

Entropy and Self-Organization in Multi-Agent Systems

H. Van Dyke Parunak

ERIM

PO Box 134001

Ann Arbor, MI 48113-4001

+1 734 623 2509

vparunak@erim.org

Sven Brueckner

ERIM

PO Box 134001

Ann Arbor, MI 48113-4001

+1 734 623 2529

sbrueckner@erim.org

ABSTRACT

Emergent self-organization in multi-agent systems appears to contradict the second law of thermodynamics. This paradox has been explained in terms of a coupling between the macro level that hosts self-organization (and an apparent reduction in entropy), and the micro level (where random processes greatly increase entropy). Metaphorically, the micro level serves as an entropy “sink,” permitting overall system entropy to increase while sequestering this increase from the interactions where self-organization is desired. We make this metaphor precise by constructing a simple example of pheromone-based coordination, defining a way to measure the Shannon entropy at the macro (agent) and micro (pheromone) levels, and exhibiting an entropy-based view of the coordination.

Keywords

Self-Organization, Pheromones, Entropy.

1. INTRODUCTION

Researchers who construct multi-agent systems must cope with the world’s natural tendency to disorder. Many applications require a set of agents that are individually autonomous (in the sense that each agent determines its actions based on its own state and the state of the environment, without explicit external command), but corporately structured. We want individual local decisions to yield coherent global behavior.

Self-organization in natural systems (e.g., human culture, insect colonies) is an existence proof that individual autonomy is not incompatible with global order. However, widespread human experience warns us that building systems that exhibit both individual autonomy and global order is not trivial.

Not only agent researchers, but humans in general, seek to impose structure and organization on the world around us. It is a universal experience that the structure we desire can be achieved only through hard work, and that it tends to fall apart if not tended. This experience is sometimes summarized informally as “Murphy’s Law,” the observation that anything that can go wrong, will go wrong and at the worst possible moment. At the root of the ubiquity of disorganizing tendencies is the Second Law of Thermodynamics, that “energy spontaneously tends to flow

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Autonomous Agents '01, May 28-June 1, 2001, Montreal, Canada.

Copyright 2001 ACM 1-58113-000-0/00/0000...\$5.00.

only from being concentrated in one place to becoming diffused and spread out.” [9]

Adding energy to a system can overcome the Second Law’s “spontaneous tendency” and lead to increasing structure. However, the way in which energy is added is critical. Gasoline in the engines of construction equipment can construct a building out of raw steel and concrete, while the same gasoline in a bomb can reduce a building to a mass of raw steel and concrete.

Agents are not immune to Murphy. The natural tendency of a group of autonomous processes is to disorder, not to organization. Adding information to a collection of agents *can* lead to increased organization, but only if it is added in the right way. We will be successful in engineering agent-based systems just to the degree that we understand the interplay between disorder and order.

The fundamental claim of this paper is that the relation between self-organization in multi-agent systems and thermodynamic concepts such as the second law is not just a loose metaphor, but can provide quantitative, analytical guidelines for designing and operating agent systems. We explain the link between these concepts, and demonstrate by way of a simple example how they can be applied in measuring the behavior of multi-agent systems. Our inspiration is a model for self-organization proposed by Kugler and Turvey [7], which suggests that the key to reducing disorder in a multi-agent system is coupling that system to another in which disorder increases. Section 2 reviews this model and relates it to the problem of agent coordination. Section 3 describes a test scenario that we have devised, inspired by self-organization in pheromone systems, and outlines a method for measuring entropy in this scenario. Section 4 reports our experimental results. Section 5 summarizes our conclusions.

2. AN ENTROPY MODEL FOR SELF-ORGANIZATION

In the context of biomechanical systems, Kugler and Turvey [7] suggest that self-organization can be reconciled with second-law tendencies if a system includes multiple coupled levels of dynamic activity. Purposeful, self-organizing behavior occurs at the macro level. By itself, such behavior would be contrary to the second law. However, the system includes a micro level whose dynamics generate increasing disorder. Thus the system as a whole is increasingly disordered over time. Crucially, the behavior of elements at the macro level is coupled to the micro level dynamics. To understand this model, we begin with an example, then abstract out the underlying dynamics, and finally comment on the legitimacy of identifying processes at this level with principles from thermodynamics.

