

How adaptable are Swarms?

Project for Adaptive Systems

Mikel Maron
Evolutionary and Adaptive Systems
University of Sussex

April 18, 2004

Abstract

The evolutionary dynamics of collective behavior are inherently different from individual behavior. This issue forms part of key questions in theoretical biology and has consequences for biologically inspired computation. Here, the classic simulation to evolve central place foraging is implemented using finite state machines, in individual and collective scenarios. These evolutionary trajectories are then analyzed in depth, with implications for the main theme. Both optimal and unexpected strategies do develop. Deviating research directions emerge from the results, to the nature of representation and generally interaction and evolvability.

1 Introduction and Problem Survey

How do the dynamics of biological and artificial evolution in cooperative agents differ from individual agents? This is a key question as the research areas of swarm intelligence, collective robotics, and socio-biology gather more inertia. The experiment described here just begins to approach the answer.

Social Insects are a core inspiration for the new kinds of computation variously known as Artificial Life, Adaptive Systems, New AI, and conversely these techniques have influenced research in the biology of eusociality. Social Insects are viewed as collective agents, individually possessing only limited local information, together solving tasks and displaying a kind of group intelligence. The whole is more than the sum of the parts, and solutions emerge systematically rather than through explicit knowledge and planning. Certainly, the algorithmic nature of insect behavior has endeared them to computer scientists and physicists looking experiment in biology.

Collective central place foraging has occupied a place of special interest, since it clearly illustrates the trade off between individual and group goals, and the sophisticated strategies that can emerge. The possibility of an individual foraging ant cheating and hoarding food for itself is balanced against the shared protection against predators and competitors afforded by the colony and the shared genetic destiny of the colony. However, foraging with recruitment does not necessarily follow from eusociality, as its largely ecologically determined by the distribution of food in the environment [3]; small amounts of dispersed food or a high degree of predation could result in individual foragers, while large clumps of persistent food sources could result in collective behavior.

Perhaps the primary source of fascination with cooperative central place foraging is stigmergy. Interactions among foraging ants is indirect, through modifying the environment by self-reinforcing pheromone trails to food. Generally, stigmergy can be chemical modification or direct physical modification, such as the highly complex nest architectures of termites in the genus *Macrotermes* [12]. Stigmergic interactions are reminiscent of cultural behavior organized by artifacts and architecture, and lead to daydreaming of common dynamics in insect and human societies. Resnick manually programmed the classic simulation of recruitment foraging, using a simple set of rules in StarLogo [9].

Stigmergy has been successfully applied to computational tasks. Artificial pheromone based algorithms, known as Ant Colony Optimization, have been designed for many graph based problems, such as the traveling salesman problem, graph coloring, and job scheduling, and practically applied to optimizing routing tables on telecommunications networks [1]. These are also generally manually constructed, sometimes quite intricate and specific algorithms, and tuning these algorithms for a problem domain by methodology, rather than black arts, is being actively researched. Artificial evolution has promise in this research, if the variation in these algorithms can be coded in a standard way.

AntFarm was an early investigation into the evolution of cooperative foraging, with ants represented as artificial neural networks [2]. The core of the experimental design has figured as the basis for other experiments that followed over the years. Ants could sense food, pheromone, and their nest, and had the ability to pickup and drop food, drop pheromone, and move in specific directions. A variety of ANN architectures were tested, but optimal foraging failed to evolve. This was attributed to the relative complexity of using pheromone trails, compared to individual foraging, and the lack of competition or predation in the stable experimental environment.

No mention is made of the possibility of ANN being less than ideal for the task, and they go further to dismiss other representations, such as lisp s-expressions and finite state machines, as not being sufficiently

complex for the task. Paradoxically, in their concurrently published paper “The Genesys System”, on evolving artificial ants to solve the simpler task of following the “John Muir Trail”, Jefferson and Collins compare ANN and FSM based ants and conclude that “ANN-ants, as a class, are somewhat narrower and less powerful computationally than our FSA-ants” [4]. So goes the twisted path of research. Koza investigated cooperative foraging in a similar experimental setup, testing lisp s-expression representations in genetic programming [5]. These evolutionary runs did result nearly optimal strategies of lisp programs, employing pheromone and stigmergy.

