Measures of Work in Artificial Life

Manoj Gambhir, Stephen Guerin, Daniel Kunkle and Richard Harris

RedfishGroup, 624 Agua Fria Street, Santa Fe, NM 87501
{manoj, stephen, daniel, rich}@redfish.com

Abstract

Definitions of autonomous agents—simple living systems—have often included the thermodynamic idea that agents need to do work to construct constraints and to maintain these constraints. We show that an ant-pheromone system is sufficiently rich to allow an explanation of its behavior in terms of thermodynamic properties, particularly the work done by the system on itself to build structure. This work is done by an inward pressure whose origin is the ant-pheromone system itself. The degree of structure formed by the system can be traced back to the behavioral rules of each of the ants—the more relevant information each ant is given in its behavioral repertoire, the more work is done by the system in compressing itself, and the tighter and more constrained is the final structure. The language introduced in this paper can serve as a framework for quantifying the propensity of a system to do work, self-organize and coalesce into life.

Introduction

Several studies of self-organizing systems have focused on the progress of these systems from an initial state of randomness, with no detectable structure, to one of structure and organization. There are many examples of systems that show this kind of behavior, and the innate tendency toward structure-formation is interesting but still little understood from a microscopic point of view. Pattern formation has also been widely studied, but, again, the viewpoint has been largely phenomenological and macroscopic rather than microscopic.

In addition to the emergence of structure in self-organizing systems, the maintenance of structure, and the apparent function served by it, is beginning to be studied. Importantly, some fairly recent definitions of the most basic kinds of living systems have focused on the thermodynamic activities they perform (Kauffman, 2000, 2003; Kugler and Turvey, 1987)

This paper extends the work of Guerin and Kunkle (2004) in which a very simple agent-based model of ant foraging was constructed and the processes leading to structure formation, structure maintenance, and structure decay were studied. It was found that these three features—commonly observed in complex system agent-based models— could be explained in terms of ideas from equilibrium and non-equilibrium thermodynamics. For example, when a system forms an organization, it appears to move from a state of high disorder, or, in thermodynamic terms, high entropy, to a state of low entropy. The second law of thermodynamics contradicts such a change, and, in the simple ant system, it was observed that an initial increase in entropy might account for the eventual drop in entropy. Such increases in entropy, which enable the formation of organization, are the mainstay of non-equilibrium thermodynamics (Atkins, 1984; Prigogine, 1962, 1984; Haken, 2000; and Schneider and Kay, 1995; Swenson and Turvey, 1992)

At the structure maintenance stage, the system has set up a positive feedback loop in which the ants, because they are very localized, deposit pheromones which are also localized, and which, in turn, cause the ants to localize further. The system requires an influx of food to maintain it, or a decay of the structure is observed. The purpose of this paper is to extend the thermodynamic analysis of the system—chiefly through the introduction of the concept of work—and to tie these ideas into a living systems framework along the lines of Kauffman (2000).

As explained previously, the ant system initially increases its disorder and then begins to organize and form a structure, decreasing its disorder in the process. We imagine a pressure acting inward on the system, causing ants to stop diffusing outward and come back to the nest and food-source, forming coherent pheromone tracks. Of course, this pressure is nothing as crude as that provided by a piston, compressing a gas in a chamber during the compression leg of an engine cycle—the ant-pheromone system is not, after all, simply squashed by our imaginary pressure; the process is closer to a condensation. However, as a first approximation, the idea of a piston compression is useful in that we are permitted to talk about the work done by the piston during its motion. By measuring the change in the motion of the ants away from their origin we show that something like an inward pressure is acting on the system. In thermodynamics, this pressure is multiplied—in infinitesimal steps—by the induced volume change to give the work done.