2.1 An Example: Pheromones

The parade example of such a system is the self-organization of an insect colony (such as the construction of minimal spanning tree networks among nests and food sources by ants, or the erection of multi-storied structures with regularly spaced pillars and floors by tropical termites), through pheromone-based coordination [1, 11]. Pheromones are scent markers that insects use in two ways. First, they deposit pheromones in the environment to record their state. For example, a foraging ant just emerging from the nest in search of food might deposit nest pheromone, while an ant that has found food and is carrying it will deposit food pheromone. ([15] documents use of multiple pheromones by insects.) Second, they orient their movements to the gradient of the pheromone field. In the example of foraging ants, those seeking food climb the gradient of the food pheromone, while those carrying food climb the gradient of the nest pheromone. The most realistic models of the ants' pheromone-climbing behavior incorporates a stochastic element in their movement. That is, they do not follow the gradient deterministically, but use its strength to weight a roulette wheel from which they determine their movement.

The environment in which pheromones are deposited plays a critical role in such a system. It is not passive, but active, and performs three information-processing functions with the pheromones.

1. It *aggregates* deposits of the same flavor of pheromone from different ants, thus providing a form of *data fusion* across multiple agents at different times, based on their traversal of a common location.
2. It *evaporates* pheromones over time, thus forgetting obsolete information. This dynamic is usefully viewed as a novel approach to *truth maintenance*. Conventional knowledge bases remember every assertion unless there is cause to retract it, and execute truth maintenance processes to detect and resolve the conflicts that result when inconsistent assertions coexist. Insect systems forget every assertion unless it is regularly reinforced.
3. Evaporation provides a third function, that of *disseminating* information from the location at which it was deposited to nearby locations. An ant does not have to stumble across the exact location at which pheromone was deposited in order to access the information it represents, but can sense the direction from its current location to the pheromone deposit in the form of the gradient of evaporated pheromone molecules.

2.2 The Model in Detail

In the Kugler-Turvey model, ants and their movements constitute the macro level of the system, while pheromone molecules constitute the micro level. The purposeful movement of ants, constructing minimal paths among their nests and food sources, achieve a reduction in disorder at the macro level, made possible because the agents at this level are coupled to the micro level, where the evaporation of pheromone molecules under Brownian motion results in an overwhelming growth in disorder. As a result, the disorder of the overall system increases, in keeping with the Second Law, in spite of the emergence of useful order at the macro level.

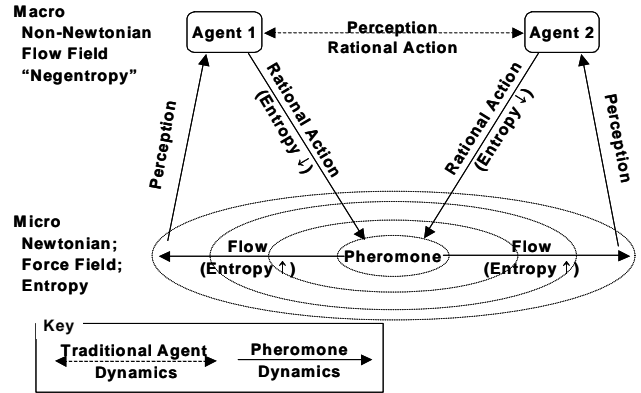


Figure 1. Comparison of Conventional and Pheromone-Based Models of Coordination

Figure 1 illustrates the interplay among these processes, and how this model of agent coordination differs from more classical views. Classically, agents are considered to perceive one another directly, reason about this perception, and then take rational action. The Kugler-Turvey model views coordination as being mediated by an environment that agents change by their actions (e.g., depositing pheromones), a process known as "stigmergy" [4]. Processes in the environment generate structures that the agents perceive, thus permitting ordered behavior at the agent level. At the same time, these processes increase disorder at the micro level, so that the system as a whole becomes less ordered over time.