Since these experiments in the early nineties, swarms of variations on the theme of evolving central place foraging have appeared in conference proceedings of Artificial Life, ECAL, From Animals to Animats, and Artificial Life and Adaptive Behavior journals. There should be pause to ponder the point of all these simulations. What’s so interesting about stitching together pre-bundled behaviors, into an already predicted strategy? It is the dynamics of the evolutionary process in collective agents that is of interest, as well as a successful end product. Really, this process is only slightly better understood in individual agents.

Maynard Smith and Szathmary list the evolution of cooperation, especially in eusocial insects, as one of the major transitions in evolution [8]. They discuss the selection steps and trade offs that lead to cooperation, and the resulting increase of “evolutionary material” and diversity. They then skip to discussion of human societies and the evolution of language, without discussing the characteristics of the new evolvability enabled by insect eusociality. It is this gap that the central place foraging simulation described here aims to address.

2 Experimental Design

The experiment evolves extended finite state machines, representing the ants, to enact central place foraging. Multiple trials were performed, with a variable number of agents in each simulation, and the resulting evolutionary trajectories and peak strategies are analyzed and compared. The design is discussed here, and some of the many possible alternatives are explored later in the conclusion.

Finite state machines (FSM) are among the simplest models of computation, capable of categorizing strings as members of regular languages [6]. FSM consist of a finite number of states, and transitions between those states responding to particular inputs. FSM begin in an initial state, and if the computation on a string ends in one of the designated final states, the string is accepted as a member of the class. Extended FSM in this experimental context do not embody regular languages; rather they specify actions in response to agent input. A callback action is added to each transition, and there is no concept of final state; the correspondence between these FSM and the series of inputs received in simulation is derived from the fitness function of the genetic algorithm.

FSM are not a directly biologically plausible representation for ant foragers, but are somewhat suitable description of observed behavioral repertoires. In discussing the colony castes and the division of labor, Wilson describes the construction of an ethogram, “which incorporates not only the repertory of a caste, but also the transition probabilities connecting individual acts and the time distributions in each act” [3]. He goes on to discuss the theoretical and practical difficulties in behavioral observation, which roughly corresponds to the gap between the actual neural and physiological structures of the ants, and the behavioral model.

An ethogram resembles closely an FSM, except that transitions are probabilistic; deterministic transi-

tions were chosen to simplify these initial experiments. Other computational models might represent more closely the actual ant, and may have other properties more suitable to applications. Again, FSM are a suitably simple starting point, and yet create a sufficiently large solution space for interesting evolutionary exploration. With n states, i inputs, and c callbacks, there are i^{n*c} possible FSM. For even a small number of states, inputs and callbacks, there are astronomical numbers of solutions.

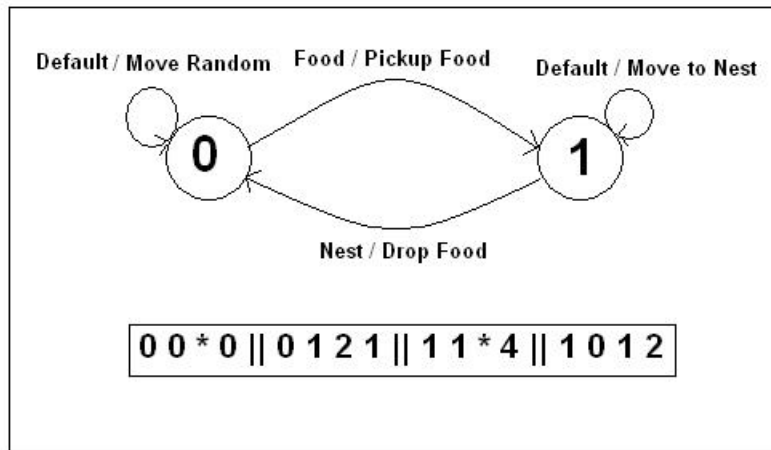


Figure 1: An example FSM encoding and actualization. This is the Basic Individual Strategy

FSM are represented as individuals in the genetic algorithm by a series of transitions, each as four integers. These integers are the In State, Out State, Input, and Callback. For this experiment, there are four states, four inputs, and six callbacks, and a maximum of sixteen transitions per genome. Additionally, a transition can be flagged as default, meaning that if the input presented to the FSM is not represented, this default transition is followed. Crossover always occurs between transitions, preserving individual “genes”. Mutation occurs on individual integers, by replacing that integer with a random value in its permissible range. Additionally, there is a probability that mutation can add a random gene or delete an existing gene.