The source of this inward pressure has to reside in the ant-pheromone system itself, given that there are no other external components to our model. It is, therefore, in the behavioral rules of the ants and pheromones, along with any relevant initial conditions, that we should find the propensity of the system to exert this pressure on itself, and thereby do work on itself. We alter the rules of the ants to test this idea. Two rules that were added to the behavioral repertoire of the ants in the model of Guerin and Kunkle (2004) are successively removed and the results on the character of the final structures formed, are observed. We find that, for the rule changes imposed, the pressure—and, therefore, the amount of work done by the system—is strongly affected. Effectively, the more informed the agents are about the task they are supposed to perform (as dictated by us, the model builders, and as reflected in the rules each agent is given) the more work the system is capable of performing in constructing the ant-pheromone structure, and the more efficiently the system can perform the task of bringing food back to the nest.

The formation of structure can also be thought of as the construction of constraints, and the system can be seen to be doing work to construct constraints. In the language of Kauffman, these constraints then enable the system to perform a task, to do at least one thermodynamic work cycle and to construct further constraints. Here is a definition of an autonomous agent or living system. Our model, in its simplicity, seems to permit an analysis very much along these lines and may provide a framework to further elucidate the ways in which systems of distributed, simple agents coalesce into life.

Experimental Setup

The following simulation of food gathering ants is presented for the purpose of calculating statistical and thermodynamic measures that help characterize phases of self-organization. The spirit of this model is an extension of the work of Parunak and Brueckner (2001) and Gutowitz (1993).

The ant system described here is discrete; the positions of all objects in the system are specified by a 2-tuple of integers (x, y). The space of positions is a square grid. The three types of objects are nests, food, and ants. Ants and the environment are modeled as active agents in the simulation. Additionally, each position in the space can contain some amount of nest pheromone and food pheromone, which are deposited by the ants as they move.

A basic simulation is set up as follows: A nest and some amount of food are placed in the space. A fixed number of ants is initially placed at the nest. All positions have zero levels of both food and nest pheromones. The system evolves as the ants move, drop pheromone, and transport food. The model is flexible to later allow experimentation of initially placing ants at any location in the space and to allow any type of object (nests, food, ants) to be introduced at any time or position.

An ant can hold one unit of food at a time and can take one of three actions: 1) move to one of eight adjacent locations (includes diagonal moves), 2) pick up a unit of food, and 3) drop a unit of food at a nest. The following pseudo-code describes what actions an ant will take on each time step:

```
if ant has food then
  drop one unit of food pheromone
  if at nest then
    drop food
  else
    follow nest pheromones
  end if
else
  drop one unit of nest pheromone
  if at food then
    pick up food
  else
    follow food pheromones
  end if
end if
```

Also, each time step some percentage of the pheromone present at each position “evaporates”, or is removed. Pheromone evaporation allows adaptation to changes in food location. For example, if there were two food sources present, *A* and *B*, and the ants were exploiting *A* for a period of time a strong trail of pheromones would be laid between the nest and *A*. Once the food at *A* is gone the ants should no longer follow that trail, but should rather explore again to find *B*. If the pheromones leading to *A* do not evaporate this cannot occur. The decay or forgetting of constructed constraints (pheromone trails) allows the system to be adaptive.

The ants have directionality. They can only travel to their forward five positions instead of choosing from all eight adjacent positions. This local directionality is present regardless of the state of the system and is an example of a *context-free constraint* (Juarrero, 1999). The direction of an ant is calculated after each step based on the previous and current positions. At time zero each ant chooses a random direction.

In addition to this directionality, ants are given a behavior that turns them around when they reach their goal—either the food-source or the nest. In the main set of experiments performed for this paper, the turnaround behavior and the ant directionality are successively turned off to observe the result on the ant-pheromone structure formed when the system reaches a steady state. The turnaround and directionality behaviors were originally included to make the system more efficient at reaching a structured steady state, but these behaviors constitute information about the environment and it is central to this paper to observe the macroscopic effect of removing this information from the microscopic actors.

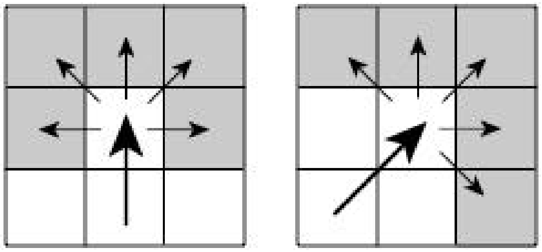

Figure 1. Examples of possible next steps (in gray) for an ant according to the last step taken. Ants have hard context-free constraints preventing backwards movement.

