the s-bot has both an extendible gripper capable of attaching to another s-bot and a fixed length gripper. These grippers allow for extended s-bot structures (rigid and flexible) to be created.

The project, demonstrated the feasibility of the integration of swarm intelligence, reinforcement learning and evolutionary computation paradigms for the implementation of self-assembling and self-organizing metamorphic robots by constructing a *swarm-bot* prototype. The working prototype achieved the following three sets of objectives:

Dynamic shape formation/change:

A *swarm-bot*, composed of at least twenty *s-bots* randomly distributed on the floor, self-assembled into a number of different planar and 3D geometric configurations, for example like those found in ant colonies and in patterns of differential adhesion by developing cells. These configurations were closed shapes with internal structure, such as:

1. centre/periphery figures (for example, all *s-bots* with a given set of sensors will stay on the outer perimeter whereas all other *s-bots* will remain inside);

2. checker-board;

3. split (each half of the assembly will contain *s-bots* with similar characteristics);

Transitions between shapes were also tested. A long-term goal, but not necessary for the success of this project, was to achieve *emergent* expulsion of "dead bodies", that is *s-bots* that malfunction.

Navigation on rough terrain:

A *swarm-bot*, composed of at least twenty *s-bots*, was capable of autonomously moving across the terrain guided by sensory information gathered by individual *s-bots*. The following objectives were achieved:

1. light following while maintaining the original shape (for example, one of those described above);

2. light following through narrow passages and tunnels that require dynamic reconfiguration of the *swarm-bot*;

3. passing over a hole or through a steep concave region that could not be passed by a single s-bot;

4. moving from point A to B (for example, on a shortest possible trajectory) on rough terrain.

A major scenario for evaluation of the project was based on a search and rescue concept where a swarm of robots must locate and retrieve a heavy object and take it to a goal location. The scenario is graphically represented in Figure 33. The s-bots have to deal with unknown, rough terrain containing obstacles and holes. It should be noted that teaming is required in order to cross the holes in the terrain. Control algorithms for the s-bots were generated by inducing the descriptions of neural networks using genetic algorithms. The simulator was used extensively in order to design the continuous time recurrent neural network controllers.

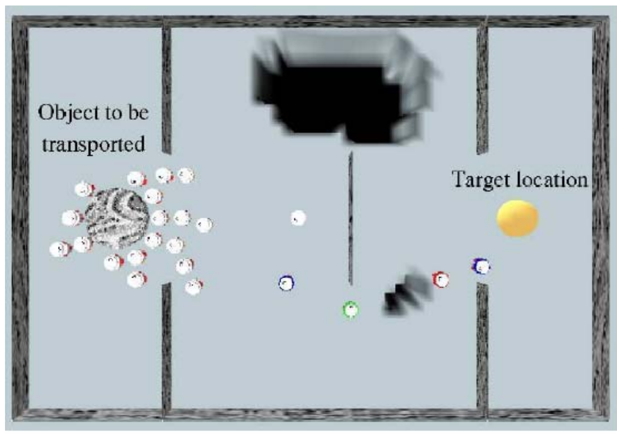

Figure 33: Search and Recover Scenario

A summary of the behaviours demonstrated includes:

- hole/obstacle avoidance
- finding an object or a goal
- adaptive division of labour
- pattern formation:
    - co-ordinated motion
    - aggregation, self-assembling, grasping
    - passing over a hole
    - moving on rough terrain
    - cooperative transport of an object

In experiments with small object retrieval, rewarding robots based on success and failure automatically categorized them into three categories: forager, undecided and loafers. Modifications to the abilities of some robots were reflected in their specializations; e.g. more speed increased the likelihood of becoming a successful forager.

Strategies for cooperation that were evolved could be quite simple and direct; e.g. two robots both pushing on the same side of something, or a little more complex; e.g. chain-pulling formations. The experimental results showed that the neural nets evolved for object transport also extended well to larger groups for larger objects and different shapes and sizes did not seem to drastically change the effectiveness of the evolved smart bots.

Figure 34: Crossing a Trench

As the figure above shows, rigid structures could be created thereby enabling trench crossing [181]. The publications page of the Swarm-bots project (http://www.swarm-bots.org/index.php?main=2) provides a wide range of reports and papers on the results of the project. Similar aggregates of s-bots were observed "holding hands" for the traversal of uneven ground where several of the s-bots would be seen with their wheels off the ground at various points on the terrain.

Specifically, scaling properties of swarm-bots were found to reasonable (30-35 s-bots) and that desired shapes could be reliably reproduced. Secondly, the time taken to move objects to a goal location was reasonable and robust with respect to the removal of individual s-bots.

5.5.3.2 Summary

This project represents a significant advance in the understanding of swarm robotic systems. While the s-bots are still simple in design – far from the complexity required of battlefield hardware – their collaborative behaviour is impressive. Furthermore, the automated design of continuous time domain recurrent neural network controllers using genetic algorithms demonstrates that emergent behaviour can be engineered. This European Union project (EU IST-2000-31010), was considered highly successful (http://www.cordis.lu/ist/fet/press.htm#success). This author would strongly recommend monitoring follow-on projects as they clearly have significant value from a military perspective.

5.6 MECHATRONICS

Mechatronics is the discipline of building reconfigurable robots. An excellent resource on the subject can be found at Colorado State (Section 8.5, reference number 12). Robots are made out of modules, which could crudely be described as intelligent Lego bricks. Plugging the bricks (or modules) together in particular ways allows a mechatronic robot to more or less effectively solve a problem such as moving over terrain of a given class; e.g. swamp or very rocky. In the mechatronic domain stigmergy is represented as perception of self. While the Swarm-bot project can be thought of as fitting into this category, mechatronic research focuses on the assembly, re-assembly and reconfiguration of simpler units. Continuing with the comparison with the Swarm Bot project, mechatronic research is concerned with the construction of an s-bot rather than the swarm-bot.

Stigmergy in this area is typically sematectonic – the robot/module configurations being used to drive the configuration process.

Noteworthy work here includes the self-reproducing machine work of Lipsen of Cornell (see http://www.mae.cornell.edu/ccsl/research/selfrep/). Here, mechatronic modules within a robot know how to reproduce in same way as a human cell has a blueprint of the individual. Mechatronic robots constructed using Lipsen's modules know how to incorporate new "blank" modules for reproduction. Stigmergy in this system is represented by the interactions between modules; i.e. a module knows the connections that it has to other modules and can tell a blank module how to configure itself.

This work clearly has military significance in that blank modules could potentially be dropped onto the battlefield, located by existing robots and then used to duplicate (or repair) the robots present.

The Polybot project from Xerox Parc (Section 8.5, reference number 14) has developed a number of sophisticated prototypes that use local interactions between multiple, identical modules in order to solve tasks. This project does not appear to have been active since 2001. The importance of this work is that the modules, when connected, allow for locomotion and can be reconfigured for movement over a wide variety of terrains. Coupled with Lipsen's work, the potential for repairable, reproducible battlefield robots capable of autonomous activity seems plausible.

5.7 AMORPHOUS COMPUTING

The Amorphous Computing project at MIT (Section 8.5, reference number 15) is included here as it represents an analog approach to swarm system design. This project is referenced by Lipsen in his work.

In amorphous computing systems, a colony of cells cooperates to form a multi-cellular organism under the direction of a program (loosely called a genetic program) that is shared by all members of the colony.

The objective of amorphous computing is the creation of algorithms and techniques for the understanding of programming materials. Essentially, amorphous computing seems to incorporate the biological mechanisms of individual cells into systems that exhibit the expressive power of digital logic circuits. Stigmergy in such systems can be either marker-based or sematectonic and be either scalar or vector in extent. An amorphous computing medium is a system of irregularly placed, asynchronous, locally interacting computing elements. The medium is modelled as a collection of “computational particles” sprinkled irregularly on a surface or mixed throughout a volume. In essence, the computational assembly forms an ad hoc network. Research into self-healing structures, circuit formation, programmable self-assembly and self-organizing communication networks are a small sample of the work undertaken.

In principle, if successful, amorphous computing would allow smart materials to be programmable. An example of a programmable material would be one that would sense the surroundings and adaptively camouflage the wearer.

5.8 MILITARY APPLICATIONS

Swarming is not new to the military; however, understanding the importance of large scale exploitation of stigmergy is. Altarum’s work described in the next section is the beginnings of the use of sensor fusion applied to several pheromones. The work describes the use of marker-based stigmergy for target acquisition and tracking. Other work by Altarum’s Dr. Parunak [234] provides insight into how marker-based stigmergy can be more generally applied to military problems. Dr. Parunak’s group should be considered to be the ***leading authority*** on the application of stigmergy to military problems.

The intelligent minefield is an example of sematectonic stigmergy – the mines themselves being the stimulus (or lack of stimulus) to cause mine reconfiguration after minefield breach.

The Autonomous Negotiating Team section is included simply to document the fact that bottom up reasoning, using simple agents that negotiate locally, is now being researched for the purpose of making organizations capable of more rapid decision making; something that has been a problem for the military historically.

5.8.1 TARGET ACQUISITION AND TRACKING

This section adapted from Parunak’s presentation at the Conference on Swarming and C4ISR, Tyson’s Corner, VA, 3rd June, 2003.

Altarum’s research has concentrated on applications of co-fields modeled rather closely on the pheromone fields that many social insects use to coordinate their behaviour. They have developed a formal model of the essentials of these fields, and applied them to a variety of problems. Altarum’s view is to allow the integration of multiple pheromones, using the fused sensor readings to drive the movement of agents in the space being monitored and controlled.

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

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

These dynamics can be modeled in a system of difference equations across a network of “places” at which agents can reside and in which they deposit and sense increments to scalar variables that serve as “digital pheromones,” and these equations are provably stable and convergent [195]. They form the basis for a “pheromone infrastructure” that can support swarming for various C4ISR functions, including path planning and coordination for unpiloted vehicles, and pattern recognition in a distributed sensor network.

Path Planning.—Ants construct networks of paths that connect their nests with available food sources as described in Section 4.9.1.1. Mathematically, these networks form minimum spanning trees, minimizing the energy ants expend in bringing food into the nest. Graph theory offers algorithms for computing minimum spanning trees, but ants do not use conventional algorithms. Instead, this globally optimal structure emerges as individual ants wander, preferentially following food pheromones and dropping nest pheromones if they are not holding food, and following nest pheromones while dropping food pheromones if they are holding food.

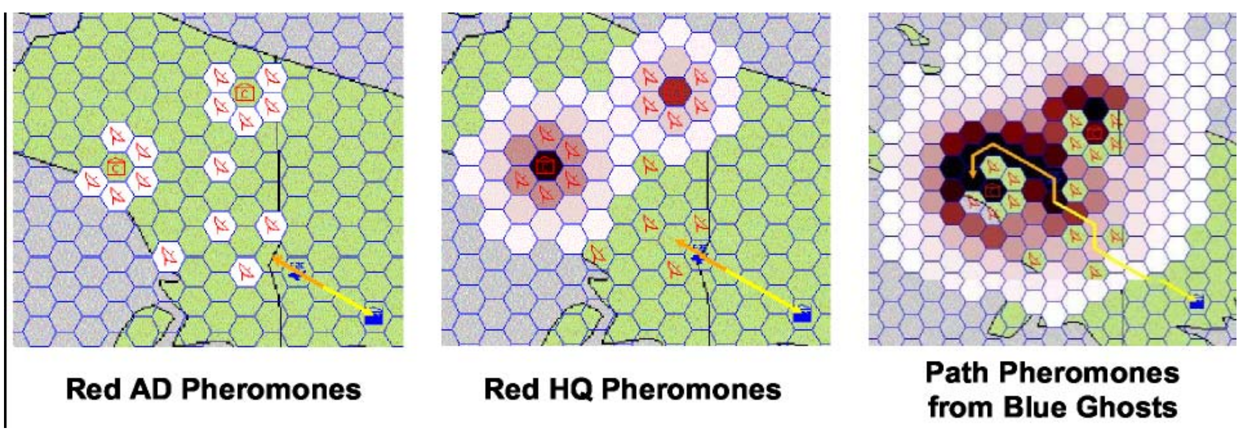

Figure 35: Digital Pheromones for Path Planning

Altarum have adapted this algorithm to integrate ISR into a co-field that then guides unpiloted vehicles away from threats and toward targets [224]. The battlespace is divided into small adjoining regions, or “places,” each managed by a “place agent” that maintains the digital pheromones associated with that place and serves as a point of coordination for vehicles in that region. The network of place agents can execute on a sensor network distributed physically in the battlespace, onboard individual vehicles, or on a single computer at a mission command center. When a Red entity is detected, a model of it in the form of a software agent is initiated in the place occupied by the Red entity, and this agent deposits pheromones of an appropriate flavour indicating the presence of the entity. The agent can also model any expected behaviours of the Red entity, such as movement to other regions. Blue agents respond to these pheromones, avoiding those that represent threats and approaching those that represent targets, and depositing their own pheromones to coordinate among themselves. (The distinction between threat and target may depend on the Blue entity in question: a SEAD resource would be attracted to SAM's that might repel other resources.) The emergence of paths depends on the interaction of a large number of Blue entities. If the population of physical resources is limited, a large population of software only "ghost agents" swarms through the pheromone landscape to build up paths that the physical Blue agents then follow. Figure 36 shows repulsive and attractive Red pheromones, and the resulting co-field laid down by Blue ghost agents that forms a path for a strike package to follow. This mechanism can discriminate targets based on proximity or priority, and can plan sophisticated approaches to highly-protected targets, approaches that centralized optimizers are unable to derive.

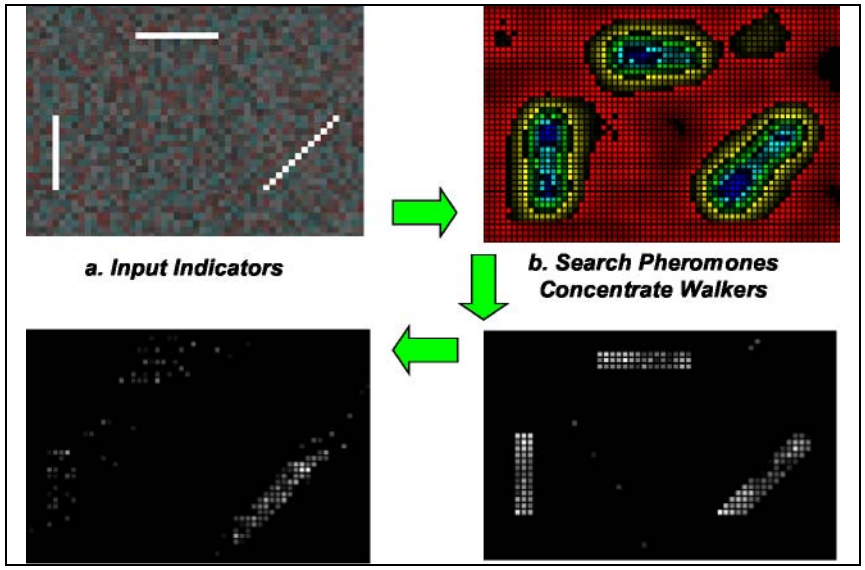

Figure 36: Multiple species of software agents swarming over a sensor network can enable the network to detect patterns without centralizing the data.

d. Pattern Pheromones Select Patterns

c. Find Pheromones Localize Behaviors

Vehicle Coordination.—The algorithms developed in our path planning work were incorporated into a limited-objective experiment conducted by SMDC for J9 in 2001 [203], [229]. In this application, up to 100 UAV's coordinated their activities through digital pheromones. UAV's that had not detected a target deposited a pheromone that repelled other UAV's, thus ensuring distribution of the swarm over the battlespace. When a UAV detected a target, it deposited an attractive pheromone, drawing in nearby vehicles to join it in the attack. This capability enabled the deployment of many more vehicles without an increase in human oversight, and yielded significant improvements in performance over the baseline, including a 3x improvement in Red systems detected, a 9x improvement in the system exchange ratio, and an 11x improvement in the percentage of Red systems killed.

Pattern Recognition.—The Army's vision for the Future Combat System includes extensive use of networks of sensors deployed in the battlespace. Conventional exploitation of such a network pipes the data to a central location for processing, an architecture that imposes a high communication load, delays response, and offers adversaries a single point of vulnerability. Altarum has demonstrated an alternative approach in which pattern recognition is distributed throughout the sensor network, enabling individual sensors to recognize when they are part of a larger pattern [198]. The swarming agents are not physical, but purely computational, and move between neighbouring sensors using only local communications. Figure 36a shows an example distribution of sensors (a 70x70 grid). With a global view, we can quickly identify the sensors with high readings (plotted as white), but individual sensors do not have this perspective and cannot be sure whether they are high or low.

One species of swarming agents compares each sensor's readings with a summary of what it has seen on other sensors to estimate whether the current sensor is exceptional, and deposits search pheromones (Figure 36b) to attract its colleagues to confirm its assessment. Each agent has seen a different subset of the other sensors, so a high accumulation of find pheromone on a sensor (Figure 36c) indicates that the sensor really is high in comparison with the rest of the network, and it can call for appropriate intervention. A second species of agents moves over the sensors both spatially and (through stored histories of recent measurements) chronologically. The movement of this species is not random, but embodies a spatio-temporal pattern, and its pheromone deposits highlight sensors that are related through this pattern (in Figure 36d, an orientation from SW to NE).