Evolutionary runs began by constructing 50 random individuals, each with five transitions. Selection was by tournament selection, with 5% elitism. There was a 15% chance of mutation on each gene, so with variable length genomes there would be an average of a single mutation, and crossover occurred at a single point. When constructing the FSM from its genetic representation, there are often multiple transitions specified for a particular In State and Input, states that are completely unreachable, and transitions which would logically never be followed; for instance, transitions responding to pheromone input would not be followed if there are no callbacks for laying pheromone. These genes are ignored, and exist somewhat like non-coding regions of biological genomes; evolutionary processes at other loci may later activate these genes.

The fitness simulation occurs on a 25x25 toroidal arena. Each cell can have some quantity of food or pheromone, or be designated as the nest. The nest is set to the center, and three piles of 33 units of food are placed equidistant from the nest. A number of ants are initialized and placed in random locations on the arena. The ants consisted of the genetically determined FSM, its x-y location on the arena, its

heading (N, S, E, W), and a flag set to on when the ant is carrying food. The number of ants was 1, 10, or 30, and each variation was run at least twice. The simulation was run for 250 time steps for 10 or 30 ants, and 2500 time steps for one ant; this length was the least where the ants could theoretically collect all available food.

At each time step, all the ants are updated according to their current input. Input could have four possible values: 1, if the ant was located directly above the nest; 2, if food was present in any of the three cells the ant was facing; 3, if any pheromone was present in any of those three cells, and 0, if none of these were present. If any of these were present in combination, the input with the strongest “pull” was presented to the ant; the nest has more effect than food, which has more effect pheromone. This strict hierarchy was chosen for simplicity, over viable alternatives like more input values corresponding to combinations or probabilistic transitions. Pheromone was not evaporated; with the small simulation length it is unlikely that a food pile would be exhausted and so evaporation will have little effect on the result.




Figure 2: Initial configuration of the arena. The nest is central, and food piles are yellow.

The six callback actions are Move Random, Pickup Food, Drop Food, Drop Pheromone, Move to Nest, and Move to Pheromone. In Move Random, the ant moves forward to one of the three cells it’s facing, by its heading. If it moves to the facing left or facing right cell, its heading is rotated in that direction; this is also true for the other Move actions. In Pickup Food, a single unit of food is removed from the arena and assigned to the ant. If there is more than one facing cell containing food, the cell is chosen randomly, and if there is no food at all, the action is essentially a non-operation (no-op). The ant can only hold one unit of food, so if it’s already holding food, again the action is a no-op. Drop Food places a unit of food, if the ant is holding, directly on the arena location of the ant. If that location is a nest, the food is no longer present in the arena. Drop Pheromone increases the amount of pheromone on the ant’s arena location by one. Move to Nest moves the ant one cell closer to the nest, among the three facing cells. Move to Pheromone moves randomly to a facing cell containing pheromone; if there is no

pheromone, then this is a no-op.

Each action itself is composed of many sub-steps and abilities in real ants. For instance, ants can determine the direction of their nest by a combination of path integration and landmark navigation [11]. So, the purpose of this evolution is to stitch together these already complicated behaviors into a complete task of central place foraging. This may seem a simplistic challenge, with all the components given to the ants. It's straightforward to manually construct an efficient strategy from these actions. However it will be shown that the evolutionary trajectories are not straightforward, and unimagined strategies are also efficient.

Fitness is primarily scored by the amount of food stored in the nest, with the small addition of one point if any ant simply picks up food. This small evolutionary helping hand was added for expediency. For any food to be stored at all requires transitions where the ant moves, picks up the food, and drops the food on the nest. This basic strategy does always evolve without the helping hand, but can take a long time; without any fitness distinction among early strategies, evolution is essentially a random search. Some biological justification for this fitness point is envisioned by the energy requirements of individual foragers. While there is no explicit metabolism in this simulation, ants could not realistically forage for the colony if their individual requirements were not met.