Research in synthetic pheromones [2, 12, 13] draws directly on this model of coordination, but the model is of far broader applicability. In a multi-commodity market, individual agents follow economic fields generated by myriad individual transactions, and self-organization in the demand and supply of a particular commodity is supported by an environment that distributes resources based on the other transactions in the system. The movement of currency in such a system provides similar functions to those of pheromones in insect systems. More broadly, we hypothesize that a coupling of ordered and disordered systems is ubiquitous in robust self-organizing systems, and that the lack of such a coupling correlates with architectures that do not meet their designers' expectations for emergent cohesiveness.

2.3 A Caveat

At this point, readers with a background in physics and chemistry may be uneasy. These disciplines formulated the Second Law within a strict context of processes that result in energy changes. The fundamental physical measures associated with the second law are temperature T , heat Q , and (thermodynamic) entropy S , related by the definition

$$\text{Equation 1} \quad dS = \frac{dQ}{T}$$

Statistical mechanics identifies this macroscopic measure with the number Ω of microscopically defined states accessible to the system by the relation

$$\text{Equation 2} \quad S \equiv k \ln \Omega$$

where k is Boltzmann's constant, $1.4E-16$ erg/deg.

Thus defined, thermodynamic entropy has strong formal similarities [10] to information entropy [14]

Equation 3
$$S = -\sum_i p_i \log p_i$$

where i ranges over the possible states of the system and p_i is the probability of finding the system in state i . These formal similarities have led to a widespread appropriation of the notion of “entropy” as a measure of macro-level disorder, and of the Second Law as describing a tendency of systems to become more chaotic. Our approach participates to this appropriation.

It has been objected [8] that such appropriation completely ignores the role of energy intrinsic to both thermodynamic definitions (via T and dQ in the macro definition and k in the micro definition). Such an objection implicitly assumes that energy is logically prior to the definition, and that unless information processes are defined in terms of energy changes, it is illegitimate to identify their changes in entropy with those of thermodynamics. An alternative approach to the question would argue that in fact the prior concept is not ergs but bits, the universe is nothing but a very large cellular automaton with very small cells [3, 6], and physics and chemistry can in principle be redefined in terms of information-theoretic concepts. Our approach is sympathetic with this view. While we are not prepared at this point to define the precise correspondence between ergs and bits, we believe that physical models are an under-exploited resource for understanding computational systems in general and multi-agent systems in particular. The fact that the thermodynamic and information approaches work in different fundamental units (ergs vs. bits) is not a reason to separate them, but a pole star to guide research that may ultimately bring them together.

3. EXPERIMENTAL SETUP

We experiment with these concepts using a simple model of pheromone-based behavior. In this section we describe the experiment and how one measures entropy over it.

3.1 The Coordination Problem

Consider two agents, one fixed and one mobile, who desire to be together. Neither knows the location of the other. The mobile agent, or walker, could travel to the destination of the stationary one, if it only knew where to go. The stationary agent, or target, deposits pheromone molecules at its location. As the pheromone molecules diffuse through the environment, they create a gradient that the walker can follow.

Initially, the walker is at (30,30) and the target is at (50,50) in a 100x100 field. Every time step, the target deposits one molecule at (50,50). Both the walker and the pheromone molecules move by computing an angle $\theta \in [0, 2\pi]$ relative to their current heading and taking a step of constant length (1 for the walker, 2 for the pheromone molecule) in the resulting direction. Thus both molecules and walkers can be located at any real-valued coordinates in the field. Molecules move every cycle of the simulation and the walker every five cycles, so altogether the molecules move ten times as fast as the walker. Molecules fall off of the field when they reach the edge, while the walker bounces off the edges.

Molecules choose the heading for their next step from a uniform random distribution, and so execute an unbiased random walk. The walker computes its heading from two inputs.

1. It generates a gradient vector \vec{G} from its current location to each molecule within a specified radius ρ , with magnitude

Equation 4
$$|\vec{G}| = \sum_{r_i < \rho} \frac{g}{r_i^2}$$

where r_i is the distance between the walker and the i th molecule and g is a “gravitational constant” (currently 1).