5.8.2 INTELLIGENT MINEFIELDS

The intelligent minefield project is an example of a self-repairing sensor network. An intelligent mine field is self-deploying and self repairing, as shown in the figure below.

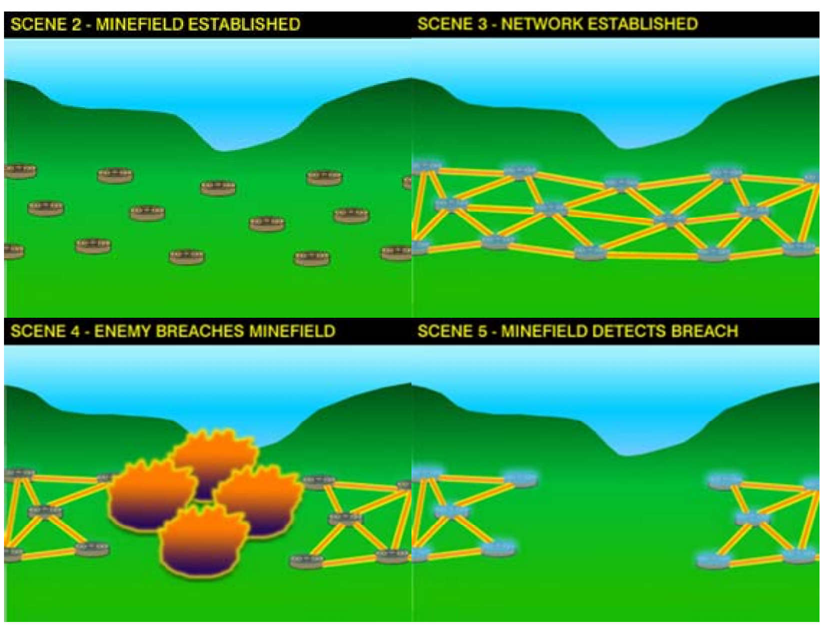

Figure 37: Intelligent Minefield

In the intelligent minefield scenario shown above, mines are deployed and determine their nearest neighbours, subsequently creating links to them. Stigmergy here is represented by knowledge of local connections. In scene 4 a minefield breach is created by enemy activity.

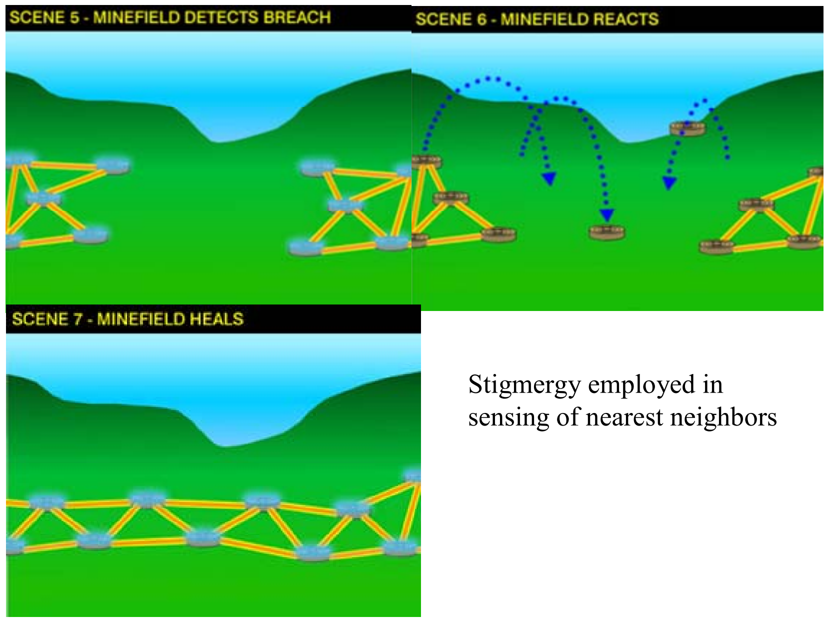

Stigmergy employed in sensing of nearest neighbors

Figure 38: Self-repairing Minefield

In scene 5 in the above figure, the mines detect the breach. Messages are then routed throughout the sensor network in order to determine which mines to redeploy, shown in scene 6. Redeployment then occurs to recreate the minefield; links are regenerated and the minefield is fully connected once more (scene 7).

5.8.3 AUTONOMOUS NEGOTIATING TEAMS

The Autonomous Negotiating Teams (ANTS) project has the goal of autonomously negotiating the assignment and customization of resources – such as weapons (or goods and services) to their consumers, such as moving targets. The goal is timely and near-optimal decision making.

ANTS using real-time negotiation using dynamically constructed organizations. An ANTS system works in a bottom-up fashion; each entity has an ant associated with it. Examples of entities are brigades, soldiers, rifles, radios etc. Ants discover each other; negotiate resources, authorizations, capabilities, actions and plans using only local interactions.

This project does not make clear how stigmergy is employed. It is included here as an example of the interest in bottom-up decision making. This project should be monitored for progress as it may be possible to ascertain its use of stigmergy at some future date.

6 FUTURE RESEARCH AND TECHNOLOGY ASSESSMENT

6.1 INTRODUCTION

This section deals with potential future research that might facilitate the introduction of stigmergic principles into military and security systems. Before proceeding with a discussion of research and a scenario that supports it, an assessment of the research and technology presented in this report is provided.

6.2 ASSESSMENT

This report has covered a wide range of material that spans tutorial material: "what is stigmergy", to trials that have occurred: Altarum's Digital Pheromone target tracking system. A wide range of technology levels have been observed, from TRL 1 through to TRL 8. Altarum's Digital Pheromone target tracking system should be considered the most mature technology at TRL 7/8 having flown in operational experimental conditions. The Swarm-bots project – arguably the most exciting project from a robotics perspective – is assessed at TRL 4/5. Mechatronic research is generally at TRL 4. The algorithms derived from the Ant Colony Optimization metaheuristic ("Smart Algorithms") should be rated at TRL 2/3 (only because physical systems are not generated in this environment). The MANET routing algorithm research should be rated at TRL 3. Routing algorithms for sensor networks would also attract a TRL rating of 3. Sensor technology achieves the rating of TRL 5/6. The tables below provide an extended assessment over the timeframe 2005-2030.

Table 4: Technology Readiness Assessment

**** Technology Readiness Levels***

Code	***Colour***	***Definition***
1		Basic principles observed and reported
2		Technology concept and/or application formulated
3		Analytical and experimental critical function and/or characteristic
4		Component and/or breadboard validation in laboratory environment
5		Component and/or breadboard validation in relevant environment
6		System/subsystem model or prototype demonstration in a relevant environment
7		System prototype demonstration in a operational environment
8		Actual system completed and 'flight qualified' through test and demonstration
9		Actual system 'flight proven' through successful mission operations

TECHNOLOGY FORECAST
Stigmergy

a. Sensor Networks Hardware

b. Sensor Networks Algorithms

c. Digital Pherome Applications

d. Smart Robots

e. Smart Materials

f. Smart Applications

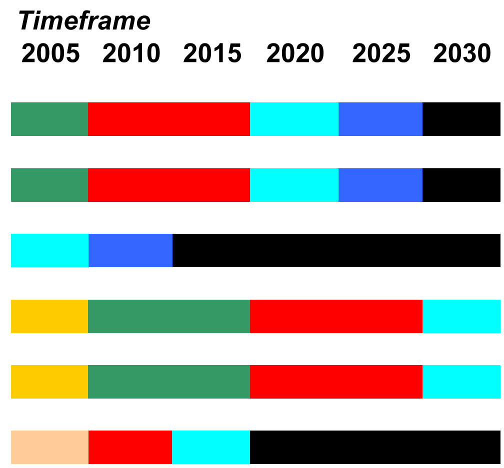

Technology Readiness
(Colour codes defined above) *

Definitions:

a. Sensor Network Hardware

Sensors capable of deployment in harsh military environments are the goal
The sensors should be deployable by ground troops, or by UAVs

b. Sensor Network Algorithms

Algorithms refer to routing protocols. Stability has to be proven.
Algorithms for information dispersal.

c. Digital Pheromone Applications

Tracking and target surveillance applications with superior performance to traditional solutions.

d. Smart Robots

Robots that are capable of threat assessment and self-deployment.
Robots should support reconfigurability and replication.

e. Smart Materials

Materials with large numbers of embedded processors capable of autonomous computation.

f. Smart Applications

Optimization applications for military logistical problems

6.3 DISCUSSION

The effectiveness of a stigmergic system depends upon the sophistication of the environment. Sophistication is measured across 2 dimensions: the ease with which information can be stored and retrieved, and the dynamics associated with the environment. While insects deal with the environment presented to them; using pheromones and naturally occurring structures as signals and principles of physics and chemistry as controlling dynamics, humans are not limited in the same way. Sensor networks represent *enabling technology* for stigmergic systems and can be constructed to contain arbitrarily complex signals and dynamics. Clearly, then, the future of sensor networks as enabling technology for stigmergic applications is key. This author *strongly* recommends that sensor network research be monitored for key developments. Pister's smart dust has already shown that millimetre-scale sensors are possible and even nanoscale sensors are being discussed. As such, the smart dust sensors are TRL 5/6. Within a decade, possibly slightly more, cost-effective construction of sensor networks in battlefield scenarios will be viable, where sensors are scattered on the ground by soldiers moving over the terrain or from unmanned autonomous vehicles. The sensor networks will then communicate with unmanned autonomous vehicles flying overhead for purposes of enemy target tracking and surveillance. The deployed sensor networks will be able to determine when breaches occur in the network – like the self-healing minefield – and reconfigure themselves in order to cover the terrain being monitored. Finally, sensor networks will interact with autonomous robots on the ground to seek and destroy enemy targets.

The robots used on the ground – once again interacting with a sensor network and using marker-based stigmergy -- will doubtless utilize the task allocation and specialization algorithms described earlier in order to best meet the threat presented to them through the sensor network. These specialization algorithms will almost certainly be augmented as other examples of teaming behaviour are analyzed for social insect systems. For this reason research in the appropriate area of theoretical biology should be monitored for developments. Ground-based systems will be capable of reconfiguration as mechatronic research matures. Damaged or malfunctioning units will be capable of disassembly/reassembly by the remaining units – using the sensor network to call for assistance – or autonomously. The European Union Swarm Bots project should be used as model for research undertaken here in Canada. Clearly, a multi-year project to produce military grade insect-like robots along with simulators and associated swarm intelligence learning algorithms is required here.

Soldiers of tomorrow will also benefit from stigmergy. First, soldiers' uniforms will consist of smart materials; materials that contain embedded sensors – another sensor network. Programmable materials will adaptively camouflage the wearer as he moves across the terrain. Soldiers' vital functions will be continuously monitored and communicated via the battlefield sensor network. The same sensor network will be used to relay orders that ensure optimal deployment of available resources, possibly using task allocation and division of labour algorithms as indicated above. Should a soldier become injured or incapacitated, instructions will automatically be relayed in order to have the remaining soldiers redeployed to meet the existing threat. These instructions will *not* come from a central commander rather from processing resulting from the removal of a soldier from the resources available on the battlefield. Finally, the injured soldier's sensor network will notify the sensor network that he is injured. A medic will then be guided by the sensor network to the location of the injured soldier using marker-based stigmergy.

So, what is required to achieve the above?

6.4 THE FUTURE

Very little theory exists for swarm-based systems. No robust systems should be deployed before we understand fundamental properties of stigmergic systems.

First, sensor network simulation tools need to be constructed; theoretical analysis should occur in parallel possibly providing bounds on performance when analytical closed-form solutions can not be easily obtained. The existing body of sensor network routing algorithms are either incompletely specified or analyzed; considerable work remains to be done. Scenario generators should be built in order to evaluate the effectiveness of the sensor network – the environment – in conjunction with agents whose behaviour is stigmergically-driven. In order to achieve this, an extensible, reusable agent framework should be developed that captures the patterns documented in this report, suitably augmented with existing intelligent agent algorithms for military applications. Research into the problem of combining stigmergic signals – sensor fusion – also needs to be conducted. Dr. Van Parunak's recent work along with the more general work on pheromone chemistry by Dr. White [157] may be a useful starting point. Furthermore, other stigmergic patterns should be captured and added as research in theoretical biology provides insight into other social insect behaviours.

Second, technologies for wide-spread cost-effective sensor networks need to be developed. Here, work is well advanced. Pister's work on smart dust should be leading the way here and the Centre for Embedded Network Sensors (CENS) the leading institution.

Third, intelligent materials research needs to be undertaken. Sensors woven into the fabric of clothing are relevant here. Also the work on Amorphous Computing may be of interest as it provides the potential for materials capable of self repair. Self-repairing materials have obvious applications in the autonomous repair of unmanned autonomous vehicles, for example.

Finally, research into reconfigurable and self-reproducing robots should be supported. The goal should be to understand, fabricate and deploy modules in the battlefield setting that can be used as building blocks for the repair and reproduction of unmanned autonomous vehicles in situ.

7 SUMMARY

This report has provided a taxonomy of stigmergic systems and provided the reader with several exemplar systems (patterns) that could be used to generate military applications. Owing to the importance of sensor networks in the battlefield of the future, an in depth review of routing algorithms for ad hoc networks has been provided. Brief descriptions of the more influential projects in collective robotics that use stigmergic principles have been included. Among those described the recently completed Swarm Bots project shows significant promise for the engineering of future robot swarms.

The most mature military systems using stigmergic principles – rated at TRL 8 – have been described and demonstrate conclusively that marker-based stigmergy ensures very good information fusion and processing in a battlefield scenario. Related work – referenced but not described – indicates that the systems evaluated are stable, can be effectively simulated and scale to large number of unmanned autonomous vehicles.

To conclude, the body of work on swarm intelligence found in the literature and social insects observed in nature, indicate that robust, scalable and engineering solutions to military problems can be created. What remains is the problem of developing a detailed research agenda and then funding it.

After Material: this is provided for reference only.

8 SOURCES OF INFORMATION

This section provides information on the principal people and projects related to swarm intelligence. The ordering of information in the various sections in no way indicates importance. A section on journals where information is often published is also included along with important books in the area. The area is somewhat interdisciplinary and hence books on biology are included.

8.1 PEOPLE

1. H. Van Dyke Parunak. Refer to: http://www.erim.org/~vparunak/. Involved in multi-agent system research and development; specifically, exploitation of pheromone-based communication for information management. He has also generated agent-based solutions in distributed manufacturing.

2. S. Brueckner. Refer to: http://www.erim.org/~sbrueckner/. Involved in multi-agent system research and development; specifically, exploitation of pheromone-based communication for information management. He has also generated agent-based solutions in distributed manufacturing. Ph.D. topic was pheromone-based control for distributed manufacturing.

3. Marco Dorigo. Refer to: http://iridia.ulb.ac.be/~mdorigo/ACO/ACO.html. Works on Ant Colony Optimization; inventor of the Ant Colony Metaheuristic. Project coordinator for the Swarm Bot project (see projects section). Refer to: http://www.swarm-bots.org/.

4. Thomas Stützle. Refer to: http://www.intellektik.informatik.tu-darmstadt.de/~tom/. Works on Ant Colony Optimization.

5. Eric Bonabeau. Refer to: http://www.icosystem.com. Widely quoted on swarm intelligence; principal researcher with the Icosystem corporation. Has contributed to a number of significant models in the area of swarm intelligence.

6. Jean-Louis Deneubourg. Refer to: http://www.ulb.ac.be/rech/inventaire/chercheurs/0/CH1480.html. Primary interest is biology; however, he is responsible for several important models used with the swarm intelligence community.

7. Guy Theraulaz. Refer to: http://cognition.ups-tlse.fr/~theraulaz/. Primary interest is biology and social insect modelling; however, he is responsible for several important models used with the swarm intelligence community.

8. Vittorio Maniezzo. Refer to: http://www3.csr.unibo.it/~maniezzo/. Works with Ant Colony Optimization.

9. Alcherio Martinoli. Refer to: http://swis.epfl.ch/people/alcherio/. Works on collective robotics. Martinoli's Ph.D. thesis is an *excellent* introduction to swarm robotics.

10. Ronald Kube. Refer to: http://www.cs.ualberta.ca/~kube/. Works on collective robotics. Kube has written a large number of papers on collective robotics using stigmergic principles.

11. Owen Holland. Refer to: http://cswww.essex.ac.uk/staff/holland.htm. Works on robotics; specifically robots that are inspired by principles from biology.

12. Maja Mataric. Refer to: http://www-robotics.usc.edu/~maja/. Mataric directs the USC robotics group, which is one of the top groups in the world for distributed multi-robot control.

13. Julia Parrish. Refer to: http://www.fish.washington.edu/people/parrish/. Generates models of aggregation behaviour; specifically, fish and birds. These models often motivate algorithms used in collective robotics.

14. Eiichi Yoshida. Refer to: http://staff.aist.go.jp/e.yoshida/. Works on reconfigurable robots; i.e. mechatronics.

15. Luca Maria Gambardella. Refer to: http://www.idsia.ch/~luca/. Principal investigator for the BISON project (routing in Ad Hoc and sensor networks using ant-based techniques). Works extensively with Ant Colony Optimization.