3 Experimental Results and Analysis

Several evolutionary runs were conducted, with variable number of ants. Highly successful strategies evolved in each case. Only the simulation with 10 ants evolved strategies that employed using pheromone, and then not on every run. Some strategies were completely unexpected. Here, a single run with 10 ants is examined in detail, making notes of important shifts in the evolutionary process, and then three runs with 1, 10, and 30 ants are compared.

3.1 A Run Resulting in Stigmergy

The strategies described here are illustrated in Figures 3 and 4.

In this run, by **generation 2**, an ant evolved with fitness 1; that is, one of the ants simply picked up food. In fact, that is the whole of the strategy; the ants try to repeatedly pick up food, no matter the input. Only when an ant is initially randomly placed next to a food pile are they able to pick up any food. Immediately, this shows the close relationship between the evolutionary trajectory and the experimental design. If ants were initially all located at the nest, this strategy would never be selected.

In **generation 3**, the best strategy is essentially the same, with the difference being that ants that end up facing nothing will lay pheromone. Laying pheromone gives no advantage to these ants, but it also does no harm. The appearance of pheromone behaviors early in the evolutionary process is key for their incorporation into more complex strategies later. Without an “acclimation period”, when neutral pheromone behaviors hitch a genetic ride, later generations are unable to spontaneously adopt these actions. In other words, there are wide valleys in the fitness landscape.

In **generation 6**, they best strategy no longer relies on random chance in initial placement. A key gene, where the ant stays in state 0 moving randomly until it finds something, persists to the final generations. Also, two key combinations of input and action appear, with food picked up if found, and food dropped on the nest. However, in this FSM, these actions will never occur in the proper order; once in state 1