Each time step ants measure a local gradient then choose a direction to step. Pheromone levels are read from the forward five positions. The probability of moving to position *j* is given by:

$$p_j = \frac{m_j^a + b}{\sum_{n=1}^{N} m_n^a + b} \tag{1}$$

where μ_j is the pheromone level at position *j*; α is a scaling exponent; β is a random base; and the denominator represents the total pheromone level in all possible next positions. In this case *N* is five. The scaling exponent α increases the probability that the next position will be the one with the greatest pheromone level, whereas the random base β has the opposite effect. Typical values used in our experiments were $\alpha = 3$, $\beta = 1$, and a maximum pheromone level of 511 at any one position. These parameters can be adjusted to tune the likelihood that an ant will explore for new food versus exploiting a found food source.

So, the movement of an ant is constrained by a measure of change not an absolute strength. One can think of the gradient as the spatial first derivative of the pheromone field.

Figure 2 graphically depicts four phases of the typical evolution of this ant system. In the next section, measures of constraint and spatial entropy are defined as tools for examining the construction and destruction of constraints in this self-organizing system.

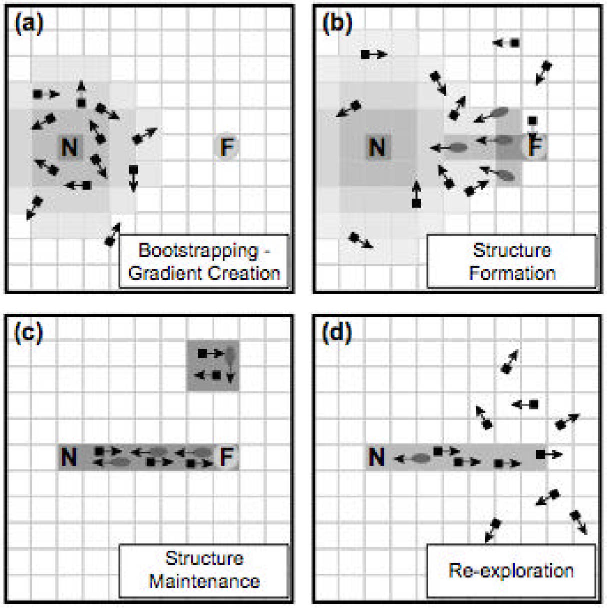

Figure 2. Typical evolution of the ant system. (a) *Bootstrapping – Gradient Creation:* Ants move randomly out from the nest, creating a gradient of nest pheromones. (b) *Structure Formation:* Some ants find the food and begin following the nest pheromones while dropping food pheromones that food-seeking ants begin to follow. (c) *Structure Maintenance:* A stable path of both food and nest pheromones is established. As shown in the upper-right corner, cycles that do not transport food can also form. (d) *Re-exploration:* Once all of the food has been transported to the nest the pheromones begin to evaporate and the ants disperse.

Constraint and Spatial Entropy Measures

As a system self-organizes, components of the system are expected to lose degrees of freedom through the emergence of context-sensitive constraints (Juarrero, 1999). In this system, ants lose directional degrees of freedom as they are informed by a gradient. We measure this constraint in our model with a directional entropy. An ant that sees no pheromone gradient, which is an equal level of pheromone in all possible next positions, is said to be maximally ignorant with an ignorance level of 1. An ant that has no choice but to move to one specific position on the next step would have an ignorance of 0, though this never occurs here because of the random base added to seach pheromone level as described in Eq. 1. The Shannon entropy (Shannon, 1948a, 1948b) of the probabilities of moving to each of the possible positions on the next step defines the ignorance for each ant

$$S = \frac{-\sum_{n=1}^{N} p_n \log p_n}{\log N} \tag{2}$$

where p_n is the probability of moving to position *n*, and *N* is the number of possible next positions, in this case five. The denominator is used to normalize the value to the range [0,1]. The ignorance of a group of ants is defined as the average ignorance of all ants.