16. Rodney Brooks. Refer to: http://people.csail.mit.edu/u/b/brooks/public_html/. Prominent robotics research; inventor of the subsumption architecture; co-founder of I-Robot Corporation.

8.2 PROJECTS

1. Swarm Bots project. This project deals with the implementation of self-organizing and self-assembling artefacts. Refer to: http://www.swarm-bots.org/.

2. Intelligent Autonomous Systems Laboratory. Projects ongoing relate to robot sorting and robot self-organization. Refer to: http://www.ias.uwe.ac.uk/.

3. USC Robotics Laboratory. Refer to: http://www-robotics.usc.edu/~agents/. Projects relate to the organization of multi-robot systems using stigmergic principles. Maja Mataric is the principal investigator here.

4. Biologically inspired robotics at the University of Essex, UK. Refer to: http://cswww.essex.ac.uk/essexrobotics/index.html. This group is led by Owen Holland, formally of Caltech and the University of the West of England.

5. Swarm Intelligence Laboratory. Refer to: http://swis.epfl.ch/. This group has inherited the work of the CORO group from Caltech. It focuses on the design, modeling, control, and optimization methodologies for self-organized, collectively intelligent, distributed systems.

6. Autonomous Systems Laboratory, EPFL. Refer to: http://lsa.epfl.ch/. This group undertakes research in the area of mechatronics.

7. Caltech Center for Neuromorphic Research. Refer to: http://www.cnse.caltech.edu/. Two aspects of the group's research are modelling swarm-based, distributed robotic manipulation and distributed exploration and coverage.

8. BISON. Refer to: http://www.idsia.ch/see_prj?all=7&typ=current. This project exploits ant colony optimization and immune networks in order to understand and improve self-organization in ad-hoc and networks and grid computing systems.

9. Anthill. Refer to: http://www.cs.unibo.it/projects/anthill/. This project is devoted to the design of a framework to support the *design*, *implementation* and *evaluation* of P2P applications based on ideas such as multi-agent and evolutionary programming borrowed from ant-based systems. An Anthill system consists of a dynamic network of peer nodes; societies of adaptive agents travel through this network, interacting with nodes and cooperating with other agents in order to solve complex problems. Agents in the system interact locally, using simple behaviours; hence the systems are swarm-based.

10. NASA Autonomous NanoTechnology Swarm. Refer to: http://ants.gsfc.nasa.gov/index.html. This project is devoted to the creation of next generation deep space exploration technology.

8.3 JOURNALS, PERIODICALS

1. The Journal on Collective Robotics
2. The Journal of Robotics and Autonomous Systems
3. The Journal of Autonomous Robots
4. The Journal of Artificial Life
5. The Journal of Adaptive Behaviour
6. The Journal of Mechatronics
7. IEEE Journal on Systems, Man and Cybernetics

8.4 BOOKS

1. Dorigo M. and Stutzle T., Ant Colony Optimization, MIT Press, ISBN 0262042193, July 2004.

2. Blum C., Theoretical and Practical Aspects of Ant Colony Optimization, IOS Press, ISBN 1586034322, November 2004.

3. Kennedy J. and Eberhart R. C., Swarm Intelligence, Morgan Kaufman, ISBN 1558605959, 2001.

4. Bonabeau E., Dorigo M. and Theraulaz G., Swarm Intelligence, Oxford University Press US, ISBN 0195131584, September 1999.

5. Resnick M., Turtles, Termites, and Traffic Jams, MIT Press, ISBN 0262181622, 1997.

6. Camazine S., Deneubourg J-L., Franks N. R., Sneyd J., Theraulaz G. and Bonabeau E., Self-Organization in Biological Systems, ISBN 0691116245, Princeton University Press, November 2003.

7. Dorigo M. and Colombetti M., Robot Shaping: An Experiment in Behavior Engineering, ISBN 0262041642, MIT Press, February 2005.

8.5 WEB SITES

1. Stigmergic Systems. Refer to: http://www.stigmergicsystems.com. Good definitions and high level descriptions of stigmergy. Some good introductory material.

2. Swarm Intelligence course. Refer to: http://www.scs.carleton.ca/~arpwhite/courses/5900z. Lecture notes for a graduate course cover definitions of swarm systems, emergent computation, biological models, ant colony optimization and swarm-based problem solving. Related materials on reinforcement learning and agent-based modelling are also available.

3. Swarm Intelligence course: Refer to: http://swis.epfl.ch/teaching/SC741/. This course is taught by Alcherio Martinoli (see people section). It covers principles of swarm intelligence and self-organizing systems; collective movements in animal and human societies; foraging, trail-laying and –following, task allocation and division of labour, aggregation and segregation, self-assembling, and collaborative transportation in social insects. Microscopic and macroscopic modeling methodologies for collective systems. Machine-learning methodologies for design and optimization of collective systems are described. Combinatorial optimization algorithms (Ant Colony Optimization and Particle Swarm Optimization) based on swarm principles. Applications in automotive engineering, civil engineering, telecommunication, and operational research. Selected topics in swarm robotics and sensor networks.

4. Caltech course on Swarm Intelligence (no longer offered). Refer to: http://www.coro.caltech.edu/Courses/EE141/. A variation on this paper is: http://www.coro.caltech.edu/Courses/EE150/course.html. There is considerable overlap in the materials from [2] with this course. The authors of this course have moved on: Owen Holland is at the University of Essex; Alcherio Martinoli is now at EPFL with the Swarm Intelligence group.

5. Swarm-bots project page. Refer to: http://www.swarm-bots.org. The publications pages describe the progress of the project to date.

6. Useful site on robotics generally: http://essexrobotics.essex.ac.uk/roboticssites.html.

7. A site dedicated to robotics and mechatronics research, including people, projects, and conferences. Refer to: http://www.iee.org/OnComms/pn/mechatronics/.

8. Boids. Refer to: http://www.red3d.com/cwr/boids/. The site related to the original flocking simulation. Variations of this work still appear from time to time.

9. The Swarm Intelligence group at EPFL maintains a useful set of links to researchers and projects in the area of Collective Robotics. Refer to: http://www5.epfl.ch/swis/page1282.html.

10. Useful site on swarm intelligence maintained by Payman Arabshahi of NASA. Refer to: http://dsp.jpl.nasa.gov/members/payman/swarm/. This site contains references to people (the list overlaps with that provided in this document), introductory material, and application-related papers.

11. Useful site on sensor network research, the Center for Embedded Network Sensing: http://research.cens.ucla.edu/portal/page?_pageid=59,43783&_dad=portal&_schema=PORTAL

12. Mechatronic Research: http://www.engr.colostate.edu/~dga/mechatronics/resources.html

13. Lipsen's work on self-replication (in the mechatronic context): http://www.mae.cornell.edu/ccsl/research/selfrep/

14. PolyBot project page: http://www2.parc.com/spl/projects/modrobots/chain/polybot/

15. Amorphous Computing project page: http://www.swiss.ai.mit.edu/projects/amorphous

8.6 CONFERENCES

1. Workshop of Ant Colony Optimization and Swarm Intelligence. Held every 2 years (last was in 2004) in Brussels; organized by Marco Dorigo. Focus is Ant Colony metaheuristic, models of ant-like behaviour and applications of stigmergy. Refer to: http://iridia.ulb.ac.be/~mdorigo/ACO/conferences.html.

2. IEEE Symposium on Swarm Intelligence. Held every 2 years (next is 2005). Refer to: http://www.ieeeswarm.org/.

3. Engineering Self-organization Applications. Workshop held annually in conjunction with the Autonomous and Multi-agent System conference. Refer to: http://esoa.unige.ch/esoa05/esoa05-cfp.html for details of the 2005 workshop.

4. The Genetic and Evolutionary Computation Conference. Held annually. Contains tracks on swarm intelligence and ant colony optimization. Refer to http://www.isgec.org.

5. The Congress on Evolutionary Computation. Held annually. Contains tracks on swarm intelligence ant colony optimization. Refer to: http://www.dcs.ex.ac.uk/~dwcorne/cec2005/ for details of 2005 conference.

6. IEEE/RSJ/GI International Conference on Intelligent Robots and Systems. Held annually. Refer to: http://www.iros2005.org/ for details on the 2005 conference.

7. International Conference on Simulation of Adaptive Behaviour. Held every 2 years (last was 2004). Refer to: http://www.isab.org.uk. Conference details are linked off of this page.

8.7 COMPANIES

1. AntOptima. Refer to: http://www.antoptima.ch/. A consulting company generating optimization solutions.

2. Icosystem. Refer to: http://www.icosystem.com/. A consulting company generating optimization solutions. Has done significant work on military projects (DARPA funded) for unmanned vehicle coordination and target location.

3. Bios Group. Refer to: http://www.biosgroup.com/. A consulting company generating optimization solutions.

4. Altarum. Refer to: http://www.altarum.org/. A consulting company that generates distributed systems solutions (simulation and actual) that employ stigmergic principles. Have been involved with military projects (DARPA funded).

5. I-Robot Corporation. Refer to: http://www.irobot.com/home.cfm. Has developed robots for military usage that have been programmed using swarm intelligence.

9 BIBLIOGRAPHY

9.1 STIGMERGY

The following papers were annotated by Dylan A. Shell and use the term stigmergy in some way, either directly or through a reference.

[1] Carl Anderson. Self-organized behaviour: Case studies. Book review of Camazine et al. [18]. Complexity, 7(2):1415, 2001.
[2] Carl Anderson. The fuzzy boundaries of three self-organization-like mechanisms. Regensburg University. http://www.mbl.edu/CASSLS/andersonwh.pdf, 2002.
[3] Carl Anderson. Self-organization in relation to several similar concepts: Are the boundaries to self-organization indistinct? Biological Bulletin, 202:247255, June 2002.
[4] Ruth Aylett, Kerstin Dautenhahn, Jim Doran, Michael Luck, Scott Moss, and Moshe Tennenholtz. Can models of agents be transferred between areas? Knowledge Engineering Review, 15(2):197-203, 2000.

On page 2: Where agents are using stigmergy communication via the environment.

[5] Ralph Beckers, Jean-Louis Deneubourg, and Simon Goss. Trails and u-turns in the selection of the shortest path by the ant lasius niger. Journal of Theoretical Biology, 159:397-415, 1992.
[6] Ralph Beckers, Owen E. Holland, and Jean-Louis Deneubourg. From local actions to global tasks: Stigmergy and collective robotics. In Artificial Life IV. Proc. Fourth International Workshop on the Synthesis and Simulation of Living Systems, pages 181-189, Cambridge, Massachusetts, USA, July 1994.

The authors write (on page 1), The principle is that of stigmergy, recognised and named by the French biologist P.P Grassé [58] during his studies of nest building in termites. Stigmergy is derived from the roots stigma(goad) and ergon(work), thus giving the sense of incitement to work by the products of work. It is essentially the production of a certain behaviour in agents as a consequence of the eects produced in the local environment by previous behaviour. Also:The use of stigmergy is not conned to building structures. On page 5: The stigmergic coupling operates as follows: if a robot adds pucks to a cluster, or removes pucks from it, the consequent change in size and shape alters the probability that a subsequent random path taken by that (or any other) robot will strike the cluser frontally or tangentially, thereby affecting the probability of adding or removing further pucks in the future. On page 6:...their behaviour is influenced by the previous behaviour of the others via stigmergy, mediated through the configuration of pucks and clusters. On page 8: Perhaps stigmergy is best regarded as the general exploitation of the environment as an external memory resource; ...

[7] Claus N. Bendtsen and Thiemo Krink. Phone routing using the dynamic memory model. In

Proc. Fourth Congress on Evolutionary Computation (CEC02), pages 992-997, Honolulu, Hawaii, USA, May 2002. IEEE Press.

On page 1: This phenomenon is often found in nature, for example many ant species use lay and follow pheromone trails in the process of foraging. The indirect interaction happens by stigmergy, i.e. communication through the environment and other agents act on changes [11].

[8] Katie Bentley. Artificial forms. EASy MSc, simulation of adaptive behaviour mid-term paper, University of Sussex, UK, August 2002.

On page 3, There are many species that construct complex architectures, social insects can be seen to generate hugely intricate patterns and structures when nest building. The possible organisational mechanism put forward by Grassé in [58] to explain how this can occur is stigmergy. The basic idea is that the coordination of individuals tasks depends not on any communication between them but on the nest structure itself [13]. On page 4: Stigmergic swam intelligence occurs when the innate rules require that the agents give off pheromone and are also attracted to areas marked with pheromones [12].

[9] Greg Biegel. Cooperation through the Environment: Stigmergy in CORTEX, chapter 4, pages 3138. Information Society Technologies Research Project IST-2000-26031, March 2002.

On page 2 and 31 (in the introductory chapter, and abstract to chapter 4), he writes: The technical term stigmergy was coined by a biologist who used it to describe the co-ordination of populations of insects without direct communication between the individuals. On page 33, Stigmergy, or coordination of actions through the environment, was rst observed in colonies of insects [63]. Stigmergy as a coordination mechanism is characterised by a lack of planning using explicit communication between entities, a fact that makes it extremely flexible and robust in large systems. On page 33 (in section 4.2, Stigmergy Defined) The term stigmergy was first introduced by Pierre Paul Grassé , a French entomologist, in 1959. He used the term to describe the coordination of the activities of ants in carrying out complex activities, such as nest building, without direct communication amongst themselves. It is evident that stigmergy describes a form of asynchronous interaction and information interchange between entities mediated by an active environment [130].

[10] Eric Bonabeau, Guy Theraulaz, Jean-Louis Deneubourg, Serge Aron, and Scott Camazine. Self-organization in social insects. Trends in Ecology & Evolution, 12(5):188-193, May 1997.

This paper describes many of the folklore instances of collective phenomena in insects. This article did not use the word stigmergy; and does not discuss methods of communication, but rather a self-organisational framework from which to view these results.

[11] Eric Bonabeau, Marco Dorigo, and Guy Theraulaz. Swarm Intelligence: From Natural to Artificial Systems. Oxford University Press, Inc., New York, USA, 1999.

[12] Eric Bonabeau, Marco Dorigo, and Guy Theraulaz. Inspiration for optimization from social insect behaviour. Nature, 406:39-42, July 2000.

[13] Eric Bonabeau, Sylvain Guerin, Dominique Snyers, Pascale Kuntz, and Guy Theraulaz. Three dimensional architectures grown by simple stigmergic agents. Biosystems, 56(1):157, March 2000.

[14] Eric Bonabeau. Editor's introduction: Stigmergy. Artificial Life, 5(2):95-96, 1999.

[15] Josh C. Bongard. Evolved sensor fusion and dissociation in an embodied agent. In Proc. EPSRC/BBSRC International Workshop Biologically-Inspired Robotics: The Legacy of W. Grey Walter, pages 102109, Bristol, UK, August 2002.

The authors only use stigmergy in reference to the work of Holland and Melhuish [63] and Dorigo and Di Caro [36].

[16] Sven A. Brueckner and H. Van Dyke Parunak. Swarming agents for distributed pattern detection and classication. In Proc. Workshop on Ubiquitous Computing, First Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 02), Bologna, Italy, August 2002.

On page 1, one of the most powerful global coordination mechanisms in distributed biological systems is stigmergy [58], from the Greek words stigma sign and ergos work. The work performed by the agents in the environment in turn guides their later actions a feedback loop that establishes dynamic information flows across the population and guides its operation. On page 3, Marker-based stigmergy in social insect colonies uses chemical markers (pheromones) that the insects deposit on the ground in specific situations (e.g. food found). Multiple deposits at the same location aggregate in strength. The paper uses pheromone computational techniques.

[17] O.H. Bruinsma. An analysis of building behaviour of the termite Macrotermes subhyalinus (Rambur). PhD thesis, Leandbouwhoge school, Wageningen, Netherlands, October 1997.

[18] Scott Camazine, Jean-Louis Deneubourg, Nigel R. Franks, James Sneyd, Guy Theraulaz, and Eric Bonabeau. Self-Organization in Biological Systems. Princeton University Press, 2001.

[19] Y. Uny Cao, Alex S. Fukunaga, and Andrew B. Kahng. Cooperative mobile robotics: Antecedents and directions. Autonomous Robots, 4(1):723, March 1997.

Mentions stigmergy in the footnote on page 2 presents the definition from Beckers et al. [6], then further states: This is actually a form of cooperation without communication, which has been the stated object of several foraging solutions since the corresponding formulations become nearly trivial if communication is used.

[20] Fabrice Chantemargue and Beat Hirsbrunner. A collective robotics application based on emergence and self-organization. In Proc. Fifth International Conference for Young Computer Scientists (ICYCS 99), Nanjing, China, August 1999.

On footnote on page 2: To out knowledge, Beckers et al. [6] were the first to exploit stigmergic coordination between robots. Stigmergic coordination means literally incitement to work by the products of the work.

[21] Vincent A. Cicirello and Stephen F. Smith. Ant colony control for autonomous decentralized shop floor routing. In Proc. Fifth International Symposium on Autonomous Decentralized Systems (ISADS 01), pages 383390, Dallas, Texas, USA, March 2001.