Figure 3: Generations 2, 3, 6, 11, 12, 20

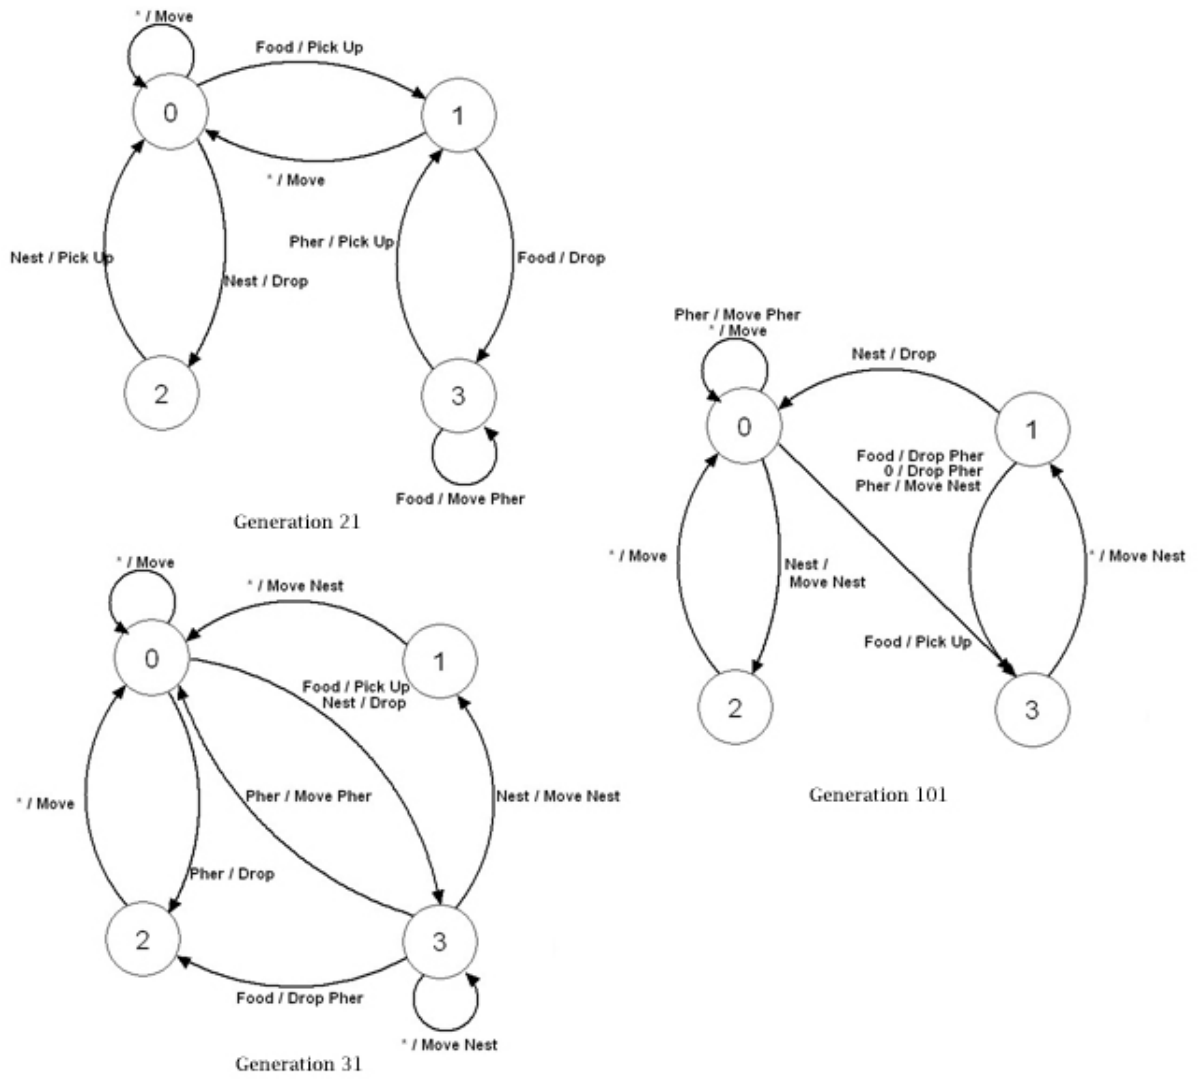


Figure 4: Generations 21, 31, 101

or 3, the ant is effectively paralyzed. The transitions leading from state 3 are not logically traversable; once in state 3, the ant will only be able to perceive the presence of the nest.

In **generation 11**, a single new transition enables actual central place foraging to occur. From state 1, reachable by finding and picking up food, there's a default transition back to state 0, with random movement. Then in state 0, the ant moves randomly to find the nest, drops the food, and becomes paralyzed. Generation 12 improves slightly, by mutating the callback in the transition from state 1 to 0 to move towards the nest. This one step towards the nest is beneficial enough for this strategy to dominate for several generations. This essential basic strategy is refined in generation 20, when the ant can continue random search after dropping food in the nest.

Generation 21 is the first of many quite odd strategies, which on first examination shouldn't work at all. The ants move randomly until food is found and pick it up. They then drop the food away from the pile, in response to the food input, and become stuck looking for pheromone. Only when another passing ant picks up the single piece of food and lays pheromone, does the trapped ant have a means of escape. This is the first appearance of cooperation! With any set of behaviors in these FSM, there are many pitfalls, where the ant ends up paralyzed. Adding pheromones adds to these paralysis dangers, which are only overcome by the ants aiding each other. This generation also demonstrates the tight reliance on implementation details. Ants pick up from facing cells, but drop food on their own location. This detail will be returned to later.

Generation 31 shows similar pitfalls of pheromones. The ants wander, pick up food from piles and drop pheromone, then continue wandering. If the ant finds the nest, it drops food, but it also drops food it finds pheromone. However, if an ant finds a single piece food, dropped by another in response to pheromone, the ant goes into another mode and directly returns to nest to deposit the food. A better strategy, where the ants follow pheromone to pick up food, can seemingly only evolve if there are cooperative methods to overcome problems with pheromone. Also note, the return to nest behavior traverses three different transitions; this makes this single behavior very susceptible to breakage in mutation and crossover operations, and hinders the long term viability of the strategy.

Finally, in **Generation 101**, after 70 generations of complex, fragile, and cooperation dependent strategies, a robust stigmergic strategy evolved. From the initial state, ants wander randomly or follow pheromone trails. When food is found it is picked up, and the ant heads directly for the nest, depositing pheromone along the way, in a two state cycle. This strategy is in fact slightly better than the one envisioned before the start of the experiment; if pheromone is already present on the path to the nest, the ant doesn't deposit pheromone, thereby increasing its speed. The transition through state 2 enables ants to move away from the nest, and to wander into another pheromone trail.

Figure 5 is a snapshot of the arena during a stigmergic run, illustrating the pheromone trails. Figure 6 graphs best and average fitness over 200 runs.