2. It generates a random vector \vec{R} with random heading and length equal to a temperature parameter T .

The vector sum $\vec{G} + \vec{R}$, normalized to the walker’s step length (1 in these experiments), defines the walker’s next step. Including \vec{R} in the computation permits us to explore the effectiveness of different degrees of stochasticity in the walker’s movement, following the example of natural pheromone systems.

The state of the walker defines the macro state of the system, while the states of the molecules define the micro state. This model can easily be enhanced in a number of directions, including adding multiple walkers and multiple targets, and permitting walkers and targets to deposit pheromone molecules of various flavors. The simple configuration is sufficient to demonstrate our techniques and their potential for understanding how the walker finds the target.

3.2 Measuring Entropy

Computing the Shannon or Information Entropy defined in Equation 3 requires that we measure

1. the set of states accessible to the system and
2. the probability of finding the system in each of those states.

3.2.1 Measuring the Number of System States

In most computational systems, the discreteness of digital computation makes counting system states straightforward (though the number of possible states is extremely high). We have purposely defined the movement of our walker and molecules in continuous space to highlight the challenge of counting discrete system states in an application embedded in the physical world (such as a robotic application). Our approach is to superimpose a grid on the field, and define a state on the basis of the populations of the cells of the grid.

We can define state, and thus entropy, in terms either of location or direction. Location-based state is based on a single snapshot of the system, while direction-based state is based on how the system has changed between successive snapshots. Each approach has an associated gridding technique.

For location-based entropy, we divide the field with a grid. Figure 2 shows a 2x2 grid with four cells, one spanning each quarter of the field. The state of this system is a four-element vector reporting the number of molecules in each cell (in the example, reading row-wise from upper left, $\langle 1, 1, 3, 2 \rangle$). The number of possible states in an $n \times n$ grid with m particles is n^{2m} . The parameters in location-based gridding are the number of divisions in each direction, their orientation, and the origin of the grid.




Figure 2. Location-based gridding.

Rectangular grids are easiest to manage computationally, but one could also tile the plane with hexagons.

For direction-based entropy, we center a star on the previous location of each particle and record the sector of the star into which the particle is found at the current step. Figure 3 shows a four-rayed star with two particles. The state of the system is a vector with one element for each particle in some canonical order. Counting sectors clockwise from the upper left, the state of this example is $\langle 2,3 \rangle$. The number of possible states with an n -pointed star and m particles is mn . The parameters in direction-based gridding are the number of rays in the star and the rotation of the star about its center.

In both techniques, the analysis depends critically on the resolution of the grid (the parameter n) and its origin and orientation (for location) or rotation (for direction).

To understand the dependency on n , consider two extremes. If n is very large, the chance of two distributions of particles on the field having the same state is vanishingly small. For N distributions,




Figure 3. Direction-based gridding.

each will be a distinct state, each state will have equal probability $1/N$, and the entropy will be $\log(N)$. This state of affairs is clearly not informative. At the other extreme, $n = 1$, all distributions represent the same state, which therefore occurs with probability 1, yielding entropy 0, again not informative. We choose the gridding resolution empirically by observing the length scales active in the system as it operates.

To understand the dependency on origin/orientation or rotation, consider two particles in the same cell. After they move, will they still be in the same cell (keeping entropy the same) or in different cells (increasing entropy)? Exactly the same movements of the two particles could yield either result, depending on how the grid is registered with the field. We follow Gutowitz's technique [5] of measuring the entropy with several different origins and taking the minimum, thus minimizing entropy contributions resulting from the discrete nature of the grid.

3.2.2 Measuring the Probabilities

In principle, one could compute the probability of different system states analytically. This approach would be arduous even for our simple system, and completely impractical for a more complex system. We take a Monte Carlo approach instead. We run the system repeatedly. At each step in time, we estimate the probability of each observed state by counting the number of replications in which that state was observed. The results reported here are based on 30 replications.