On page 2, they write: Stigmergy is a process by which the common environment is altered by individuals of the ant colony and this dynamically changing environment is used for self-organization and coordination within the colony. Also on the same page, they organize such activities through indirect communication known as stigmergy by which individuals alter the environment and this ever-changing environment is used for coordinating the colony's activities.

[22] Antonio DAngelo and Enrico Pagello. Using stigmergy to make emerging collective behaviours. In Proc. Eighth Conference of the Italian Association for Artificial Intelligence, Siena, Italy, September 2002.

On page 6 (note, the text has been reproduced accurately) to this aim we must require the acting flow of an individual robot can partially enter another individual as a perceptual flow, namely, both robots should interact each other allowing the action of the former to affect the behaviour of the latter. In biology literature such kind of interaction is known as stigmergy, namely, cooperation without communication or cue-based communication [84], where the action of an individual robot is a suggestion for another issue some specific behaviour.

[23] Arindam K. Das, Robert J. Marks II, Mohamed El-Sharkawi, Payman Arabshahi, and Andrew A. Gray. The minimum power broadcast problem in wireless networks: An ant colony system approach. In Proc. IEEE CAS Workshop on Wireless Communications and Networking, Pasadena, California, USA, September 2002.

On page 1: ..stigmergy, or communication through the environment. An example is pheromone laying on trails followed by ants.... This causes an autocatalytic reaction, i.e. one that is accelerated by itself.

[24] Kerstin Dautenhahn and Bruce Edmonds. Social embeddedness origins, occurrence and opportunities. A half-day tutorial presented at the Seventh International Conference on Simulation of Adaptive Behaviour (SAB2002) at the University of Edinburgh, August 2002.

On page 27, the concept of stigmergy describes a class of mechanisms mediating animal-animal interactions... Also, Stigmergy is based on indirect communication, communication via the environment, and an example of collective behaviour.

[25] Kerstin Dautenhahn. Embodiment and interaction in socially intelligent life-like agents. In Chrystopher L. Nehaniv, editor, Computation for Metaphors, Analogy and Agent, volume 1562 of Springer Lecture Notes in Artificial Intelligence, pages 102142. Springer, 1999.

On page 6: Deneubourg and his colleagues [28] give an impressive example where a group of robots ant-like robots collectively solves a sorting task. Their work is based on a mode of how ants behave, using the principle of stigmergy which is defined as the production of a certain behaviour in agents as a consequence of the effects produced in the local environment by previous behaviour [6]. On page 7: Behaviour based research on the principle of stigmergy is not using explicit representations of goals, the dynamics of group behaviour are emergent and self-organizing.

[26] Kerstin Dautenhahn. Evolvability, culture and the primate social brain. In Proc. Evolvability Workshop, Seventh International Conference on the Simulation and Synthesis of Living Systems (Artificial Life 7), pages 23-26, Portland, Oregon, USA, August 2000.

On page 3 (the typographical errors are reproduced), The concept stigmergy was first developed by the French zoologist Pierre-Paul Grassé in order to understand the emergence of regulation and control in social insect societies. Stigmergy is a class of mechanisms mediating animal-animal interactions [126]. According to Bonabeau et al. [11] and Theraulaz and Bonabeau [126] two of such mechanisms are quantitative stigmergy and self-organised dynamics and qualitative stigmergy and self-assembling dynamics. Generally, the behaviour of each insect can be described as a stimulus-response (S-R) sequence (even for solitary species). If animals to not distinguish between products of others' activities and their own activity, then individuals can respond to and interact through stimuli. She then provides examples of quantitative stigmergy [where] stimuli in the S-R sequence [are] different quantitatively using the standard examples of pheromones (in this case termite nest construction from pheromone impregnated soil pellets). The qualitative stigmergy [wherein] we have a discrete set of stimuli types example is of wasp nest construction. The elements of positive feedback (and negative feedback, and amplification of fluctuations) are considered as elements of self-organisation (as consistent with Bonabeau et al. [11]) she also presents a view stressing stigmergy when combined with self-organisation as a keystone of swarm intelligence.

[27] Kerstin Dautenhahn. Reverse engineering of societies - a biological perspective. In Proc. AISB Symposium Starting from Society - the application of social analogies to computational systems, a symposium at Time for AI and Society2000 Convention of the Society for the Study of Artificial Intelligence and the Simulation of Behaviour (AISB-00), pages 23-26, University of Birmingham, UK, April 2000.

[28] Jean-Louis Deneubourg, Simon Goss, Nigel R. Franks, Ana B. Sendova-Franks, Claire Detrain, and Laeticia Chretien. The dynamics of collective sorting robot-like ants and ant-like robots. In Proc. First International Conference on Simulation of Adaptive Behaviour (SAB1990), pages 356-363, Paris, France, September 1990.

This paper does not use the word stigmergy, but is referred to as using a stigmergic approach. Furthermore, on page 7, another example of how such indirect communication can organise the activity of a group of non-communicating agents can be seen in the way the foragers of some species of ants set up individual and non-overlapping foraging territories.

[29] Reimundo Deveza, David Thiel, R. Andrew Russell, and Alan Mackay-Sim. Odor sensing for robot guidance. International Journal of Robotics Research, 13(3):232-239, 1994.

[30] Gianni Di Caro and Marco Dorigo. Antnet: A mobile agents approach to adaptive routing. Technical Report IRIDIA 97-12, Institut de Recherches Interdisciplinaires et de Developpements en Intelligence Articielle, Universite Libre de Bruxelles, 1997.

Page 1 has the following: The term stigmergy was first introduced by Grassé [58] to describe the indirect communication taking place among individuals through modifications induced in their environment.

[31] Gianni Di Caro and Marco Dorigo. An adaptive multi-agent routing algorithm inspired by ants behaviour. In Proc. Fifth Annual Australasian Conference on Parallel and Real-Time Systems (PART98), pages 261272. Springer-Verlag, 1998.

On page 1, This effort is mediated by stigmergic communication [11; 58], that is, a form of indirect communication of information on the problem structure ants collect while building solutions. On page 4, The key concept in the cooperative aspect lies in the indirect and non-coordinated way communication among ants happens (stigmergy) [58].

[32] Gianni Di Caro and Marco Dorigo. Ant colonies for adaptive routing in packet-switched communications networks. Lecture Notes in Computer Science, 1498:673-682, 1998.

Page 1: ...stigmergy [58], that is, the indirect communication taking place among individuals through modifications induced in their environment.

[33] Gianni Di Caro and Marco Dorigo. Antnet: Distributed stigmergetic control for communications networks. Journal of Artificial Intelligence Research, 9:317-365, 1998.

In the abstract: The communication among the agents is indirect and asynchronous mediated by the network itself. This form of communication is typical of social insects and is called stigmergy. On page 2, ...the notion of stigmergy [58], that is, the indirect communication taking place among individuals through modifications induced in their environment.

[34] Gianni Di Caro and Marco Dorigo. Mobile agents for adaptive routing. In Proc. 31 st Hawaii International Conference on System Science (HICSS-31), Big Island of Hawaii, USA, January 1998.

On page 1, ...the notion of stigmergy, introduced by Grassé [58] to describe the indirect communication taking place among individuals through modifications induced in their environment. On page 4(grammar preserved as is): The information locally stored and updated at each network node defines the agent input state. Each agent uses it to realise the next node transition and, at the same time, it will modify it, modifying in this way the local state of the node as seen by future agents. This specific form of indirect communication through the environment with no explicit level of agents' coordination is called stigmergy [58; 114; 120]. Active stigmergy occurs when an agent alters the environment so as to affect the input of another agent, passive stigmergy occurs when an agent

alters the environment in such a way that the effect of the actions of the other agents is no more the same.

[35] Gianni Di Caro and Marco Dorigo. Two ant colony algorithms for best-effort routing datagram networks. In Proc. Tenth IASTED International Conference on Parallel and Distributed Computing and Systems (PDCS98), Las Vegas, Nevada, October 1998.

On page 1, ...the notion of stigmergy [58; 11], that is, the indirect communication taking pace among individuals through local, persistent (or slowly changing) modifications induced in their environment.

[36] Marco Dorigo and Gianni Di Caro. The ant colony optimization meta-heuristic. In David Corne, Marco Dorigo, and Fred Glover, editors, New Ideas in Optimization, pages 11-32. McGraw-Hill, London, 1999.

On the first page: The emergence of this shortest path selection behaviour can be explained in terms of autocatalysis (positive feedback) and differential path length, and it is made possible by an indirect form of communication, known as stigmergy [58] mediated by local modifications of the environment.

[37] Marco Dorigo, Gianni Di Caro, and Luca M. Gambardella. Ant algorithms for discrete optimization. Artificial Life, 5(2):137172, 1999.

Page 3, ...it is only the ensemble of ants, that is the ant colony, which presents the shortest path finding behaviour. In a sense, this behaviour is an emergent property of the ant colony. It is also interesting to note that ants can perform this specific behaviour using a simple form of individual communication mediated by pheromone laying, known as stigmergy [58]. On page 3, As defined by Grassé in his work on Bellicositermes Natalensis and Cubitermes [58], stigmergy is the stimulation of workers by the performance they have achieved. Further they mention (on footnotes on pages 3-4) that Workers are one of the castes in termite colonies. Although Grassé introduced the term stigmergy to explain the behaviour of termite societies, the same term has been used to describe indirect communication mediated by modification of the environment that can be observed also in other social insects.Page 4:In fact, Grassé [57] observed that insects are capable to respond to so called significant stimuli which activate a genetically encoded reaction. In social insects, of which termites and ants are some of the best known examples, the effects of these reactions can act as new significant stimuli for both the insect that produced them and for other insects in the colony. The production of a new significant stimulus as a consequence of the reaction to a significant stimulus determines a form of coordination of the activities and can be interpreted as a form of indirect communication. For example, Grassé [58] observed that Bellicositermes Natalensis as well as Cubitermes, when building a new nest, start by a random, non coordinated activity of earth-pellet depositing. But once the earth-pellet reach a certain density in a restricted area they become a new significant stimulus which causes more termites to add earth-pellet so that pillar and arches, and eventually the whole nest, are built. Further on page 4: What characterizes stigmergy from other

means of communication is (i) the physical nature of the information released by the communicating insects, which corresponds to a modification of physical environmental states visited by the insects, and (ii) the local nature of the released information, which can only be accessed by insects that visit the state in which it was released (or some neighbourhood of that state). Accordingly, in this paper we take the state that it is possible to talk of stigmergetic communication whenever there is an indirect communication mediated by physical modifications of environmental states which are only locally accessible by the communicating agents.

[38] Marco Dorigo, Eric Bonabeau, and Guy Theraulaz. Ant algorithms and stigmergy. Journal of Future Generation Computer Systems, 16(8):851-871, June 2000.

On page 1: A particularly interesting body of work is the one that focuses on the concept of stigmergy, a particular form of indirect communication used by social insects to coordinate their activities. Page 2: The term stigmergy was introduced by Grassé [58] to describe a form of indirect communication mediated by modification of the environment that he observed in two species of termites: Bellicositermes Natalensis and Cubitermes. Grassé s original definition of stigmergy was: Stimulation of workers by the performance they have achieved. Although Grassé first introduced the term stigmergy to explain the behaviour of termites societies, the same term has later been used to indicate indirect communication mediated by modifications of the environment that can be observed also in other social insects [126].The authors then explain the nest construction as observed by Grassé [58]. Soil pellets are impregnated with pheromones, which the termites can sense. The authors mention that this example has a positive feedback mechanism, since once a critical quantity of pellets has been amassed in a pillar, workers add to that pillar, making it more salient, and hence the region of renewed activity, further increasing the saliency; and so the cycle continues. They use the term autocatalytic to describe this process. This example contains what Bonabeau et al. [11] consider the hallmarks of self-organisation. Attempting to pry apart, the components that are considered necessary for characteristics for stigmergy (as distinct self-organisation, or self-organisation with stigmergy), one may be aided in the artificial stigmergy these authors describe: on pages 3-4, the write: The implementation of ant algorithms is made possible by the use of so-called stigmergic variable, i.e., variables that contain the information used by artificial ants to indirectly communicate. In some cases ... the stigmergic variable is a specifically defined variable used by ants to adaptively change the way they build solutions to the considered problem. In other cases ... the stigmergic variable is one of the problem variables: in this case a change in its value determines not only a change in the way a solution to the problem is built, but also a direct change in the solution of the problem itself. They also refer to algorithms built using these ideas as subscribing to the stigmergic paradigm.

[39] Marco Dorigo, Gianni Di Caro, and Thomas Stutzle. Guest editorial: Ant algorithms. Journal of Future Generation Computer Systems, 16(8):v-vii, June 2000.

On page 1, As Dorigo et al. [38] point out ... the indirect stigmergic communication among ants is the key characteristic of ant algorithms. Stigmergy defines a paradigm of indirect and asynchronous communication mediated by an environment. While carrying out their own tasks, ants deposit some chemical substance (called pheromones) or induce some other physical modifications of the environment. These modifications change the way the environment (and in a way, the problem under consideration) is sensed by other ants in the colony, and implicitly act as signals triggering other ants behaviours that again generate new modification that will stimulate other ants and so on.

[40] Holly A. Downing and Robert L. Jeanne. Nest construction by the paper wasp, polistes: a test of stigmergy theory. Animal Behaviour, 36:1729-1739, 1988.

[41] Holly A. Downing and Robert L. Jeanne. The regulation of complex building behaviour in the paper wasp, polistes fuscatus (insecta, hymenoptera, vespidae). Animal Behaviour, 39:105-124, 1990.

[42] Johann Dreo and Patrick Siarry. A new ant colony algorithm using the heterarchical concept aimed at optimization of multiminima continuous functions. In Proc. Third International Workshop on Ant Algorithms (ANTS 2002), pages 216228, Brussels, Belgium, September 2002.

In the abstract: The original idea consisted in simulating the stigmergic communication, therefore these algorithms are considered as a form of adaptive memory programming. On page 2: stigmergy, which is a form of indirect communication mediated by modifications of environment.

[43] Arnaud Dury, Guillaume Vakanas, Christine Bourjot, Vincent Chevrier, and Bertrand Krat. Multi-agent simulation to test a coordination model of the prey capture in social spiders. In Proc. Thirteenth Annual European Simulation Symposium (ESS01), pages 831833, Marseille, France, October 2001.

On page 1: Stigmergy [58] can be described as the animal creates, by his activity, a structure possessing a particular stimulating value that triggers a specific response of other members of the group. This is a way to achieve the coordination without any explicit reference to the tasks being, performed by any spider of the system: past actions put traces (inscriptions) in the environment.

[44] Serge Fenet and Salima Hassas. A.N.T: a distributed problem-solving framework based on mobile agents. In Advances in Mobile Agents Systems Research (MAA2000), Proc. 12th International Conference on System Research, Informatics & Cybernetics, pages 3944, July 2000.

On page 2: ...agents do not communicate directly; only the trails perceived in the environment will influence an agents behaviour. This indirect communication mechanism, called stigmergy by P. P. Grassé [58], is the root of the emerging complex properties of systems made of unintelligent reactive agents. The system does a form of pheromone computing.

[45] Serge Fenet and Salima Hassas. A distributed intrusion detection and response system based on mobile autonomous agents using social insects communication paradigm. In Proc. First International Workshop on Security of Mobile Multi-agent Systems (SEMAS-2001), Fifth International Conference on Autonomous Agents (Agents 2001), Montreal, Canada, May 2001.

Page 2, ...swarms of simple agents that interact with and through their environment. This interaction is mainly based on a marking mechanism and is using a chemical substance called pheromone. This indirect communication mechanism, called stigmergy in Grassé [58], is the root of the emerging complex properties of the systems made of unintelligent reactive agents. Again this is a form of pheromone computation, and the authors seem to feel that pheromones are necessary.

[46] Serge Fenet and Salima Hassas. A distributed intrusion detection and response system based on mobile autonomous agents using social insects communication paradigm. In Klaus Fischer and Dieter Hutter, editors, Electronic Notes in Theoretical Computer Science, volume 63. Elsevier, 2002.

This is identical to Fenet and Hassas [45].

[47] Paola Flocchini, Giuseppe Prencipe, Nicola Santoro, and Peter Widmayer. Distributed coordination of a set of autonomous mobile robots. In Proc. IEEE Intelligent Vehicles Symposium (IV2000), pages 480485, Dearborn, Mississippi, USA, 2000.

On page 1: Moreover, there are no explicit direct means of communication: the communication occurs in a totally implicit manner, through the environment (in biology, this communication is called stigmergy [6; 58]).

[48] Terrence Fong, Illah Nourbakhsh, and Kerstin Dautenhahn. A survey of socially interactive robots. Special issue on Socially Interactive Robots, Robotics and Autonomous Systems, 42(3-4):143166, 2003.

On page 2: ..principles such as stigmergy (indirect communication between individuals via modifications made to the shared environment)...Stigmergy was first described by Grassé as a way to explain how social insect societies can collectively produce complex behaviour patterns and physical structures, whilst each individual seemingly works on its own[11], Also cites the following as experiments on stigmergy in simulated and physical robots: Deneubourg et al. [28]; Beckers et al. [6]; Kube and Bonabeau [72]; Melhuish et al. [87]; Krieger et al. [69].