3.2 Comparison among variations

Comparing the runs, it is clear that the number of ants makes a fundamental difference in evolved strategies. Generally, early histories have similar path, before diverging into well defined, number dependent "behavior basins".

Stigmergic cooperation only evolved in runs with 10 ants, and then only a portion of the time. All other times, the basic individual strategy evolved (Figure 1, where the ant wanders until it finds food, then




Figure 5: Snapshot of the arena during a stigmergic run. Pheromone in green




Figure 6: Fitness of Best and Average Ants, Stigmergy Strategy

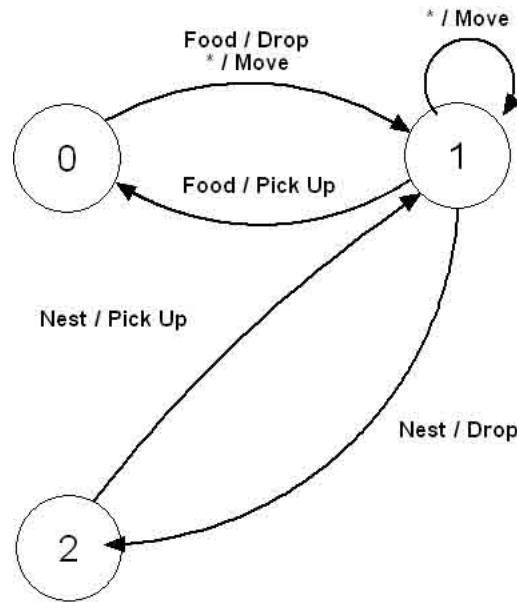


Figure 7: Best evolved strategy for single ant

heads straight back to the nest and deposits the food. Once present, without any competitive pheromone based strategies, this strategy becomes established throughout the population. Figure 9 shows the fitness history for a sample run.

The runs with a single ant initially developed a similar strategy, before discovering a much more effective strategy, shown in Figure 7. The ant wanders, and when food is found, picks it up. It then puts that piece down, picks up another and repeats. This can happen because of a small design detail; food is picked up from the cells the ant is facing, and placed on the cell where the ant is located. Effectively, the ant moves an entire pile of food one square. If the ant is positioned between the food pile and the nest location, it will move one step towards the nest and begin again. By this method, the ant only needs to find the pile of food once, to then drag the entire pile of food to the nest.

This strategy would also be effective with 10 ants, but never developed. The reason is the difference in length of the simulation; 2500 steps for the single ant, 250 for the collective ants. Dragging a single pile of 33 pieces of food the 12 cell distance to the nest takes at least $33 * 12 = 396$ steps. The time difference is implemented to promote fairness in the fitness function, since with the basic strategy one ant would take at least 10 times longer to collect the same amount of food. However, it turns out to be a key factor in evolvability.

The 30 ant simulations also initially developed the basic individual strategy, before developing another unexpected strategy, shown in Figure 8. It's essentially similar to the dragging strategy in the single ant, but it works very differently in this context. Since there are so many ants wandering the arena, a single ant dragging food will soon attract other workers to drag the new pile. The ants will line up, between the food pile and the nest, and pass pieces of food to each other, just like a bucket brigade of fire fighters passing water along a line. A snapshot of an ant bucket brigade is shown in Figure 12.