As systems self-organize, statistical measures of order increase. In this model, we capture increased order with a spatial entropy measure applied to the positions of ants. The spatial entropy measure is also a Shannon entropy of the form shown in Eq. 2 where p_n is the proportion of all ants at position *n*, and *N* is the total number of positions in the space. Note that in the case of zero ants being at a location 0 log 0 = 0. The maximum spatial entropy is achieved with an equal number of ants at each position and the minimum with all ants at a single location.

Experimental Results

Summary of Model Dynamics

As a first step in presenting our results, we summarize the dynamics of the system for typical model runs. All results given in this summary are based on an average over 20 runs with the same initial conditions. The space was 21 positions square, with a single nest at position *(7,7)*, a single food source at position *(15,15)*, and all ants starting at the nest.

An indicator of which phase the system is in (i.e. an order parameter) could be the mean path length of the ants. The path length of an ant at a given time is defined as the number of steps it has taken since it last picked up or dropped a unit of food. Figure 3 shows the four phases of development with a plot of mean path length vs. time. In the bootstrap phase the mean path length increases uniformly. The line in this case has a slope of one because each ant takes one step per unit time and no ants have yet found food. The structure formation phase begins when food is found which leads to a rapid decrease in the mean path length. During the structure maintenance phase, mean path length remains stable at a near minimum value (the shortest distance between the nest and food). When the food source is depleted, the mean path length again increases as the ants re-explore the space for alternative food sources.

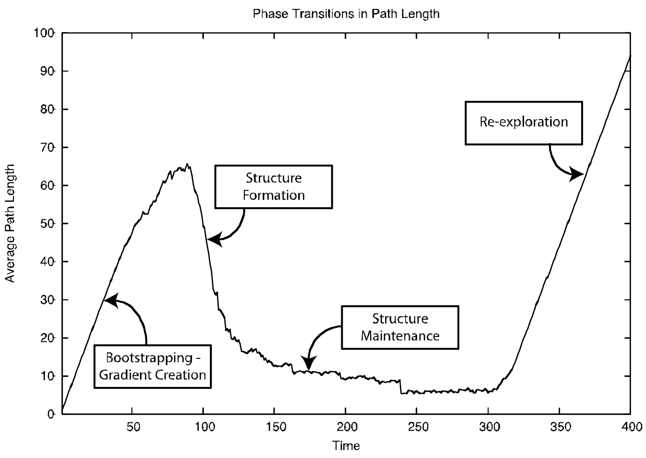

Fig. 3. The four phases of development are clearly visible in this plot of the ants' mean path length at each time step. A path length is calculated as the number of steps an ant has taken since it last picked up or dropped a unit of food.

Figure 4 displays the spatial entropy and ignorance for all ants over the first 100 time steps of an experiment with an inexhaustible food source. The bootstrapping phase occurs over approximately the first 25 time steps. In this phase, the ants' random walk from the nest causes a rapid increase in spatial entropy and the establishment of a nest pheromone gradient around the nest. In the structure formation phase ants that find food use this gradient to direct their return to the nest. The food-carrying ants' constrained movement is reflected in a reduction of overall ant ignorance as the gradient informs them to the nest location.