[49] Noria Foukia, Salima Hassas, and Serge Fenet. An intrusion response an intrusion response scheme: Tracking the source using the stigmergy paradigm. In Proc. Second International Workshop on Security of Mobile Multi-agent Systems (SEMAS-2001), First Joint Conference on Autonomous Agents and Multi-agent Systems (AAMAS 02), pages 1826, Bologna, Italy, July 2002.

Page 3, Cooperation in these systems is mediated by an efficient mechanism of communication through the inscription of task evolution in the environment. This paradigm introduced for the first time by P.P. Grassé in [58] described the way social insects communities (ants, termites, bees, ...) interact through their environment. Later on the same page, after discussing the construction of pheromone gradients, and reinforcement effects, they write: This indirect communication between different members of the colony through the environment is called stigmergy.

[50] Lewis Girod and Deborah Estrin. Robust range estimation using acoustic and multimodal sensing. In Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2001), Maui, Hawaii, October-November 2001.

Page 1, Modification of the environment for communication and collaboration, sometimes known as stigmergy, is a very powerful tool used to great advantage by some of the most successful animals on earth [58]. In the context of robotics, the idea of stigmergy can be take a step further if the robot implants active devices into the environment that can enable it to operate more efficiently.

[51] Andreas Goebels. Kommunikation in mas. University of Paderborn Talk, http://www.uni-paderborn.de/cs/ag-klbue/sta/labeo/work/publications/talks/MASKommunikation.ppt, May 2003.

This is a presentation partly in German partly in English. Slide 4, seems to indicate that stigmergy is a form of indirect communication (as is Blackboard communication). Slide 7 paraphrases Grassé 's definition as no supervisor and no direct communication necessary. Further listing Indirect communication and Environment is media. Slide 9 shows recognition of both quantitative stigmergy (mentions higher stimulation more reaction and gives ant path minimization as an example) and qualitative stigmergy (writing stimulation typereaction.)

[52] Wan Tsin Goh and Zhengwen Zhang. Multi-agent system for dynamic production scheduling and optimization. In Proc. Third International Symposium on Multi-Agent Systems, Large Complex Systems and E-Businesses (MALCEB03), Germany, October 2002.

Page 3, stigmergy [101], i.e. a coordination mechanism found in insect society.

[53] Dani Goldberg and M. Mataric. Interference as a guide to designing efficient group behaviour. Technical Report CS-96-186, School of Computer Science, Brandeis University, 1996.

Uses the word in reference to the foraging done by Holland and Melhuish [63].

[54] Claudia V. Goldman and Jeffrey S. Rosenschein. Emergent coordination through the use of cooperative state-changing rules. In Proc. of the Twelfth National Conference on Artificial Intelligence, pages 408413, Seattle, Washington, USA, July 1994.

This paper does not make use of the word stigmergy, but is cited in [121] because it talks of agents exhibiting cooperative behaviour by altering the environment to make it easier for others to get their jobs done.

[55] David Gordon. Ant-based path finding. Final Year Undergrad Projects, School of Computing, Leeds University, 2002.

Pages 14-15 have a discussion of definition of stigmergy. He explains the process of implicit communication in the nest-building behaviour of termites as observed by Grassé [58]. The discussion follows closely to that presented in Bonabeau et al. [11]. Earth pellets have a dual purpose in this scheme they are both the raw materials of the construction, and a means of communication...Also mentions (the first place in written form I am aware of) the uses of paths by humans as being stigmergy. On page 16, he mentions the wasp division of labour presented in Bonabeau et al. [11] as a form of qualitative/discrete stigmergy, and thus recognises the existence also of quantitative stigmergy. The author also uses stigmergically as an adverbial form of the word.

[56] David Gordon. Collective intelligence in social insects. AI Depot online Web Essay, http://aidepot.com/Essay/SocialInsects.html, January 2003.

Page 1, The termite's actions ... rely on how the termites world appears at any given moment. [The termite] just needs to invoke a simple behaviour dependent on the state of its immediate environment. Grassé termed this stigmergy, meaning incite to work.

[57] Pierre-Paul Grassé . Les insects dans leur univers. Ed. du Palais de la decouverte, 1946.

[58] Pierre-Paul Grassé . La Reconstruction du nid et les Coordinations Inter-Individuelles chez Bellicositermes Natalensis et Cubitermes sp. La theorie de la Stigmergie: Essai dinterpretation du Comportement des Termites Constructeurs. Insectes Sociaux, 6:4181, 1959.

[59] Pierre-Paul Grassé . Termitologia, tome II. Fondation des societes. Construction, 1984.

[60] Mesut Gunes, Udo Sorges, and Imed Bouazizi. Ara the ant-colony based routing algorithm for manets. In International Conference on Parallel Processing Workshops (ICPPW 02), Vancouver, B.C., Canada, August 2001.

On page 1, they write: The interesting point is, that the ants do not need any direct communication for the solution process, instead they communicate by stigmergy. The notion of stigmergy means the indirect communication of individuals through modifying their environment.

[61] Marc Heissenbuttel and Torsten Braun. Ants-based routing in large-scale mobile ad-hoc networks. In Proc. of Kommunikation in verteilten Systemen (KiVS 03), pages 9199, Leipzig, Germany, February 2003.

Page 2: Ants were shown to find shortest paths through a process called stigmergy, which could be described as indirect communication between individuals through the environment. Ants returning from a food source to the

nest lay down pheromones (a chemical substance) behind them. Other ants are attracted by these pheromones trails and in turn reinforce them even more. While mention is made of the positive feedback loop, no explicit mention is made of it as a requirement for stigmergy.

[62] Marc Heissenbuttel. Ants-based routing in large-scale mobile ad-hoc networks. In University of Berne Summer School of the Computer Networks and Distributed Systems Research Group, pages 7-9, Vira-Gambarogno Ticino, Switzerland, August 2002.

Page 1: Ants were shown to find shortest paths through a process called stigmergy, which could be described as indirect communication between individuals through the environment. He further gives and example of ant path following through pheromones.

[63] Owen E. Holland and Chris Melhuish. Stigmergy, self-organization, and sorting in collective robotics. Artificial Life, 5(2):173-202, 1999.

[64] Owen E. Holland. Multi-agent systems: Lessons from social insects and collective robotics. In Sandip Sen, editor, Working Notes for the AAAI Symposium on Adaptation, Co-evolution and Learning in Multi-agent Systems, pages 57-62, Stanford University, CA, USA, March 1996.

[65] Istvan Karsai. Decentralized control of construction behaviour in paper wasps: An overview of the stigmergy approach. Artificial Life, 5(2):117136, 1999.

[66] Hadeli Karuna, Paul Valckenaers, Constantin Bala Zamrescu, Hendrik Van Brussel, Bart Saint Germain, Tom Holvoet, and Elke Steegmans. Self-organising in multi-agent coordination and control using stigmergy. In Engineering Self-Organising Applications (ESOA 2003), pages 5361, Melbourne, Australia, July 2003.

On pages 2-3: In general, there exist two major classes of co-ordination between agents; coordination by direct communication, and coordination within dissipative fields (indirect interaction). This research uses coordination within dissipative fields. It utilizes an approach inspired by the way in which ant colonies propagate information while foraging for food; this is called stigmergy by biologists [126]. Stigmergy describes the use of asynchronous interaction and information exchange between agents mediated by an active environment. Investigations of social insect societies show that they coordinate themselves by producing such a dissipative field in their environment.

[67] Ioannis N. Kassabalidis, Mohamed A. El-Sharkawi, Robert J. Marks II, Payman Arabshahi, and Andrew A. Gray. Swarm intelligence for routing in communication networks. In Proc. IEEE Globecom 2001, San Antonio, Texas, USA, November 2001.

On page 1: Swarm intelligence, in particular, uses stigmergy (i.e. communication through the environment) for agent interaction [58; 11; 30; 75].On page 2: The main principle behind these interactions is called stigmergy, or communication through the environment. An example is pheromone laying on trails followed by ants... This causes an autocatalytic effect, i.e., one that is accelerated by itself. Later on the same page: Another form of stigmergy alters the environment in such a manner as to promote further similar action by the agents. This process is

dubbed task-related stigmergy. An example is sand grain laying by termites when constructing nests [11]. On page 4:. .. and stigmergy, or communication through the environment.

[68] Nupur Kothari, Vartika Bhandari, and Dheeraj Sanghi. Query localization using pheromone trails: A swarm intelligence based approach. In Proc. of Indian National Conference on Communications (NCC03), Chennai, India, 2003.

On page 2, ...individuals communicate with each other through their environment, via the mechanism of stigmergy [58].

[69] Michael J. B. Krieger, Jean-Bernard Billeter, and Laurent Keller. Ant-like task allocation and recruitment in cooperative robots. Nature, 406:992995, August 2000.

This document is been cited by Fong et al. [48] in the contest of a stigmergic approach. This document never uses the word stigmergy.

[70] Krishnan Krishnaiyer and S. Hossein Cheraghi. Ant algorithms: Review and future applications. In Proc. Industrial Engineering Research Conference (IERC 02), Orlando, Florida, USA, May 2002.

On page 2, ..individuals communicate with each other through their environment, via the mechanism of stigmergy [58].

[71] Lars Kroll Kristensen. Aintz: Collective problem solving by artificial ants. Masters thesis Dept. of Computer Science, University of Aarhus, Denmark, 2000.

On page 22: Stigmergy was a concept introduced by Grassé [58] to explain the apparent coordination of termite building behaviour. Grassé defines Stigmergy like this: The coordination of tasks and the regulation of constructions does not depend directly on the workers, but on the constructions themselves. The worker does not direct his work, but is guided by it. It is to this special form of stimulation that we give the name STIGMERGY (stigma: Wound from a pointed object; ergon: Work, product of labour = stimulating product of labour).On page 23, Stigmergy is thus a conceptual model for explaining cooperation achieved by simple agents working in a (complex) environment.Tthe explanation is primarily based on the concept of the work guiding the workers. On page 24, the concept of stigmergy can be divided into qualitative and quantitative stigmergy. The difference between the two types lies in what the agents respond to. Later, on page 26, she mentions, that the environment is then used as a global master memory. This implicit use of the environment as memory is very important to the concept of stigmergy...

[72] Claus Ronald Kube and Eric Bonabeau. Cooperative transport by ants and robots. Robotics and Autonomous Systems, 30(1/2):85-101, 2000.

Page 8: Coordination in collective transport seems to occur through the item being transported: a movement of one ant engaged in group transport is likely to *modify the stimuli perceived by the other group members, possibly producing, in turn, orientational or positional changes in these ants. This is an example of stigmergy [58], the coordination of activities through indirect interactions. Page 16: A termite's nest ... is constructed through a series of building steps. Each construction phase is thought to be governed by a building program with step transition specified as stimulus cues. In fact, this communication through the environment is the basis of Grassé s Stigmergy Theory [58].On Page 22: Stigmergy, a term coined by French biologist P. Grassé , which means to incite work by the effect of previous work [58] is a principle finding its way from the field of social insects to collective robotics [6; 125].On Page 32, Stigmergy as proposed by Grassé is a model used to explain the regulation of building behaviour in termites [58]. Stigmergy theory holds that transitions between a sequence of construction steps is regulated by the effect of previous steps. In more general terms, the theory has been used to explain and describe the process by which task activity can be regulated using only local perception and indirect communication through the environment as applied to algorithms for coordinating distributed building behaviour [125] and foraging tasks by multi-robot systems [6]. In the box-pushing task the results support the use of indirect communication through the environment as proposed by stigmergy theory. However, Downing and Jeanne found that stigmergy theory does not explain the use of additional cues, not dependent on previous steps, in regulating task execution in nest construction by papers wasps [40].*

[73] Claus Ronald Kube. Collective Robotics: From Local Perception to Global Action. PhD thesis, University of Alberta, 1997.

The thesis asks (and answers) the question, is explicit communication necessary for cooperation. On pages 22, [The termite nests] construction, through a linear series of building steps, is hypothesized to be the result of a building program and stimulus cues used to switch between construction steps, and forms the basis of Grassé 's Stigmergy Theory [58]. While he does not mention the word stigmergy directly, in other places (pages 37-40, 93, etc) he mentions various biological (insect) examples of cooperation (e.g. nest building, cooperative transport.) Regarding these examples he writes on page 93: ...results in predictable global action without resorting to directly communicating building intentions between ants. Rather, indirect communication through the task itself is sufficient to produce a coherent behaviour. On page 95: Stigmergy, a term coined by French biologist P. Grassé , which means to incite work by the effect of previous work [58] is a principle finding its way from the field of social insects to collective robotics [6; 125] ... which result from common task coordination that does not appear to depend on interaction between the agents by rather on the object they act upon. On page 115, Data presented in this study also agrees in certain aspects with other studies in which stigmergy is used as the task coordinating mechanism. Stigmergy as proposed by Grassé is a model used to explain the regulation of building behaviour in termites [58]. Stigmergy theory holds that transitions in a sequence of construction steps are regulated by the effect of previous steps. In more general terms, the theory has been used to *explain and describe the process by which task activity can be regulated using only local perception and indirect communication through the environment as applied to algorithms for coordinating distributed building behaviour [125] and foraging tasks by multi-robot systems [6]. In the box-pushing task the results support the use of indirect communication through the environment as proposed by stigmergy theory. He further explains that it has been found (in Downing and Jeanne [40]) that this previous step implies next step type description does not succeed for construction by paper wasps (they use perceptual cues from other stimuli). A processing hierarchy for the multiple cues reduces the number of cues that need to be evaluated at the same time [41]. On page 116: Hence, stigmergy theory would have to be expanded to include both additional and multiple cues which may adapt to the environment as proposed in Downing and Jeanne [40].*

[74] Jesper Bach Larsen. Specialization and division of labour in distributed autonomous agents. Masters thesis, Dept. of Computer Science, University of Aarhus, Denmark, 2001.

[75] Steen Lipperts and Birgit Kreller. Mobile agents in telecommunications networks a simulative approach to load balancing. In Proc. Fifth Information Systems Analysis and Synthesis (ISAS99), Orlando, Florida, USA, July 1999.

This paper does not appear to use the word, but is referenced by Kassabalidis et al. [67], which claims that this paper uses a stigmergic approach.

[76] Michael L. Littman. Algorithms for Sequential Decision Making. PhD thesis, Brown University, 1996.

The reference to stigmergy is in the footnote on page 135. Here, while in section entitled Register Memory he makes mention of saving information in non-volatile memory, and for retrieving this information at a later time, claiming that this type of memory can be viewed as a form of stigmergy [6]. The idea behind stigmergy is that the actions of an agent change the environment in a way that affects later behaviour resulting in a form of external memory.

[77] Alcherio Martinoli and Francesco Mondada. Collective and cooperative group behaviours: Biologically inspired experiments in robotics. In Proc. Fourth International Symposium on Experimental Robotics (ISER-95), Stanford, California, USA, June-August 1995.

On page 2, they write: Beckers et al. [6] made the same experiment with robots... the collective behaviour was analysed on the basis of the stigmergy principle, which signifies incitement to work by the products of the work. It consists in essentially the production of certain behaviour in the agents as a consequence of the effects produced in the environment by previous actions.

[78] Alcherio Martinoli and Francesco Mondada. Probabilistic modelling of a bio-inspired collective experiment with real robots. In Proc. Third International Symposium on Distributed Autonomous Robotic Systems (DARS-98), Karlsruhe, Germany, May 1998.

Pages 2 and 9, and mention stigmergic communication as a mechanism that probabilistic modelling may aid in understanding.

[79] Alcherio Martinoli, Masakazu Yamamoto, and Francesco Mondada. On the modelling of bio-inspired collective experiments with real robots. In E-Proc. Fourth European Conference on Artificial Life (ECAL-97), Brighton, UK, July 1997.

This paper uses text very similar to [77], but is missing one component of the definition. Pages 1-2, Beckers and collaborators [6] did the same experiment with robots...the collective behaviour was analysed on the basis of the stigmergy principle. Essentially, it consists in the production of certain behaviour in agents as a consequence of the effects produced in the environment by previous actions.

[80] Alcherio Martinoli, Auke J. Ijspeert, and Francesco Mondada. Understanding collective aggregation mechanisms: From probabilistic modelling to experiments with real robots. Robotics and Autonomous Systems, Special Issue on Distributed Autonomous Robotic Systems, 29:51-63, 1999.

On page 4, The only possible interactions among robots are the reciprocal avoidance of collisions and an indirect form of communication through modifications of the environment (stigmergic communication).It should be noted that the authors are referring to Khepera robots, whose only way of modifying the environment would be through the movement of seeds (the experiment is in clustering of seeds), which is actually work required for the task being performed.

[81] Alcherio Martinoli. Invited book review of Bonabeau et al. [11]. Artificial Life, 7(3):315-319, 2001.

On page 2, (spelling of Grassé s first name has been preserved as spelt in this document) The concept of stigmergy, a form of indirect communication among team mates through the environment, was introduced in 1959 by Pier-Paul Grassé [58]. He further says that: Fish societies, for instance, can consist of thousands of individuals ... that can be communicated in a stigmergic way by generating pressure waves.

[82] Zachary Mason. Programming with stigmergy: Using swarms for construction. In Artificial Life VI: Proc. of the Eighth International Conference on Artificial Life, pages 371-374, 2002.

