

Swarm Intelligence: Power in Numbers

Appears in *Communications of the ACM*, August 2002, 45(8), 62-67.

Peter Tarasewich

College of Computer and Information Science
Northeastern University
360 Huntington Avenue, 161CN
Boston, Massachusetts 02115
617-373-2078 (voice), 617-373-5121 (fax)
tarase@ccs.neu.edu

Patrick R. McMullen

Management Department
College of Business
Auburn University
Lowder Business Building
Auburn, Alabama 36849
334-844-6511 (voice), 334-844-5159 (fax)
pmcmullen@business.auburn.edu

ACM COPYRIGHT NOTICE. Copyright © 2002 by the Association for Computing Machinery, Inc. Permission to make digital or hard copies of part or all 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. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Publications Dept, ACM Inc., fax +1 (212) 869-0481, or permissions@acm.org.

Even with today's ever-increasing computing power, there are still many types of problems that are very difficult to solve. Particularly, combinatorial optimization problems continue to pose challenges. An example of this type of problem can be found in product design. Take as an example the design of an automobile based on the attributes of engine horsepower, passenger seating, body style, and wheel size. If we have three different levels for each of these attributes (such as those in Table 1), there are 3^4 , or 81, possible configurations to consider. For a slightly larger problem with 5 attributes of 4 levels each, there are suddenly 1024 combinations. Typically, an enormous amount of possible combinations exist, even for relatively small problems. Finding the optimal solution to these problems through enumeration of all possible combinations is usually impractical. Fortunately, search heuristics have been developed to find good solutions to these problems in a reasonable amount of time.

Table 1 – Automobile Design Example	
Attribute	Levels
Engine horsepower	150, 180, 210
Passenger seating	4, 5, 6
Body style	coupe, sedan, wagon
Wheel size	14", 15", 16"

Heuristics and the Natural Sciences

Over the past decade or so, several heuristic techniques have been developed that build upon observations of processes in the physical and biological sciences. Examples of these are Genetic Algorithms and simulated annealing. These techniques have been widely and successfully used to solve a variety of problems, including those of a combinatorial nature.

Genetic Algorithms (GA) use biological methods such as reproduction, crossover, and mutation to quickly search for solutions to complex problems. GA start with a set (usually random) of possible solutions to a given problem, each solution represented by a string of bits or characters. A fixed number of solutions that are the most “fit,” i.e. the best current solutions to a problem, are carried over to the next generation and chosen to reproduce. In other words, the best current solutions are saved, and are used to generate new solutions. These strings (current solutions) can be altered to produce child (new) solutions using genetic operators. One of these genetic operators is crossover, in which a randomly chosen sequence of bits is exchanged between two strings to produce two new strings. Another operation is mutation, which randomly alters bits within a single string to produce a new string. The mutation function is included to keep from becoming trapped at a local optimum. The new population is evaluated, a new reproducing population is chosen, and the process of crossovers and mutations repeats for a predetermined number of iterations or until an acceptable evaluation level is reached. More information on this approach can be found in [11].

Simulated annealing (see [11] for a good introduction) borrows from the physical annealing of solids – heating a solid to a high temperature and slowly cooling it until desired properties of the solid are obtained. When simulated annealing begins, an initial solution to a problem is generated and used as the first current solution. The “temperature” is then systematically decreased, and neighboring solutions to the current solution are found. If the objective function value is superior to that of the current solution, the neighboring solution becomes the current solution. If the neighboring solution provides an objective function value that is inferior to that of the current

solution, the neighboring solution may still become the current solution if a certain acceptance criterion is met. This occasional acceptance of inferior solutions allows the search to move to a different location on the continuum of feasible solutions in an effort to reach the global optimum. The process of finding and evaluating neighboring solutions is repeated until some stopping criteria is met.

Insect Behavior

More recently, scientists have been turning to insects for ideas that can be used for heuristics. Many aspects of the collective activities of social insects, such as ants, are self-organizing, meaning that complex group behavior emerges from the interactions of individuals who exhibit simple behaviors by themselves. Examples of these collective activities among ants are finding food and building a nest. The results of self-organization are global in nature, but come about from interactions based entirely on local information. To achieve this, self-organization relies on several components: positive feedback, negative feedback, amplification of fluctuations, and multiple interactions [4].

Positive feedback comes from basic behavioral rules, such as the recruitment of other insects to forage a food source, which create the necessary structure for collective behavior. Negative feedback is provided by limitations on behavior caused by events such as the depletion of a food source. Amplification of fluctuations refers to the necessity of random events, such as an ant getting lost but finding a new source of food to exploit. Self-organization also requires the multiple interactions of individuals. These interactions can be direct (e.g., physical, visual, or chemical contact) or indirect. Indirect

contact can take the form of stigmergy, where one individual performs an action based on what was performed earlier by a different individual. An example of this can be seen from the construction of wasp nests, where certain configurations of existing cells trigger the creation of a new cell.