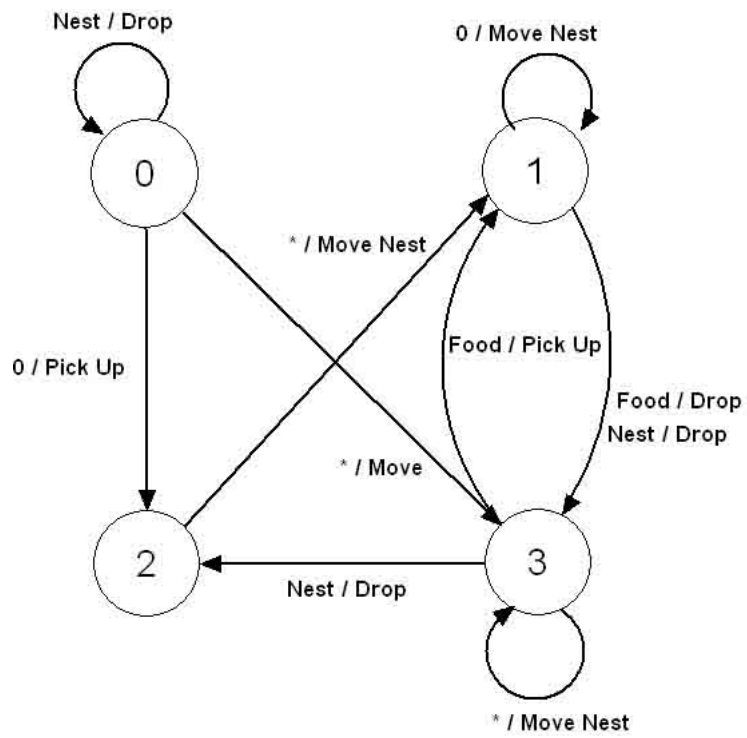


Figure 8: Best evolved strategy for 30 ants

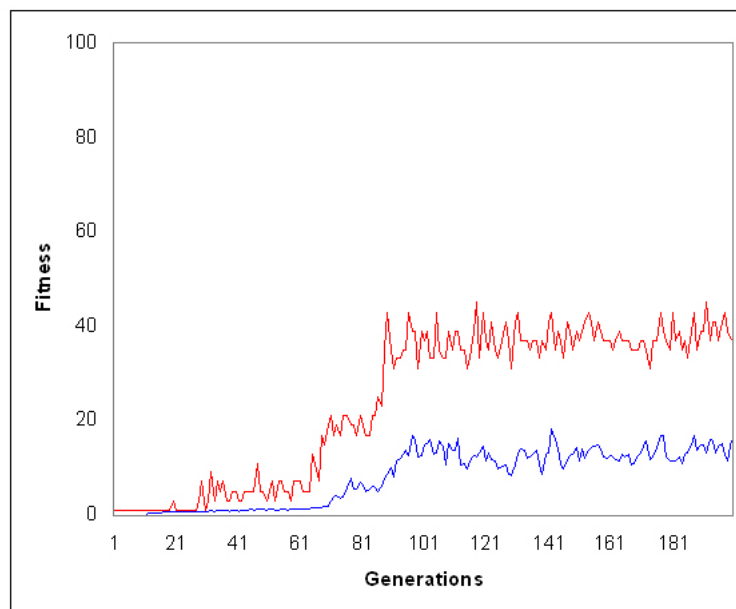


Figure 9: Best and Average Fitness for 10 ants, Basic Strategy




Figure 10: Best and Average Fitness for a Single Ant




Figure 11: Best and Average Fitness for 30 ants




Figure 12: A snapshot of a bucket brigade

4 Conclusions and Future Studies

So, how adaptable are swarms? The results don't present any straightforward answers, but do suggest additional questions and methodologies. Definitely, it's a thrill to carry around population of evolving ants on a laptop.

Different strategies evolved in experiments with different number of clones, effectively opening the possible solutions to the problem space. Seth investigated this in [10], without providing any explicitly collective behaviors to the agents and found "behaviors adapted to a group situation may be suboptimal when expressed by an isolated individual". This applied to intermediate evolutionary steps in FSM foraging, but was not precisely true of end product strategies. Any of the end product strategies would apply to the other, and the individual and 30 ant strategies were essentially the same, though actualized different behavior. Intermediate steps, especially involving pheromone strategies, were only successful through cooperation. This conclusions are judged solely by analysis; it would be worthwhile to follow Seth, who explicitly tested evolved strategies in different scenarios and recorded detailed statistics.

Each run variation followed a trajectory through the "basic strategy". This is interesting, in comparing biological and artificial evolution. Eusocial ants evolved from eusocial wasps with the ability to survive, to forage, build nests, defend themselves, brood care. This didn't include stigmergy for food foraging, since of course wasps forage by flight. In contrast, artificial evolution generally proceeds from random, without any survival requirements, to enable evolution with the least possible bias. Does this approach have anything to inform evolutionary theory, and what role does organismic history play in evolvability? The emergence of a common strategy in all trajectories does lead some credence to this application, a sort of inevitability in the problem domain.