On the first page: Termites and many social insects interact stigmergically that is, communication is mediated through changes in the environment rather than direct signal transmission. This paper also presents examples of the difference between quantitative and qualitative stigmergy, but, provides the termite pillar construction, and ant cemeteries as qualitative stigmergy, and wasp nest construction as quantitative. This is likely to just be a minor typographical error.

[83] Maja J Mataric. Designing and understanding adaptive group behaviour. Adaptive Behaviour, 4(1):51-80, December 1995.

Only uses the word in reference to Beckers et al. [6].

[84] Maja J Mataric. Issues and approaches in the design of collective autonomous agents. Robotics and Autonomous Systems, 16:321-331, December 1995.

On page 3: Direct communication is a purely communicative act, one with the sole purpose of transmitting information such as a speech act, or a transmission of a radio message. More specifically, directed communication is direct communication aimed at a particular receiver. Such communication can be one-to-one or one-to-many, in all cases to identified receivers. In contrast, indirect communication is based on the observed behaviour, not communication, of other agents. This type of communication is referred to as stigmergic in biological literature, where it refers to communication based on modifications of the environment rather than direct message passing.

[85] Maja J Mataric. Using communication to reduce locality in distributed multi-agent learning. Journal of Experimental and Theoretical Artificial Intelligence, special issue on Learning in Distributed AI Systems, 10(3):357-369, July-September 1998.

On page 3, very similar to that in Mataric [84], In contrast, indirect communication is based on the observed behaviour, not communication, of other agents, and its effects on the environment. This type of communication is referred to as stigmergic in biological literature, where it refers to communication based on modifications of the environment rather than direct message passing.

[86] Maja J Mataric. Great expectations: Scaling up learning by embracing biology and complexity. In NSF Workshop on Development and Learning, Michigan State University, Michigan, USA, April 2000.

On page 4, another important but overlooked form of communication is stigmergy, the ability to react to the effects of actions of others in the environment. She cites [63] (actually, due to an error, she cites something else, but this was the intended citation), and further mentions: Our work has explored stigmergy in the context of coordinated movement [138; 139]...

[87] Chris Melhuish, Owen E. Holland, and Steve Hoddell. Collective sorting and segregation in robots with minimal sensing. In Proc. Fifth International Conference on Simulation of Adaptive Behaviour (SAB1998), Zurich, Switzerland, August 1998.

This paper, discusses robot experiments; the authors refer the reader to Beckers et al. [6] for a description of the collective aspects. They do make the interesting distinction (made, they point out, in Holland [64]) between active and passive stigmergy; on page 10, Active stigmergy occurs when the effect of an environmental change due to one agent is to influence the choice of behaviour of the second agent... Passive stigmergy occurs when the environmental change affects only the outcome the behaviour of the second agent.

[88] Chris Melhuish, Owen E. Holland, and Steve Hoddell. Using chorusing for the formation of travelling groups of minimal agents. In Proc. Fifth International Conference on Intelligent

Autonomous Systems (IAS-5), pages 572-578, Sapporo, Japan, June 1998.

This paper does not use the word stigmergy, but is cited in Vaughan et al. [135] as doing so. The work involves forming what the authors term pseudoswarms through the use of a local signal, which influences the motions of others (and their chirping). Work is done in simulation, but real robots are intended to use ultrasound.

[89] Theodoros Michalareas and Lionel Sacks. Stigmergic techniques for solving multi-constraint routing for packet networks. In Proc. First International Conference on Networking, pages 687-697, Colmar, France, July 2001.

On page 2, The technique used by ant colonies to locate and transfer food supplies into their nest, or even construct complex structures has been termed stigmergy [58]...Also on page 2: The proper definition of stigmergy by Grassé is the following: Stimulation of workers by the performance they have achieved. So, basically stigmergy is a positive feedback mechanism using chemical substances like pheromones to attract agents, that themselves depose these chemicals. The paper presents a pheromone computational solution to a routing problem.

[90] Alberto Montresor. Anthill: a framework for the design and analysis of peer-to-peer systems. In Proc. Fourth European Research Seminar on Advances in Distributed Systems, Bertinoro, Italy, May 2001.

On page 2: This reflects the behaviour of real ants that cooperated through stigmergy, i.e. the capability to communicate among individuals through modifications induced in the environment.

[91] Scott Moss and Kerstin Dautenhahn. Hierarchical organization of robots: A social simulation study. In Proc. 12th European Simulation Multi-conference Simulation Past, Present and Future (ESM 1998), pages 400404, Manchester, UK, June 1998.

On page 1: ...using the principle of stigmergy which is defined as The production of a certain behaviour in agents as a consequence of the effects produced in the local environment by previous behaviour [6] and later on the same page: Behaviour based research on the principle of stigmergy is not using explicit representations of goals, the dynamics of group behaviour are emergent and self-organising.

[92] Divine T. Ndumu and Joseph M. H. Tah. Agents in computer-assisted collaborative design. AI in Structural Engineering, 1454:249270, 1998.

On page 3, stigmergic where implicit non-language communication proceeds through the mutual interactions of the agents in their shared environment.

[93] Ozalp Babaoglu, Hein Meling, and Alberto Montresor. Anthill: A framework for the development of agent-based peer-to-peer systems. In Proc. 22nd International Conference on Distributed Computing Systems (ICDCS 02), Vienna, Austria, July 2002.

On page 3: Ants do no communicate directly with each other; instead, they communicate indirectly by leaving information ... in the visited nests. This form of indirect communication, used also by real ants, is known as stigmergy [58].

[94] Claudia Pahl-Wosl and Eva Ebenhoh. Komplexe adpative systeme. Technical Report ISSN:1433-3805, Institute of Environmental Systems Research, University of Osnabruck, July 2003.

This document is in German. I present only an overview. Grassé [58] is presented as the origin of the word, and his recognition of coordination between the workers that comes from the construction process. On page 7, the separation between quantitative and qualitative stigmergy is discussed wherein a difference in degree verses the difference of types of pheromones are used as examples. Also mentions that Anderson [2] considers there to be a fuzzy boundary between self-organisation and qualitative stigmergy.

[95] H. Van Dyke Parunak, Sven A. Brueckner, and John Sauter. Erims approach to fine-grained agents. In Proc. NASA/JPL Workshop on Radical Agent Concepts (WRAC2001), Greenbelt, Maryland, USA, September 2001.

He claims that stigmergy is a pattern of interaction wherein the agents interact through the environment rather than considering it as passive. On page 2: Stigmergy [58], from the Greek words stigma sign and ergos work: the work performed by the agents in the environment in turn guides their later actions. Almost all the discussion on pages 2-3, are specific to the placement of actual items in the environment (e.g. beads, chemicals or markers). Their work uses synthetic pheromones.

[96] H. Van Dyke Parunak, Sven A. Brueckner, John Sauter, and Jedd Posdamer. Mechanisms and military applications for synthetic pheromones. In Proc. Workshop on Autonomy Oriented Computation (AOC), Fifth International Conference on Autonomous Agents (Agents 2001), Montreal, Canada, May 2001.

In the abstract, the authors claim that artificial pheromones can imitate stigmergetic dynamics of insects. Stigmergy is one of the document keywords; it is not used in the document itself. This work takes a very specific view, the experiments are in pheromone computation.

[97] H. Van Dyke Parunak, Robert Savit, Sven A. Brueckner, and John Sauter. A technical overview of the aorist project. Technical Report CS-95-187, Environmental Research Institute of Michigan (ERIM), April 2001.

On page 1, ...the environment state variables to which agents respond in their decision making are affected by the actions of the agents themselves, a process called stigmergy. This takes a broader view than Parunak et al. [95], Parunak et al. [97] or Parunak et al. [96]. Still on the first page, they have (as a caption for a figure) Stigmergy Agent behaviours both influence and are influenced by the external environment.

[98] H. Van Dyke Parunak, Sven A. Brueckner, Mitch Fleischer, and James Odell. Co-x: Defining what agents do together. In Proc. Workshop on Teamwork and Coalition Formation, First Joint Conference on Autonomous Agents and Multi-agent Systems (AAMAS 02), Bologna, Italy, August 2002.

On page 4, Indirect decentralized flows occur when peers make and sense changes to endogenous environmental variables. This class of Coordination is called stigmergy, [58], from the Greek words stigma sign and ergos work: the work performed by the agents in the environment in turn guides their later actions. Also: A particularly common form of stigmergy is resource Competition, which occurs when agents seek access to limited resources. For example, if one agent consumes part of a shared resource, other agents accessing that resource will observe its reduced availability, and may modify their behaviour accordingly. Even less directly, if one agent increases its use of resource A, thereby increasing its maintenance requirements, the loading on maintenance resource B may increase, thereby decreasing its availability to other agents who would like to access B directly. In the latter case, environmental processes contribute to the dynamics of the state variables involved. Interestingly, the footnote on page 4 states that ... the term is too well established in the research community to warrant suggesting an alternative. They further claim (on page 5) that the Minority Game is an excellent example of Stigmergy.(Here they are likely referring to a weapons platform example of the game presented in Parunak et al. [97].) Also, it is interesting to note that the authors call the information flow in stigmergic systems indirect because there is non-message interaction (Table 1, page 4).

[99] H. Van Dyke Parunak, Sven A. Brueckner, and John Sauter. Synthetic pheromone mechanisms for coordination of unmanned vehicles. In Proc. First Joint Conference on Autonomous Agents and Multi-agent Systems (AAMAS 02), Bologna, Italy, August 2002.

The word only appears in the keywords. It is a pheromone computing paper.

[100] H. Van Dyke Parunak. Making swarming happen. In Proc. Conference on Swarming and Network Enabled Command, Control, Communications, Computers, Intelligence, Surveillance and Reconnaissance (C4ISR), McLean, Virginia, USA, January 2003.

On page 4: Centralized command and control ... treats the centralized commander as the main locus of intelligence. Classical AI mechanisms seek to endow the individual entity with local intelligence, while stigmergic mechanisms generate system-level intelligent behaviour through the interactions among entities that individually may not exhibit high levels of intelligence. On page 6, Stigmergy is a term coined in the 1950s by the French biologist Grassé [58] 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 works stigma sign and ergon action, and captures the notion that an agent's actions leave signs in the environment, signs that it and other agent's sense and that determine their subsequent actions. Different varieties of

stigmergy can be distinguished. One distinction concerns whether the signs consist of special markers that the agents deposit in the environment (marker--based stigmergy) or whether agents base their actions on current state of the solution (sematectonic stigmergy). Another distinction focuses on whether the environmental signals are a single scalar quantity, analogous to a potential field (quantitative stigmergy) or whether they form a set of discrete options (qualitative stigmergy).The author claims that these two distinctions are orthogonal, they provide the following examples, of differing combinations: Marker-Based Quantitative stigmergy: gradient following in a single pheromone eld; Marker-Based Qualitative stigmergy: decisions based on combinations of pheromones; Sematectonic Quantitative stigmergy: ant cemetery clustering; Sematectonic Qualitative stigmergy: wasp nest construction.

[101] Patrick Peeters, Hendrik Van Brussel, Paul Valckenaers, Jo Wyns, Luc Bongaerts, Martin Kollingbaum, and Tapio Heikkila. Pheromone based emergent shop floor control system for flexible flow shops. Artificial Intelligence in Engineering, 15(4):343-352, October 2001.

Page 2: This term was introduced by Grassé in 1959 [58] to characterize the way social insects like ants interact. It describes a form of indirect and asynchronous interaction between individuals by using the environment as a means of information transfer. Indirect communication is taking place between individuals of an insect society by local modifications induced in their environment. According to the term Stigmergy, a sign (stigma) in the environment triggers an action (ergon, work) within the ant society.

[102] Andres Perez-Uribe and Beat Hirsbrunner. Learning and foraging in robot-bees. In International Society for Adaptive Behaviour (SAB2000) Proc. Supplement Book, pages 185194, Paris, France, September 2000.

Page 1: Insect societies exhibit division of labour and cooperate through direct and indirect (stigmergy) communication schemes [10].

[103] Leonid M. Peshkin, Nicolas Meuleau, and Leslie P. Kaelbling. Learning policies with external memory. In Proc. 16th International Conference on Machine Learning (ICML-99), pages 307 314, San Francisco, CA, 1999. Morgan Kaufmann.

In the abstract, they write: ...we explore a stigmergic approach, in which the agents actions include the ability to set and clear bits in an external memory, and the external memory is included as part of the input to the agent. On pages 1-2,The term is defined in the Oxford English Dictionary [116] as The process by which the results of an insects activity act as a stimulus to further activity, and is used in the mobile robotics literature [6] to describe activity in which an agents changes to the world affect its future behaviour, usually in a useful way. Further (page 2), one form of stigmergy is the use of external memory devices. We are all familiar with practices such as making grocery lists, tying a string around a finger, or putting a book by the door at home so you will remember to take it to work. In each case, an agent needs to remember something about the past and

does so by modifying its external perceptions in such a way that a memoryless policy will perform well.

[104] Leonid M. Peshkin. Architectures for policy search, July 2000. Proposal for PhD thesis, Brown University.

This work has the same text (in Chapter 3, pages 22-23) as that in Peshkin et al. [103] on pages 1-2.

[105] Vitorino Ramos and Ajith Abraham. Swarms on continuous data. In Proc. Fifth Congress on Evolutionary Computation (CEC03), Canberra, Australia, December 2003. IEEE Press.

The word is used identically to in Ramos et al. [107].

[106] Vitorino Ramos and Juan J. Merelo. Self-organized stigmergic document maps: Environment as a mechanism for context learning. In Proc. First Spanish Conference on Evolutionary and Bio-Inspired Algorithms (AEB2002), pages 284-293, Merida, Spain, February 2002.

[107] Vitorino Ramos, Fernando Muge, and Pedro Pina. Self-organized data and image retrieval as a consequence of inter-dynamic synergistic relationships in artificial ant colonies. In Frontiers in Artificial Intelligence and Applications, Soft Computing Systems - Design, Management and Applications, Second International Conference on Hybrid Intelligent Systems (HIS-2002), volume 87, pages 500509, Santiago, Chile, December 2002.

On page 1, One well known example is provided by the emergence of self-organization in social insects, via direct (mandibular, antennation, chemical or visual contact, etc) or indirect interactions. The latter types are more subtle and defined by Grassé as stigmergy [58] to explain task coordination and regulation in the context of nest reconstruction in Macrotermes termites.... In other words, stigmergy could be defined as a typical case of environmental synergy. The authors discuss termite nest construction, and ant cemetery clustering. Recruitment, trail laying division of labor and finally collaborative prey-transportation is mentioned.

[108] Olivia Rossi-Doria, Michael Sampels, Mauro Birattari, Marco Chiarandini, Marco Dorigo, Luca M. Gambardella, Joshua Knowles, Max Manfrin, Monaldo Mastrolilli, Ben Paechter, Luis Paquete, and Thomas Stutzle. A comparison of the performance of dierent metaheuristics on the timetabling problem. In Proc. Fourth International Conference on the Practice and Theory of Automated Timetabling IV (PATAT02), pages 115-119, Gent, Belgium, August 2002.

On page 2, mentions in passing: ...stigmergic information in the form of a pheromone level...

[109] R. Andrew Russell. Laying and sensing odor markings as a strategy for assisting mobile robot navigation tasks. IEEE Robotics & Automation Magazine, 2(3):39, September 1995.

This paper is about using olfactory sensors in mobile robots, including potential applications. It does mention that insects use this sense, but does not make use of the word stigmergy.

[110] Kurt Schelfthout, Tim Coninx, Alexander Helleboogh, Tom Holvoet, Elke Steegmans, and Danny Weyns. Agent implementation patterns. In Proc. Workshop on Agent-Oriented Methodologies, 17th Annual ACM Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA02), pages 119130, Seattle, Washington, USA, November 2002.

This paper makes use of a stigmergy metaphor for discussing a particular form of communication. All uses of the word are closely related to pheromones, for example on page 6: persistence is the time the message hangs in the air. Typically, a spoken word will disappear quickly, while a pheromone (in the stigmergy metaphor) will last a certain while.

[111] Miguel Schneider-Fontan and M. Mataric. The role of critical mass in multi-robot adaptive task division. Technical Report CS-95-187, School of Computer Science, Brandeis University, October 1996.

They make use of the word by saying, that Beckers et al. [6] describe a stigmergic approach to multi-robot foraging.

[112] Miguel Schneider-Fontan and M. Mataric. A study of territoriality: The role of critical mass in adaptive task division. In Proc. Fourth International Conference on Simulation of Adaptive Behaviour (SAB1996), pages 553-561, North Falmouth, Massachusetts, USA, September 1996. MIT Press/Bradford Books.

The word is used only regarding the stigmergic approach employed in Beckers et al. [6], this is identical to Schneider-Fontan and Mataric [111].

[113] Miguel Schneider-Fontan and M. Mataric. Territorial multi-robot task division. IEEE Transactions of Robotics and Automation, 14(5):815-822, October 1998.

The usage is identical to [111].

[114] Ruud Schoonderwoerd, Owen E. Holland, Janet L. Bruten, and Leon J. Rothkrantz. Ant-based load balancing in telecommunications networks. Adaptive Behaviour, 5(2):169-207, 1996.