Shannon entropy has a maximum value of $\log(N)$ for N different states, achieved when each state is equally probable. To eliminate this dependence on N , we normalize the entropies we report by dividing by $\log(N)$ (in our case, $\log(30)$), incidentally making the choice of base of logarithms irrelevant.

4. EXPERIMENTAL RESULTS

We report the behavior of entropy first in the micro system, then in the unguided and guided macro system, and finally in the complete system.

4.1 Entropy in the Micro System

Figure 4 shows locational entropy in the micro system (the pheromone molecules), computed from a 5x5 grid. Entropy increases with time until it saturates at 1. The more molecules enter the system and the more they disperse throughout the field, the higher the entropy grows. Increasing the grid resolution has no effect on the shape of this increase, but reduces the time to saturation, because the molecules must spread out from a single




Figure 4. Micro Entropy x Time (5x5 Grid)




Figure 5. Unguided Walker Path. Axes are location in the (100x100) field.

location and the finer the grid, the sooner they can generate a large number of different states.

Directional entropy also increases with time to saturation. This result (not plotted) can be derived analytically. The molecule population increases linearly with time until molecules start reaching the edge. Then the growth slows, and eventually reaches 0. Let M be the population of the field at equilibrium, and consider all M molecules being located at (50,50) through the entire run. Initially, all are stationary, and each time step one additional molecule is activated. Then the total number of possible system states for a 4-star is $4M$, but the number actually sampled during the period of linear population growth is $4t$, since the stationary molecules do not generate any additional states. Thus the entropy during the linear phase is $\log(4t)/\log(4M)$. As the growth becomes sublinear, the entropy asymptotically approaches 1, as with locational entropy.

4.2 Entropy in the Unguided Macro System

Figure 5 shows the path of a walker uncoupled to the micro system (when the target is emitting no pheromone molecules). With no coupling to the micro field, the walker is just a single molecule executing a random walk. Figure 6 shows that locational entropy for this walker increases over time, reflecting the increased number of cells accessible to the walker as its random walk takes it farther from its base. The grid size (15 divisions in each direction) is chosen on the basis of observations of the guided walker, discussed in the next section.

The directional entropy (not plotted) is constant at 1, since the walker chooses randomly at each step from all available




Figure 6. Unguided Walker Locational Entropy (15x15 Grid)




Figure 7. Guided Walker Path ($\rho = 20, T = 0$)

directions.

4.3 Entropy in Guided Macro System

Now we provide the walker with a micro field by emitting pheromone molecules from the target. Figure 7 shows the path followed by a typical walker with radius $\rho = 20$ and $T = 0$. This path has three distinct parts.

- Initially, the walker wanders randomly around its origin at (30,30), until the wavefront of molecules diffusing from (50,50) encounters its radius. In this region, the walker has no guidance, because no molecules are visible.
- Once the walker begins to sense molecules, it moves rather directly from the vicinity of (30,30) to (50,50), following the pheromone gradient.
- When it arrives at (50,50), it again receives no guidance from the molecules, because they are distributed equally in all directions. So it again meanders.

The clouds of wandering near the start and finish have diameters in the range of 5 to 10, suggesting a natural grid between 20x20 and 10x10. We report here experiments with a 15x15 grid.

Because of their initial random walk around their origin, walkers in different runs will be at different locations when they start to move, and will follow slightly different paths to the target (Figure 8).

The dots in Figure 9 and Figure 10 show the directional and locational entropies across this ensemble of guided walkers as a function of time. The solid line in each case plots the normalized median distance from the walkers to the target (actual maximum




Figure 8. Ensemble of Guided Walkers ($\rho = 20, T = 0$)




Figure 9. Guided walker: dots = directional entropy (4 star), solid line = median distance to target (max 28), dashed line = median visible molecules (max 151).

28), while the dashed line plots the normalized median number of molecules visible to the walkers (actual maximum 151). The lines show how changes in entropy and reduction in distance to the target are correlated with the number of molecules that the walker senses at any given moment.