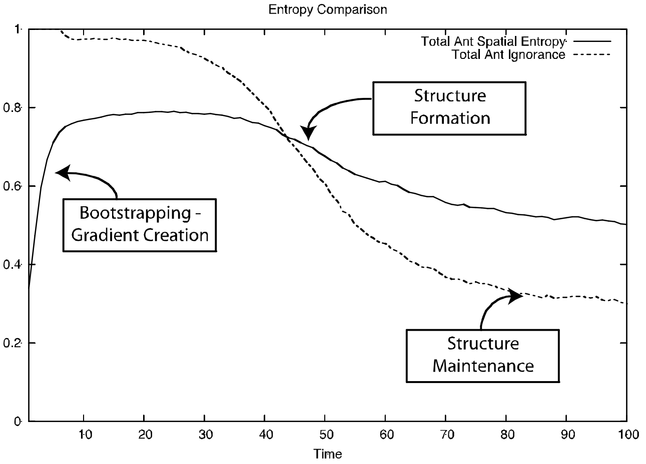

Figure 4. A comparison of the spatial entropy and ignorance of ants over time. In the *bootstrapping – gradient creation* phase the spatial entropy of the ants increases, which establishes a gradient of pheromones around the nest. In the *structure formation* phase the ignorance of the ants decreases as they find food and begin to follow the nest pheromones back. The spatial entropy of the ants also begins to decrease when a path is formed between the nest and food. The *structure maintenance* phase begins when the spatial entropy and ignorance of the ants becomes relatively constant.

System Pressure

During the bootstrapping phase, when ants are moving away from the nest in all directions, we can chart the tendency of the system to expand by looking at a specific measure. If we refer to the nest and food locations as goals—given that each ant is seeking either food or nest—we can, at each time-step, calculate the distance of each ant from its goal. This distance is obtained by a simple Euclidian calculation. In a time-step, this distance will change for each ant: we call an increase in the distance positive and a decrease negative. The sum of these distance-changes at each time-step constitutes our measure *P*, the system pressure. When the system has self-organized, almost every ant will decrease its distance from its goal at every time-step and, therefore, the measure will be large and negative. At the beginning of each model run, most ants will increase their distance from their goals at every time-step and so the measure will be positive. The behavior of *P* follows our intuition for the self-imposed pressure of the system: positive at first, and finally, after the structure has formed, large and negative.

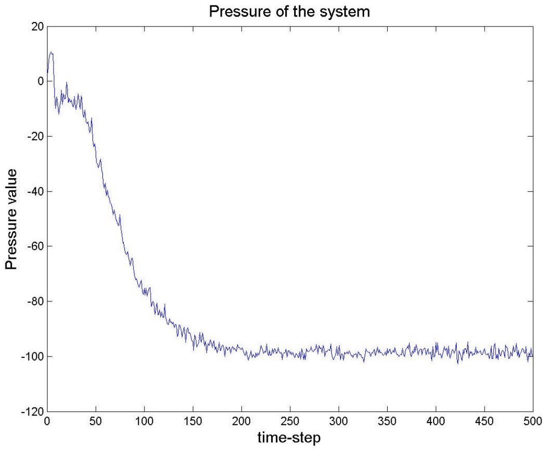

Figure 5 The evolution of the pressure in the system. Initially, the ants are mostly moving away from their goals, albeit randomly, giving a positive overall system pressure. When the structure has fully formed—around step 200 here—the pressure is strongly negative: the ants are nearly always moving toward their goals.

Total Work Done by the System

By changing the rules governing the ant behaviors we can investigate the influence of these rules on the final self-organized structure of the system and the work done in forming this structure. We look at three rule settings:

1) Basic ants
2) Basic ants with directional motion
3) Basic ants with directional motion and turnaround behavior

The 'basic ant' indicated above is simply an ant with no directional behavior i.e. no directional rule or turnaround behavior when the ant reaches food or nest. Statistics are compiled over 1000 runs. For each case, we calculate the number of food pieces picked up during an entire model run, and the average final directional entropy of the ants. Table 1 shows these values.

Experiment	**No. Food Pieces**	**Directional entropy**	**Max. Directional Entropy**
1	8	0.090	0.49
2	2074	0.055	0.19
3	3367	0.016	0.05

Table 1 Experiments performed on the ant systems in which the rules of the ants were changed. The measured quantities are the total food picked up by ants in a model run, the average directional entropy at the end of the runs, and the maximum directional entropy.

The trend observed is clear. As we add behavioral rules to the ants, the final value of the directional entropy decreases—the pattern formed by the ant-pheromone system when the model run has settled into a steady state becomes tighter and more confined with more behavioral rules. The ants themselves are more constrained to their paths. This trend is even clearer if we look at the fourth column in Table 1, which shows the maximum value of the average directional entropy of the ants in the 1000 models runs.