On page 5: In many cases, the principle of stigmergy [58] is used. Stigmergy is a form of indirect communication through the environment....If an ants action changes the local environment in a way that affects one of these specific stimuli, this will influence the subsequent actions of ants at the location. The environmental change may take either of two distinct forms. In the first, the physical characteristics may be changed as a result of carrying to some task-related action... this type of influence has been called sematectonic [149]. In the second form, the environment is changed by depositing something which makes *no contribution to the task, but is solely intended to influence subsequent behaviour which is task related. This [is] sign-based stigmergy...Later (still on this page), they write: Some of the above behaviours have been successfully simulated with computer models, using both sematectonic stigmergy [125], and sign-based stigmergy [119] and also on robots [6; 109; 29]. The authors then go on to use, in their terminology, sign-based stigmergy.*

[115] Michael Schreyer and Gunther R. Raidl. Letting ants labeling point features. In Proc. Fourth Congress on Evolutionary Computation (CEC02), pages 1564-1569, Honolulu, Hawaii, USA, May 2002. IEEE Press.

Page 2: [Ants] deposit pheromones on the ground that influence the behaviour of following ants. Via this indirect communication, called stigmergy, a cooperative ant colony is able to efficiently determine the shortest path between its nest and a food source [5].

[116] John A. Simpson and Edmund S. C. Weiner, editors. Oxford English Dictionary. Oxford: Clarendon Press, 2nd edition, 1989.

Stigmergy: The process by which the results of an insects' activity act as a stimulus to further activity.

[117] Krzysztof Socha, Michael Sampels, and Max Manfrin. Ant algorithms for the university course timetabling problem with regard to the state-of-the-art. In Proc. Third European Workshop on Evolutionary Computation in Combinatorial Optimization (EvoCOP 2003), pages 334-345, Essex, UK, April 2003.

This paper provides no working definition for stigmergy. The most useful way to infer their intended meaning appears on page 4: The stigmergic information is in the form a matrix of pheromone values.... the pheromone values are an estimate of the utility of making the assignment, as judged by previous iterations of the algorithm.

[118] Krzysztof Socha. The influence of run-time limits on choosing ant system parameters. In Proc. Genetic and Evolutionary Computation Conference (GECCO 2003), Chicago, Illinois, USA, July 2003.

Similarly to Socha et al. [117], no definition is provided for stigmergy, and there are few hints regarding their intended meaning. Again, the only hint is really the use of a pheromone field.

[119] Tim R. Stickland, Chris M.N Tofts, and Nigel R. Franks. A path choice algorithm for ants. Naturwissenshaften, 79:567-672, 1992.

[120] Peter Stone and Manuela M. Veloso. Multi-agent systems: A survey from a machine learning perspective. Technical Report CMU-CS-97-193, School of Computer Science, Carnegie Mellon University, December 1997.

This is identical to Stone and Veloso [121].

[121] Peter Stone and Manuela M. Veloso. Multiagent systems: A survey from a machine learning perspective. Autonomous Robots, 8(3):345-383, 2000.

This paper was formerly cited as Stone and Veloso [120]. On page 17, they write: Agents can also affect each other by one of two types of stigmergy [64]. First, active stigmergy occurs when an agent alters the environment so as to affect the sensory input of another agent. For example, a robotic agent might leave a marker behind it for other agents to observe... Second, passive stigmergy involves altering the environment so that the effects of another agents actions change. For example, if one agent turns off the main water valve to a building, the effect of another agent subsequently turning on the kitchen faucet is altered. Regarding the utilization of stigmergic techniques, they cite Goldman and Rosenschein [54] and Holland [64].

[122] Thomas Stutzle and Holder H. Hoos. MAX MIN ant system. Journal of Future Generation Computer Systems, 16(8):889-914, June 2000.

On page 2, The (artificial) pheromone trails are a kind of distributed numeric information (called stigmergic information in [38]) which is modified by the ants to reflect their experience accumulated with solving a particular problem.

[123] Tarja Susi and Tom Ziemke. Social cognition, artefacts, and stigmergy: A comparative analysis of theoretical frameworks for the understanding of artefact-mediated collaborative activity. Cognitive Systems Research, 2(4):273-290, December 2001.

This work considers the problem of what is indirect vs. what is direct, through the use of the word artefacts. The discussion of stigmergy on pages 5 and 6 are more useful than most. On page 1, In the case of social insect such emergent coordination had been explain by the theory of stigmergy, which describes how individuals can detect the behaviour of others (and their own) through artefacts, i.e. the produce of their own activity (e.g., building material in the ants case).On page 2, In the 50s Grassé [58] formulated the concept of stigmergy, which is a class of mechanisms that mediate animal-animal interactions [126].

[124] Peter Tarasewich and Patrick R. McMullen. Swarm intelligence: Power in numbers. Communications of ACM, 45(8):62-67, August 2002.

The article gives an overview of a number of things that are commonly considered swarm-intelligence. The only actual use of the word is on page 3, These interactions can be direct (via physical, visual or chemical contact) or indirect. Indirect contact can take the form of stigmergy, where one individual performs an action based on what was performed earlier by a different individual. An example of this can be seen from the construction of wasp nests, where certain configurations of existing cells trigger the creation of a new cell.

[125] Guy Theraulaz and Eric Bonabeau. Coordination in distributed building. Science, 269(4):686-688, 1995.

On page 1, in many cases, the structuration of the environment caused by the colony's activities structures in turn individual behaviours, in a process that Grassé coined stigmergy [58]. He showed that task coordination and regulation of the building activity in termites do not depend on interactions between workers themselves but are mainly achieved by the nest structure: Individual behaviours are controlled and guided by previous work. Every time a worker takes a building action, it modifies the shape of the local configuration that triggered its building action. The new configuration then automatically stimulates new actions from any worker in the colony.

[126] Guy Theraulaz and Eric Bonabeau. A brief history of stigmergy. Artificial Life, 5(2):97116, 1999.
[127] Guy Theraulaz, Simon Goss, Jacques Gervet, and Jean-Louis Deneubourg. Task differentiation in polistes wasp colonies: A model for self-organizing groups of robots. In Proc. First International Conference on Simulation of Adaptive Behaviour (SAB1990), pages 346355, Paris, France, September 1990.
[128] Guy Theraulaz, Eric Bonabeau, and Jean-Louis Deneubourg. The origin of nest complexity in social insects. Complexity, 3(6):1525, 1998.

Page 3: Indeed, a single action by an insect results in a small modification of the environment that influences the actions of other insects: This forms of indirect communication through the environment is an important aspect of collective coordination and has been coined stigmergy by Grassé [58]. Further: Grassé introduced stigmergy (from the Greek stigma: sting, and ergon: work) to explain task coordination and regulation in the context of nest reconstruction in termites of the genus Bellicositermes [59; 58]. Grassé showed that the coordination and regulation of building activities do not depend on the workers themselves by it mainly achieved by the nest structure: A stimulating configuration triggers the response of a termite worker, transforming the configuration into another configuration that may trigger in turn another (possibly different) action performed by the same termite or any other worker in the colony. On page 4, the authors show a termite constructing a nest, with some number of responses for a number of states. In the caption of figure 3(on page 4) they write: These new stimuli then act on the same termite or any other worker in the colony. Such a process, where the only relevant interactions taking place among the agents are indirect, through the environment that is modified by the other agents, is also called sematectonic communication. The authors state this reflects Grassé s [59; 58] notion of stigmergy. Further on page 4 the authors state that while stigmergy is an explanation for the coordination involved in the worker-worker interactions, it is insufficient for explaining the global phenomena that occur in the form of nests from only local rules. They mention that investigations seem to indicate that self-organisation and self-assembly are two possible mechanisms. Page 5 mentions that it is both stigmergy and self-organisation that result in the renowned piles resulting from termite nest construction. On pages 8-9, they turn to self-assembly, stating: Self-assembly, which we may also call qualitative stigmergy in the context of this articles, differs from SO in that individuals interact through and respond to, qualitative stimuli: When termites build pillars, they respond to *quantitative stimuli, namely: pheromone fields and gradients. Self-assembly is based on a discrete set of stimulus types: For example, an insect responds to a type-1 stimulus with action A and responds to a type-2 stimulus with action B. In other words, qualitatively different stimuli result in qualitatively different responses.*

[129] Paul Valckenaers, Hendrik Van Brussel, Martin Kollingbaum, and Olaf Bochmann. Multi--agent coordination and control using stigmergy applied to manufacturing control. In Multi-Agent Systems and Applications, ninth ECCAI Advanced Course (ACAI 2001) and Agent Links third European Agent Systems Summer School, (EASSS 2001), pages 317-334, Prague, Czech Republic, July 2001.

Page 2: P.P. Grassé introduced the word stigmergy in 1959 [58; 126]. Stigmergy means that agents put signs, called stigma in Greek, in their environment to mutually influence each others behaviour. Such mechanism is suitable for small-grained interactions compared to coordination methods that require an explicit rendezvous amongst the agents. With stigmergy, agents observe signs in their environment and act upon them without needing any synchronization with other agents... Stigmergy belongs to that category of indirect interactions. They provide the following example: This situation is analogous to one person buying a kilo of apples in a supermarket. The display of the apples and their price constitute the signs in the environment. The agent observes these signs and decides whether he will buy these apples without direct interaction with any other agent.

[130] Paul Valckenaers, Martin Kollingbaum, Hendrik Van Brussel, Olaf Bochmann, and Constantin Zamrescu. *The design of multi-agent coordination and control systems using stigmergy.* In Proceedings of the third International Workshop on Emergent Synthesis (IWES01), Bled, Slovenia, March 2001.

On page 1: The concept of stigmergy was coined by Grassé [58], characterising the type of interaction taking place in biological insect societies. In observing ant colonies, a form of coordination mechanism can be identified that is based on the creation and placement of smelling substances (or signs/stigmas, as Grassé points out) in the environment by members of these biological communities. On pages 1-2, stigmergy describes a form of asynchronous interaction and information exchange between agents mediated by an active environment. Later on page 2, ... [they] coordinate themselves by producing a dissipative field in their environment. The paper presents a pheromone computational approach for designing manufacturing system.

[131] Robert van Kommer and Fabrice Chantemargue. A speech recognition interface to khepera robots. In Proc. First International Khepera Workshop, pages 227238, Paderborn, Germany, December 1999.

Used in footnote on page 2: To our knowledge, Beckers et al. [6] were the rst to exploit a stigmergic coordination between robots. Stigmergic coordination means literally incitement to work by the products of the work.

[132] Griselda Navarro Varela and Mark C. Sinclair. Ant colony optimisation for virtual-wavelength-path routing and wavelength allocation. In Proc. First Congress on Evolutionary Computation (CEC99), Washington DC, USA, July 1999.

On the first page, Ants communicate indirectly through environmental stimuli; this form of communication is termed stigmergy [114]. The stimuli are based on two kinds of changes in the environment. With sematectonic stigmergy, the stimuli are task-related, actions such as digging a hole of building a ball of mud. These actions change the environment, and other ants react by performing the same or related actions: they remove more material from the hole, or add more mud to the ball. This cooperation with the task is not a consequence of intelligence, but simply a response to stimuli [114]. In sign-based stigmergy, the ants deposit a volatile hormone (pheromone) to act as a stimulus to other ants. The pheromone thus serves as a signally system, acting as a means of indirect communication.

[133] Richard T. Vaughan, Kasper Sty, Gaurav Sukhatme, and Maja J Mataric. Blazing a trail: insect-inspired resource transportation by a robot team. In Proc. Fifth International Symposium on Distributed Autonomous Robotic Systems (DARS-00), Knoxville, Tennessee, USA, October 2000.

Stigmergy only used as something that is advantageous (page 9) and in reference to the work in [87] (on page 2).

[134] Richard T. Vaughan, Kasper Sty, Gaurav Sukhatme, and Maja J Mataric. Whistling in the dark: Cooperative trail following in uncertain localization space. In Proc. Fourth International Conference on Autonomous Agents (Agents 2000), Barcelona, Spain, June 2000.

Page 1 has: ..stigmergy; the production of certain behaviour in agents as a consequence of the effects produced in the local environment by previous behaviour[128]).

[135] Richard T. Vaughan, Kasper Sty, Gaurav Sukhatme, and Maja J Mataric. Exploiting task regularities to transform between reference frames in robot teams. In Proc. IEEE International Conference on Robotics and Automation (ICRA 2002), Washington DC, USA, May 2002.

On page 1: ..as have other applications of stigmergic communication [88]. But Melhuish et al. [88] does not contain the word stigmergy.

[136] Richard T. Vaughan, Kasper Sty, Gaurav Sukhatme, and Maja J. Mataric. Lost: Localization-space trails for robot teams. IEEE Transactions on Robotics and Automation, Special Issue on Multi-Robot Systems, 18(5):796-812, October 2002.

Most instructive use of the word is on page 2: ... feedback loops mediated either by direct communication between agents, or by indirect interaction by repeated sensing and modification of the environment, known as stigmergy [58].

[137] Katja Verbeeck and Ann Nowe. Stigmergy and bootstrapping: Using ant algorithms as a case study for learning in MAS. In Proc. Second International Workshop on Ant Algorithms (ANTS 2000), Brussels, Belgium, September 2000.

On page 2: The ants' main medium of communication is through the building up of the path through an artificial chemical substance called pheromone. This method of indirect communication is referred to as stigmergy [37]... The ant, which travels the shortest path, reinforces the path with more amount of pheromone, which aids others to follow... This behaviour is known as auto catalytic behaviour of the positive feedback mechanism in which reinforcement of the previously most followed route, is more desirable for future search.

[138] Barry Werger and Maja J Mataric. Robotic foodchains: Externalization of state and program for minimal-agent foraging. In Proc. Fourth International Conference on Simulation of Adaptive Behaviour (SAB1996), pages 625-634, North Falmouth, Massachusetts, USA, September 1996.

On page 8: Some robotics research has presented or reproduced particular instances of stigmergy the production of a certain behaviour in agents as a consequence of the effects produced in the local environment by previous behaviour[6] (see also [28; 127]).

[139] Barry Werger and Maja J Mataric. Exploiting embodiment in multi-robot teams. Technical Report IRIS-99-378, Institute for Robotics and Intelligent Systems, University of Southern California, 1999.

On page 2, ... stigmergy the production of a certain behaviour in agents as a consequence of the effects produced in the local environment by previous behaviour [6] (see also [28; 127]).

[140] Barry Werger and Maja J Mataric. From insect to internet: Situated control for networked robot teams. Annals of Mathematics and Artificial Intelligence, 31(1):173-198, 2001.

On page 2, ... stigmergic [6; 63; 138] insect interaction (i.e. interaction through environmental effects)...Note that while in the authors' earlier work [138], they kept clear from calling it stigmergic (and mentioned stigmergy only with respect to the work of others) they seem to have since decided that it was in fact stigmergy.

[141] Barry Werger. Cooperation without deliberation: A minimal behaviour-based approach to multi-robot teams. Artificial Intelligence, 110:293-320, 1999.

[142] Tony White and Bernard Pagurek. Towards multi-swarm problem solving in networks. In Proc. Third International Conference on Multi-Agent Systems (ICMAS98), Paris, France, 1998. IEEE Press.

The use of stigmergy is identical to that in White et al. [145], (references occur on pages 1-2), some modifications are made for formatting (e.g. some words are italicized).

[143] Tony White and Bernard Pagurek. Artificial life, adaptive behaviour, agents application oriented routing with biologically-jnspired agents. In Proc. Genetic and Evolutionary Computation Conference (GECCO 1999), Orlando, Florida, USA, July 1999.

Stigmergy is used on page 2, and is identical to that in White et al. [145].

[144] Tony White and Bernard Pagurek. Emergent behaviour and mobile agents. In Proc. Workshop on Mobile Agents in Coordination and Cooperation at the Third International Conference on Autonomous Agents (Agents 1999), Seattle, Washington, USA, May 1999.

Again, the description is identical to that in White et al. [145] (all references are on page 2).

[145] Tony White, Bernard Pagurek, and Franz Oppacher. ASGA: Improving the ant system by integration with genetic algorithms. In Proc. Third Annual Conference on Genetic Programming, pages 610-617, 1998.

The text describing stigmergy is almost identical to that in White [147], all references to the word are on page 2 (spelling is changed to US, and punctuation changed in places.)

[146] Tony White, Bernard Pagurek, and Franz Oppacher. Connection management using adaptive agents: An application of mobile agents in network management. In Proc. International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA98), pages 802-809, July 1998.

The text mentioning stigmergy is practically identical to that in White et al. [145]; all references to the word are on page 2.

[147] Tony White. Swarm intelligence and problem solving in telecommunications. Canadian Artificial Intelligence Magazine, 41:14-16, Spring 1997.

On page 1: ...actions include modification of the environment in which the agent operates. Intelligent behaviour frequently arises through indirect communication between the agents; this being the principle of stigmergy [58]. It should be stressed, however, that the individual agents have no explicit problem solving knowledge...On pages 1-2: ...two forms of stigmergy have been observed. Sematectonic stigmergy involves a change in the physical characteristics of the environment. Nest building is an example of this... The second form of stigmergy is sign-based. Here something is deposited in the environment that makes no direct contribution to the task being undertaken but is used to influence the subsequent behaviour that is task related... [this form] is highly developed in ants. And on page 2: The collective behaviour which emerges is a form of autocatalytic behaviour where the more the ants follow the trail the more likely they are to do so. The process is characterized by a positive feedback loop...