At the beginning and end of the run, when the walkers are wandering without guidance, directional entropy is 1, corresponding to a random walk. During the middle portion of the run, when the walker is receiving useful guidance from the micro level, the entropy drops dramatically. As the temperature parameter T is increased in the range 50 to 100, the bottom of the entropy well rises, but the overall shape remains the same (plot not shown).

The locational entropy presents a different story. The minimization method for avoiding discreteness artifacts has the effect of selecting at each time step the offset that best centers the cells on the walkers. At the beginning of the run and again at the end, most walkers are close together, and fall within the same cell (because we chose a cell size comparable to these clouds). Walkers leave the starting cloud at different times, since those closer to the target sense the pheromones sooner, and follow different paths, depending on where they were when the pheromone reached them. Thus they spread out during this movement phase, and cluster together again once they reach the target. The effect of raising T to 100 on locational entropy is that the right end of the curve rises until the curve assumes a similar




Figure 10. Guided walker: dots = locational entropy (15x15 grid), solid line = median distance to target (max 28), dashed line = median visible molecules (max 151).

shape (plot not shown) to Figure 6.

Comparison of Figure 6 and Figure 10 shows that though the directed portion of the walker's movement has higher entropy than the undirected portions, coupling the walker to the micro level does reduce the walker's overall entropy. Even at its maximum, the entropy of the guided walker is much lower than that of the random one, demonstrating the basic dynamics of the Kugler-Turvey model.

The different behavior of locational and directional entropy is instructive. Which is more orderly: a randomly moving walker, or one guided by pheromones? The expected location of a random walker is stationary (though with a non-zero variance), while that of a guided walker is non-stationary. In terms of location, the random walker is thus more regular, and the location entropy reflects this. However, the movement of the guided walker is more orderly than that of the random walker, and this difference is reflected in the directional entropy. This difference highlights the importance of paying attention to dynamical aspects of agent behavior. Our intuition that the guided walker is more orderly than the random one is really an intuition about the movement of this walker, not its location.

4.4 Entropy in the Overall System

Central to the Kugler-Turvey model is the assertion that entropy increase at the micro level is sufficient to ensure entropy increase in the overall system even in the presence of self-organization and concomitant entropy reduction at the micro level. Our experiment illustrates this dynamic. As illustrated in Figure 4, by time 60, normalized entropy in the micro system has reached the maximum level of 1, indicating that each of the 30 replications of the experiment results in a distinct state. If each replication is already distinct on the basis of the locations of the pheromone molecules alone, adding additional state elements (such as the location of the walker) cannot cause two replications to become the same. Thus by time 60 the normalized entropy of the entire system must also be at a maximum. In particular, decreases in macro entropy, such as the decrease in locational entropy from time 80 on seen in Figure 10, do not reduce the entropy of the overall system.

One may ask whether the reduction in macro (walker) entropy is causally related to the increase in micro entropy, or just coincidental. After all, a static gradient of pheromone molecules would guide the walker to the target just as effectively, but would be identical in every run, and so exhibit zero entropy. This argument neglects whatever process generates the static gradient in the first place. An intelligent observer could produce the gradient, but then the behavior of the system would hardly be "self-organizing." In our scenario, the gradient emerges as a natural consequence of a completely random process, the random walk of the pheromone molecules emerging from the target. The gradient can then reduce the entropy of a walker at the macro level, but the price paid for this entropy reduction is the increase in entropy generated by the random process that produces and maintains the gradient.

One may also ask whether our hypothesis requires a quantitative relation between entropy loss at the macro level and entropy gain at the micro level. A strict entropy balance is not required; the micro level might generate more entropy than the macro level loses. In operational terms, the system may have a greater capacity for coordination than a particular instantiation exploits. What is

required is that the entropy increase at the micro level be sufficient to cover the decrease at the macro level, and this we have shown.

5. SUMMARY

To be effective, multi-agent systems must yield coordinated behavior from individually autonomous actions. Concepts from thermodynamics (in particular, the Second Law and entropy) have been invoked metaphorically to explain the conditions under which coordination can emerge. Our work makes this metaphor more concrete and derives several important insights from it.