We can interpret the evolution of the population of ants—from a state in which they choose equally from the available directions of motion to one in which they are much more constrained—as a thermodynamic effect. Work has been done on the ants in reducing their directional entropies and the average entropy change, from a maximum initial value to the final value, is a measure of the work done.

The number of food pieces picked up by the ant population during the runs also increases as rules are added. The system performs a task—transferring food from the source to the nest—more efficiently with added rules. It should be noted that the ant rules are specific to this system—not just any accumulation of rules would have resulted in the observed trends. But, it can be said that the more information we give the agents, in the form of behavioral rules, the greater the potential the system has to do work to form structure.

Conclusions

The ant-pheromone system has been shown to impose on itself an inward pressure, which causes it to organize from a state of random ant motion and pheromone placement to confined ant-pheromone tracks. The motion of the ants in these tracks is very constrained and we can say that the inward system pressure has done work in constructing these constraints and confining the system.

When system specific rules are added to each ant, the structure formed becomes progressively more constrained and the system does more work to bring itself to its confined state. The extra rule information with which each agent is armed is expressed macroscopically by the greater overall work done by the system in constraining itself.

Our measure for system pressure is maintained at a high negative value as long as there is a sufficient supply of food and this may indicate that the system continues to do work to maintain its constraints. Kauffman's (2000) idea that an autonomous agent should perform work cycles is relevant here—our system, once assembled may be said to be performing work cycles to maintain itself. As a self-assembling, work cycle performing agent system, consuming food to maintain its structure, the ant-pheromone system might be said to be an archetypical living system minus the ability to reproduce itself.

Acknowledgement

The authors would like to express their gratitude for the suggestions offered from M. Salamon and members of the FRIAM Complexity Coffee Group in Santa Fe, New Mexico.

References

Atkins, P. W. (1984). *The second law*. New York: Scientific American Books.

Guerin, S. & Kunkle, D. (2004). Emergence of constraint in self-organizing systems. *Journal of Nonlinear Dynamics in Psychology and Life Sciences.* Kluwer Academic.

Gutowitz, H. A. (1993). Complexity-seeking ants. In Deneubourg & Goss (Eds.) Proceedings of Third European Conference on Artificial Life.

Haken, H. (2000). *Information and self-organization: A macroscopic approach to complex systems.* Berlin: Springer-Verlag

Juarrero, A. (1999). *Dynamics in action: Intentional behavior as a complex system.* Cambridge, MA: MIT Press

Kauffman, S. A. (2000). *Investigations*. Oxford University Press.

Kauffman, S. A. (2003). Autonomous agents. [online web video lecture] *FRIAMGroup Applied Complexity Lecture Series.* Retrieved April 10, 2003, from http://www.redfish.com/friam/. Santa Fe Institute.

Kugler, P. N. & Turvey, M. T. (1987). Information, natural law, and the self-assembly of rhythmic movement. Hillsdale, NJ: Lawrence Erlbaum.

Parunak, H. V. D. & Brueckner, S. (2001). Entropy and self-organization in multi-agent systems. *Proceedings of the Fifth International Conference on Autonomous Agents*. (pp. 124-130). Montreal: ACM Press.

Prigogine, I. (1984). Order out of chaos: Man’s new dialouge with nature. New York: Bantam.

Prigogine, I. (1962). *Non-equilibrium statistical mechanics*. Wiley Interscience, New York.

Schneider, E. D. & Kay, J. J. (1995). Order from disorder: the thermodynamics of complexity in biology. In M. P. Murphy & L. A. J. O’Neill (Eds.). *What is life: the next fifty years. reflections on the future of biology* (pp. 161-172). Cambridge: Cambridge University Press.

Shannon, C. E. (1948a). A mathematical theory of communication, Part 1. *Bell System Technical Journal*, 27, 379-423

Shannon, C. E. (1948b). A mathematical theory of communication, Part 2. *Bell System Technical Journal*, 27:623-656.

Swenson, R. & Turvey, M. T. (1991). Thermodynamic reasons for perception-action cycles. *Ecological Psychology*, 3:317-348.