[148] Tony White. Cemetery organization, sorting, graph partitioning and data analysis. Carleton

University, 95.590H Swarm Intelligence Lecture Notes,
http://www.scs.carleton.ca/arpwhite/courses/95590Y/notes/SI%20Lecture%2017.pdf, 2003.
[149] Edward Osborne Wilson. Sociobiology. Belknap Press of Harvard University Press, 1975.
[150] Weilin Zhong and David Evans. When ants attack: Security issues for stigmergic systems. Technical Report CS-2002-23, Department of Computer Science, University of Virginia, April 2002.

In the abstract: Stigmergic systems solve global problems by using indirect communication mediated by an environment. Because they are localized and dynamic, stigmergic systems are self-organizing, robust and adaptive. On page 1, stigmergic systems build these applications by using indirect communication mediated by a shared environment [58].

9.2 SWARM INTELLIGENCE

These papers refer to swarm intelligence.

[151] E. Bonabeau and G. Theraulaz, Swarm Smarts,
http://web.cs.ualberta.ca/~kube/papers/SciAmericanMarch2000SwarmSmarts.pdf

Stigmergy is mentioned throughout the paper, with numerous examples given that include routing, factory scheduling and cooperative transport.

[152] Ivers Peterson, Calculating Swarms,
http://www.sciencenews.org/articles/20001111/bob10.asp
[153] W. Knight, Military Robots to get Swarm Intelligence.
http://www.newscientist.com/article.ns?id=dn3661.
[154] S. Brueckner and H. Van Dyke Parunak, Resource-Aware Exploration of the Emergent Dynamics of Simulated Systems, Proceedings of the Second International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS03), Melbourne, Australia, July 2003, 781-788.
[155] S. Brueckner and H. Van Dyke Parunak, Information-Driven Phase Changes in Multi-Agent Coordination, Proceedings of the Second International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS03), Melbourne, Australia, July 2003.
[156] H. V. D. Parunak. Go to the Ant: Engineering Principles from Natural Agent Systems. Annals of Operations Research, 75:69-101, 1997.
[157] T. White. SynthECA: A synthetic ecology of chemical agents. Ph.D. thesis, Carleton University, Canada, August 2000.

9.3 ANT COLONY OPTIMIZATION

[158] Bonabeau E., Dorigo M., and Theraulaz G. Swarm Intelligence From Natural to Artificial Systems. Oxford University Press, New York NY, 1999.
[159] Dorigo M. and Stutzle T., Ant Colony Optimization, MIT Press, ISBN 0262042193, July 2004.
[160] Dorigo M., V. Maniezzo and A. Colorni. The Ant System: Optimization by a Colony of Cooperating Agents. IEEE Transactions on Systems, Man, and Cybernetics-Part B,

26(1):29-41.
[161] Dumitrescu A. and Mitchell J., Approximation Algorithms for Geometric Optimization Problems, in the Proceedings of the 9th Canadian Conference on Computational Geometry, Queen's University, Kingston, Canada, August 11-14, 1997, pp. 229-232.
[162] Pilat M. and White T., Using Genetic Algorithms to optimize ACS-TSP. In Proceedings of the 3rd International Workshop on Ant Algorithms, Brussels, Belgium, September 12-14 2002.
[163] Reinelt G. TSPLIB, A Traveling Salesman Problem Library. ORSA Journal on Computing, 3:376-384, 1991.
[164] Stützle T. and Dorigo M. ACO Algorithms for the Traveling Salesman Problem. In K. Miettinen, M. Makela, P. Neittaanmaki, J. Periaux, editors, Evolutionary Algorithms in Engineering and Computer Science, Wiley, 1999.
[165] White T., Kaegi S. and Oda T., Revisiting Elitism in Ant Colony Search. In Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2003), Chicago, July 2003.

9.4 COLLECTIVE ROBOTICS

[166] Bowyer, Adrian. Automated Construction using Co-operating Biomimetic Robots http://www.bath.ac.uk/~ensab/B-man/tr_11_00/ (2000)
[167] Werfel, Justin. Autonomous Multi-Agent Construction http://www.csail.mit.edu/research/abstracts/abstracts04/html/221/221.html (2004)
[168] Pollack, Jordan et al. Coevolutionary Robotics http://demo.cs.brandeis.edu/papers/nasa_eh99.pdf (1999)
[169] C.R. Kube, H. Zhang. Collective Robotics: From Social Insects to Robots http://web.cs.ualberta.ca/~kube/papers/ab93.pdf (1993)
[170] Mondada, F. et al. SWARM-BOT: A New Concept of Robust All-Terrain Mobile Robotic System http://www.swarm-bots.org/dllink.php?id=169&type=documents
[171] Erkin Bahceci, Onur Soysal and Erol Sahin.A Review: Pattern Formation and Adaptation in Multi-Robot Systems http://www.swarm-bots.org/dllink.php?id=347&type=documents
[172] Medicoff, Zach, The Way We Live, Backbone Magazine May/June 2004, pp. 16-18
[173] Parker, Zhang, Kube. Blind Bulldozing: Multiple Robot Nest Construction http://web.cs.ualberta.ca/~kube/papers/Blind_Bulldozing_IROS2003.pdf
[174] Eric Bonabeau and Guy Theraulaz. "Swarm Smarts," Scientific American, pages 82-90, March 2000. http://web.cs.ualberta.ca/~kube/papers/SciAmericanMarch2000SwarmSmarts.pdf
[175] 10. Groß R. and Dorigo M. Evolving a Cooperative Transport Behavior for Two Simple Robots, In Liardet P., Collet P., Fonlupt C., Lutton E., and Schoenauer M., editors, Artificial Evolution - 6th International Conference (EA 2003), volume 2936 of Lecture Notes in Computer Science, Springer Verlag, Berlin, Germany (2004) 305-317.
[176] C. Ronald Kube and Hong Zhang. Stagnation Recovery Behaviours for Collective Robotics, 1994 IEEE/RSJ/GI International Conference on Intelligent Robots and Systems, pages 1893-1890.
[177] Groß R. and Dorigo M. Cooperative Transport of Objects of Different Shapes and Sizes, In M. Dorigo, M. Birattari, C. Blum, L. M. Gambardella, F. Mondada, and T. Stützle, editors, Ant Colony Optimization and Swarm Intelligence, 4th International Workshop (ANTS 2004), volume 3172 of Lecture Notes in Computer Science, Springer Verlag, Berlin, Germany (2004) 107-118.

[178] Baldassarre G., Parisi D., Nolfi S. Coordination and behaviour integration in cooperating simulated robots. In S. Schaal, A. Ijspeert, A. Billard, S. Vijayakamur, J. Hallam, and J.-A. Meyer, editors, From Animals to Animats 8. Proceedings of the Eight International Conference on Simulation of Adaptive Behavior (SAB04), pages 385-394. MIT Press, Cambridge, MA, 2004.

[179] V. Trianni, E. Tuci, and M. Dorigo. Evolving functional self-assembling in a swarm of autonomous robots. In S. Schaal, A. Ijspeert, A. Billard, S. Vijayakamur, J. Hallam, and J.-A. Meyer, editors, From Animals to Animats 8. Proceedings of the Eight International Conference on Simulation of Adaptive Behavior (SAB04), pages 405-414. MIT Press, Cambridge, MA, 2004.

[180] A. Martinoli. Swarm Intelligence in Autonomous Collective Robotics: From Tools to the Analysis and Synthesis of Distributed Collective Strategies. Ph.D. Thesis Nr. 2069, October, 1999, DI-EPFL, Lausanne, Switzerland. Available online at: http://www.coro.caltech.edu/People/alcherio/thesis/AM_PhD.pdf

[181] Trianni V., Nolfi S., Dorigo M. Cooperative Hole Avoidance in a *Swarm-bot* To appear in *Robotics and Autonomous Systems*

9.5 MISCELLANEOUS

[182] Couzin, I.D., Krause, J., Franks, N.R. & Levin, S.A. Effective leadership and decision making in animal groups on the move. *Nature,* 433, 513-516, 2005.

[183] V. A. Cicirello and S. F. Smith. Wasp nests for self-configurable factories. In Proceedings of the Fifth International Conference on Autonomous Agents. ACM Press, May-June 2001.

[184] White T. and Helferty J., Emergent Team Formation: Applying Division of Labour Principles to Robot Soccer. In Lecture Notes in Computer Science Vol. 3464, 2005.

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

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

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

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

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

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

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

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

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

[194] R. A. Brooks. A Robust Layered Control System for a Mobile Robot. *IEEE Journal of Robotics and Automation,* RA-2(1 (March)):14-23, 1986.
[195] S. Brueckner. *Return from the Ant: Synthetic Ecosystems for Manufacturing Control.* Dr.rer.nat. Thesis at Humboldt University Berlin, Department of Computer Science, 2000. URL http://dochost.rz.hu-berlin.de/dissertationen/brueckner-sven-2000-06-21 /PDF/Brueckner.pdf.
[196] 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, 2003. URL http://www.erim.org/~vparunak/AAMAS03InfoPhaseChange.pdf.
[197] S. Brueckner and H. V. D. Parunak. Resource-Aware Exploration of Emergent Dynamics of Simulated Systems. In *Proceedings of Autonomous Agents and Multi-Agent Systems (AAMAS 2003),* Melbourne, Australia, 2003. URL http://www.erim.org/~vparunak/AAMAS03APSE.pdf.
[198] S. A. Brueckner and H. V. D. Parunak. Swarming Agents for Distributed Pattern Detection and Classification. In *Proceedings of Workshop on Ubiquitous Computing, AAMAS 2002,* Bologna, Italy, 2002. URL http://www. erim.org/~vparunak/PatternDetection01.pdf.
[199] B. Clough. Emergent Behavior (Swarming): Tool Kit for Building UAV Autonomy. In *Proceedings of Swarming: Network Enabled C4ISR,* Tysons Corner, VA, ASD C3I, 2003.
[200] C. Croom, Jr. C4ISR Warfighter Integration. In *Proceedings of Swarming: Network Enabled C4ISR,* Tysons Corner, VA, ASD C3I, 2003.
[201] R. Davis and R. G. Smith. Negotiation as a metaphor for distributed problem solving. *Artificial Intelligence,* 20:63-109, 1983.
[202] M. B. Dias and A. Stentz. A Free Market Architecture for Distributed Control of a Multi-robot System. In *Proceedings of The 6th International Conference on Intelligent Autonomous Systems (IAS),* Venice, Italy, 2000. URL http ://www. frc. ri. cmu. edu/projects/colony/Dias Stentz.p s.
[203] J. Dubik, R. Richards, and G. Trinkle. Joint Concept Development and Experimentation. In *Proceedings of Swarming: Network Enabled C4ISR,* Tysons Corner, VA, ASD C3I, 2003.
[204] S. J. A. Edwards. Swarming on the Battlefield: past, Present, and Future. MR-1100-OSD, RAND, Santa Monica, CA, 2000.
[205] S. J. A. Edwards. Military History of Swarming. In *Proceedings of Swarming: Network Enabled C4ISR,* Tysons Corner, VA, ASD C3I, 2003.
[206] J. T. Feddema, R. Robinett, D. Schoenwald, C. Lewis, J. Wagner, and E. Parker. Analysis and Control for Distributed Cooperative Systems. *Presentation at Swarming Entities – Joint C4ISR DSC Study Plan Conference,* Johns Hopkins University Applied Physics Laboratory, Laurel, MD, 2002.
[207] M. Fleischer. Foundations of Swarm Intelligence: From Principles to Practice. In *Proceedings of Swarming: Network Enabled C4ISR,* Tysons Corner, VA, ASD C3I, 2003.
[208] D. Frelinger, J. Kvitky, and W. Stanley. Proliferated Autonomous Weapons: An Example of Cooperative Behavior. DB-239-AF, RAND Corporation, 1998.
[209] M. Gerla. Swarm Communications in the ONR-Minuteman Project. In *Proceedings of Swarming: Network Enabled C4ISR,* Tysons Corner, VA, ASD C3I, 2003.
[210] S. Goss, S. Aron, J. L. Deneubourg, and J. M. Pasteels. Self-organized Shortcuts in the Argentine Ant. *Naturwissenschaften,* 76:579-581, 1989.
[211] S. Hackwood and G. Beni. Self-Organizing Sensors by Deterministic Annealing. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robot and Systems,* Los

Alamitos, CA, pages 1177-1183, IEEE Computer Society Press, 1991, 1991.
[212] S. Hackwood and G. Beni. Self-organization of Sensors for Swarm Intelligence. In *Proceedings of IEEE Int. Conf. on Robotics and Automation,* pages 819-29, 1992.
[213] E. Hornung and B. M. Bryan, Editors. *The Quest for Immortality: Treasures of Ancient Egypt.* Washington, DC, National Gallery of Art, 2002.
[214] D. Inbody. Swarming: Historical Observations and Conclusions. In *Proceedings of Swarming: Network Enabled C4ISR,* Tysons Corner, VA, ASD C3I, 2003.
[215] C. R. Kube and H. Zhang. Collective Robotics: From Social Insects to Robots. *Adaptive Behavior,* 2(2):189-219, 1993. URL ftp://ftp.cs.ualberta.ca/pub/kube/ab93.pdf.
[216] M. K. Lauren. Describing Rates of Interaction between Multiple Autonomous Entities: An Example Using Combat Modelling. 2001. PDF File, http: //aps. arxiv. org/pdf/nlin/0109024.
[217] L. Leonardi, M. Mamei, and F. Zambonelli. Co-Fields: Towards a Unifying Model for Swarm Intelligence. DISMI-UNIMO-3-2002, University of Modena and Reggio Emilia, Modena, Italy, 2002. URL http://polaris.ing.unimo.it/didattica/curriculum/marco/Web-Co Fields/stuff/Swarm.pdf.
[218] M. Mamei, F. Zambonelli, and L. Leonardi. A Physically Grounded Approach to Coordinate Movements in a Team. In *Proceedings of First International Workshop on Mobile Teamwork (at ICDCS),* Vienna, Austria, IEEE CS Press, 2002. URL http://polaris.ing.unimo. it/Zambonelli/PDF/Teamwork.pdf.
[219] M. Mamei, F. Zambonelli, and L. Leonardi. Distributed Motion Coordination with Co-Fields: A Case Study in Urban Traffic Management. In *Proceedings of 6th IEEE Symposium on Autonomous Decentralized Systems (ISADS 2003),* Pisa, Italy, IEEE CS Press, 2003. URL http://polaris.ing.unimo.it/Zambonelli/PDF/isads.pdf.
[220] E. Neufeld. Insects as Warfare Agents in the Ancient Near East. *Orientalia,* 49(1):30-57, 1980.
[221] H. V. D. Parunak. 'Go to the Ant': Engineering Principles from Natural Agent Systems. *Annals of Operations Research,* 75:69-101, 1997. URL http://www.erim.org/~vparunak/gotoant.pdf.
[222] H. V. D. Parunak, S. Brueckner, R. Matthews, and J. Sauter. How to Calm Hyperactive Agents. In *Proceedings of Autonomous Agents and Multi -Agent Systems (AAMAS 2003),* Melbourne, Australia, 2003. URL http://www.erim.org/~vparunak/AAMAS03Ritalin.pdf.
[223] H. V. D. Parunak, S. Brueckner, J. Sauter, and R. Savit. Effort Profiles in Multi-Agent Resource Allocation. In *Proceedings of Autonomous Agents and Multi-Agent Systems (AAMAS02),* Bologna, Italy, pages 248-255, 2002. URL www. erim. org/~vparunak/AAMAS02Effort.pdf.
[224] H. V. D. Parunak, M. Purcell, and R. O'Connell. Digital Pheromones for Autonomous Coordination of Swarming UAV's. In *Proceedings of First AIAA Unmanned Aerospace Vehicles, Systems,Technologies, and Operations Conference,* Norfolk, VA, AIAA, 2002. URL www.erim.org/~vparunak/AIAA02.pdf.
[225] K. Pister. Smart Dust: Autonomous sensing and communication in a cubic millimeter. 2001. Web Page, http://robotics.eecs.berkeley.edu/~pister/SmartDust/.
[226] J. M. Riggs. The Objective Force. In *Proceedings of Swarming: Network Enabled C4ISR,* Tysons Corner, VA, ASD C3I, 2003.
[227] 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.
[228] J. A. Sauter, R. Matthews, H. V. D. Parunak, and S. Brueckner. Evolving Adaptive Pheromone Path Planning Mechanisms. In *Proceedings of Autonomous Agents and Multi-*

Agent Systems (AAMAS02), Bologna, Italy, pages 434-440, 2002. URL www. erim. org/~vparunak/AAMAS02Evolution.pdf.

[229] SMDC-BL-AS. Swarming Unmanned Aerial Vehicle (UAV) Limited Objective Experiment (LOE). U.S. Army Space and Missile Defense Battlelab, Studies and Analysis Division, Huntsville, AL, 2001.

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

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

[232] 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.

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

[234] H. V. D. Parunak, Making Swarming Happen, In *Proceedings of Swarming: Network Enabled C4ISR,* Tysons Corner, VA, ASD C3I, 2003