Still it would be worth simulated evolution from an established strategy. Specifically with stigmergy, this would have its difficulties. Stigmergic strategies only evolved from a diverse population, with pheromone behaviors well represented. It's not impossible for stigmergy to evolve from a meta-population of the "basic strategy" but it would take a much longer evolutionary process. There is a missing element; in biological ants, direct chemical communication predated indirect communication. With the ability to emit and detect pheromones, a step where pheromone is deposited in the landscape is not too great. These abilities could be added to the FSM representation by adding the ability to sense other ant's state. This approach also has application to purely optimization oriented swarm evolution. With quite complicated, parameter dependent algorithms, evolution could tune already devised strategies.

In biology, group evolution and eusociality do not infer cooperate central place foraging. Foraging strategy is quite dependent on ecological conditions. The desert ant *Cataglyphis* is a solitary forager, each forager attending to a specific food source found by solar navigation. Recruitment would pose a danger to the entire colony, because of heavy predation and the increased detectability of large numbers of ants [3]. Other species can exhibit different strategies with flexibility, foraging individually or collectively depending on the food source [7]. Generally, allocation of colony resources, in the form of individuals taking on tasks, is highly dependent on the energetic constraints of individuals in each caste and the requirements of the colony.

This leads to a number of issues, centered on representation of agents and their environment. It suggests multitudes of variations. Ants could belong to castes, perhaps by activating different subsets of the genome, in colony evolved ratios. Food distribution is tightly tied to behavior, as is the presence of predators. These domains would be fruitful for systematic exploration. The precise nature of senses and actions turned out to have large consequences in these experiments; food dragging in individuals and bucket brigades in groups were very dependent on how food was handled. FSM seemed to be a worthwhile representation. Optimal strategies evolved after testing individuals amounting to an order of magnitude less than in genetic programming experiments. Still they're not without question. Though somewhat analogous to ethograms, behaviors in these FSM traversed multiple, reused states, in contrast to the clean divisions of ethograms. This shows FSM to be more related to neural like models, where state is a more implicit quantity. It would be worth looking a various neural network models, as well as probabilistic FSM and deterministic FSM with greater number of states.

With so many representational possibilities and such strong dependence on representation, can simulations say anything generally about evolvability of cooperation? It's possible to imagine certain representations of individuals that would select against any cooperation, seemingly invalidated any generalities on group behavior. One direction is to model agents on the most abstract ecological conception of an organism, as purely energy transfer and transformation. Further, the results suggest this line of research extends beyond modeling a single species. The core issue is evolvability with interaction, and interference, and this can certainly take place in a community or ecosystem. In any case, both abstraction and complication of the representation has dangers; it could throw the brood out with the bath water.

References

- [1] Eric Bonabeau, Marco Dorigo, Guy Theraulaz (1999). *Swarm Intelligence: from natural to artificial systems*. Oxford University Press.
- [2] Robert J. Collins, David R. Jefferson (1992). *Ant Farm: Towards Simulated Evolution*. *Artificial Life II*. Addison-Wesley.
- [3] Bert Holdobler and Edward Osborne Wilson (1991). *The Ants*. Springer.
- [4] David Jefferson, Robert Collins, Claus Cooper, Michael Dyer, Margot Flowers, Richard Korf, Charles Taylor, Alan Wang (1992). *The Genesys System: Evolution as a theme in artificial life*. *Artificial Life II*. Addison-Wesley.
- [5] Koza, John (1992). *Genetic Programming, On the Programming of Computers by Means of Natural Selection*. MIT Press.
- [6] Martin, John C. (1991). *Introduction to Languages and the Theory of Computation*. McGraw Hill.
- [7] J. L. Mercier, A. Lenoir (1999). Individual flexibility and choice of foraging strategy in *Polyrhachis laboriosa*. *Insectes Soc.*, 46, 267-272
- [8] John Maynard Smith and Eors Szathmary (1995). *The Major Transitions in Evolution*. Oxford University Press.
- [9] Resnick, Mitchel (1997). *Turtles, Termites, and Traffic Jams*. MIT Press.
- [10] Seth, Anil K. (2002). Modelling Group Foraging: Individual Suboptimality, Interference, and a Kind of Matching. *Adaptive Behavior*, 9(2), 67-90.
- [11] Rudiger Wehner, Barbara Michel, Per Antonsen (1996). Visual Navigation in Insects: Coupling of Egocentric and Geocentric Information. *J. exp. Biol.* 199 129-140.
- [12] Wilson, Edward O. (1971). *The Insect Societies*. Harvard University Press.