- This metaphor can be made quantitative, through simple state partitioning methods and Monte Carlo simulation.
- These methods show how coordination can arise through coupling the macro level (in which we desire agent self-organization with a concomitant decrease in entropy) to an entropy-increasing process at a micro level (e.g., pheromone evaporation). Our demonstration focuses on synthetic pheromones for the sake of expositional simplicity, but we believe that the same approach would be fruitful for understanding self-organization with other mechanisms of agent coordination, such as market systems.
- This confirmation of the Kugler-Turvey model encourages us as agent designers to think explicitly in terms of macro and micro levels, with agent behaviors at the macro level coupled in both directions (causally and perceptually) to entropy-increasing processes at the micro level.
- Some form of pheromone or currency is a convenient mechanism for creating such an entropy-increasing process.
- Researchers must distinguish between static and dynamic order in a multi-agent system. We have exhibited a system that appears intuitively to be self-organizing, and shown that the measure of order underlying this intuition is dynamic rather than static.

6. ACKNOWLEDGMENTS

This work is supported in part by the DARPA JFACC program under contract F30602-99-C-0202 to ERIM CEC. The views and conclusions in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the US Government.

7. REFERENCES

- [1] E. Bonabeau, M. Dorigo, and G. Theraulaz. *Swarm Intelligence: From Natural to Artificial Systems*. New York, Oxford University Press, 1999.
- [2] S. Brueckner. *Return from the Ant: Synthetic Ecosystems for Manufacturing Control*. Thesis at Humboldt University Berlin, Department of Computer Science, 2000.
- [3] E. Fredkin. Finite Nature. In *Proceedings of The XXVIIth Recontre de Moriond*, 1992.
- [4] P.-P. Grassé. La Reconstruction du nid et les Coordinations Inter-Individuelles chez *Bellicositermes Natalensis* et *Cubitermes* sp. La théorie de la Stigmergie: Essai d'interprétation du Comportement des Termites Constructeurs. *Insectes Sociaux*, 6:41-80, 1959.
- [5] H. A. Gutowitz. Complexity-Seeking Ants. In *Proceedings of Third European Conference on Artificial Life*, 1993.
- [6] B. Hayes. Computational Creationism. *American Scientist*, 87(5):392-396, 1999.
- [7] P. N. Kugler and M. T. Turvey. *Information, Natural Law, and the Self-Assembly of Rhythmic Movement*. Lawrence Erlbaum, 1987.
- [8] F. L. Lambert. Shuffled Cards, Messy Desks, and Disorderly Dorm Rooms - Examples of Entropy Increase? Nonsense! *Journal of Chemical Education*, 76:1385, 1999.
- [9] F. L. Lambert. The Second Law of Thermodynamics. 2000. Web Page, <http://www.secondlaw.com/>.
- [10] J. Lukkarinen. Re: continuing on Entropy. 2000. Email Archive, <http://necsi.org:8100/Lists/complex-science/Message/2236.html>.
- [11] H. V. D. Parunak. 'Go to the Ant': Engineering Principles from Natural Agent Systems. *Annals of Operations Research*, 75:69-101, 1997.
- [12] H. V. D. Parunak and S. Brueckner. Ant-Like Missionaries and Cannibals: Synthetic Pheromones for Distributed Motion Control. In *Proceedings of Fourth International Conference on Autonomous Agents (Agents 2000)*, pages 467-474, 2000.
- [13] P. Peeters, P. Valckenaers, J. Wyns, and S. Brueckner. Manufacturing Control Algorithm and Architecture. In *Proceedings of Second International Workshop on Intelligent Manufacturing Systems*, pages 877-888, K.U. Leuven, 1999.
- [14] C. E. Shannon and W. Weaver. *The Mathematical Theory of Communication*. Urbana, IL, University of Illinois, 1949.
- [15] Smithsonian Institution. Encyclopedia Smithsonian: Pheromones in Insects. 1999. Web Page, <http://www.si.edu/resource/faq/nmnh/buginfo/pheromones.htm>.