The expression “swarm intelligence” was originally used in the context of cellular robotic systems to describe the self-organization of simple mechanical agents through nearest-neighbor interaction. Bonabeau, Dorigo, and Theraulaz [4] recently extended its definition to include “any attempt to design algorithms or distributed problem-solving devices inspired by the collective behavior of social insect colonies and other animal societies.” (p. 7) This includes the behaviors of certain ants, honeybees, wasps, cockroaches, beetles, and termites. Solutions to many different problems have already been found using heuristics based on the behavior of ants searching for food.

Ants Looking for Food

Although the capabilities of a single ant are very limited, ants can collectively establish the shortest route between a source of food and their nest, and efficiently move the food to their home. Ants communicate with each other through the use of pheromones, chemical substances that attract other ants. As the ants move, they lay down a trail of these pheromones that other ants can follow. Ants move randomly, but when they encounter a pheromone trail, they decide whether or not to follow it. If they do so, they lay down their own pheromone on the trail as well, reinforcing the pathway. The probability that an ant chooses one path over another increases with the amount of

pheromone present. The more ants use a given trail, the more attractive that trail becomes to other ants.

If a colony of ants are presented with a short path and a long path to a source of food, they will first use both paths in equal numbers, laying down pheromone as they go. But the ants taking the shorter path will return to the nest first. The shorter pathway will then be doubly marked with pheromone, and will be more attractive to those ants that return to the food source. This is illustrated by Figure 1, which shows a distribution of ants over a set of pathways between a nest and a food source over time. Early on, the ants are equally distributed, but eventually they favor the shorter route.

There is, however, always a chance that an ant will not follow a previous or well-marked trail. This allows for randomness and exploration, which is beneficial in that it allows for discovery of shorter or alternate pathways, or new sources of food. The pheromones also evaporate over time. Given this, pheromone will become less detectable on longer trails, since these trails take more time to traverse. The longer trails will hence be less attractive, another benefit to the colony as a whole.

Problem Solving Using Ants

Ant colony optimization (ACO) is a meta-heuristic that uses artificial ants to find good solutions to difficult combinatorial optimization problems. The behavior of artificial ants is based on the traits of real ants, plus additional capabilities that make them more effective, such as a memory of past actions. Each ant of the “colony” builds a solution to the problem under consideration, and uses information collected on the problem characteristics and its own performance to change how other ants see the problem. A

more detailed explanation of ACO can be found in [7]. *ACO algorithm* refers to any instantiation of this meta-heuristic. The more informal term *ant algorithm* is used to describe algorithms that, in general, follow ACO guidelines but do not necessarily follow all aspects of the meta-heuristic.

Dorigo, Maniezzo, and Coloni [8] have used this approach to address the classic traveling salesman problem. In this problem, a person must travel among a set of cities, visiting each city only once. A colony of artificial ants is created which first randomly travels complete circuits which contain each city. During this first step, local travel to closer cities is favored. After a complete circuit is determined, “pheromone” is deposited on each of its links. The amount of pheromone is inversely proportional to the length of the circuit; shorter distances receive more pheromone. The colony is then released to travel circuits again, but this time ants favor links with higher concentrations pheromones in addition to links that are shorter. The pheromone evaporates at a constant rate, and links that are not part of good overall circuits eventually fall out of favor. The ant approach to this problem also provides the advantage of backup routes. Since the ants are continuously exploring different paths, alternative routes already exist if the link between two cities becomes unusable (e.g., bad weather makes travel impossible between two cities).

The closely related vehicle routing problem was investigated by Bullnheimer, Hartl, and Strauss [5] using an ant approach. The goal of this problem is to minimize the cost for vehicles going from a central depot to a series of customers. Each customer must be visited exactly once by only one vehicle, and each vehicle begins and ends its journey at the depot. There are also restrictions on vehicle capacity and total distance

traveled by each vehicle. Artificial ants travel along pathways from the depot to the cities and back again, with pheromones influencing which pathways they choose next. The pheromone levels on the pathways are updated according to the quality of the solutions found by the ants.

Ant-colony optimization in conjunction with tabu search was used by Bland [3] to address the problem of space planning. The problem specifically involves trying to optimally locate activities within a given space. For example, an organization might want to locate administrative functions and employees among a set of offices in a building. In this case, it might be most efficient to minimize the amount of movement (of both people and paperwork) within the building.

An ant colony heuristic was developed by Bauer, Bullnheimer, Hartl, and Strauss [2] for the single machine total tardiness problem. Here, a series of jobs to be performed on one machine must be scheduled. The goal is to minimize the total tardiness of the jobs, where tardiness is defined as the difference between the job's completion time and its due date (or zero if the job was completed on time). A set of ants independently choose successive jobs until all jobs are scheduled. Jobs are chosen based on how good the job seems to be (based on its due date) and how good a choice the job actually was in former runs (based on pheromone levels left behind). Each sequence of jobs is evaluated, and the best solution is used to update the pheromone trails.

Song, Chou, and Min [12] implemented a search algorithm based on ant foraging to solve the power economic dispatch problem. In this problem, total fuel cost at thermal plants must be minimized while meeting total power demand and producing electricity within the production limits of each plant's generator. Artificial ants move from plant to

plant, distributing the total power demand as they go, until all plants have been utilized. The choice of the next generator on the ant's path is based on the pheromone levels left behind by other ants and the perceived cost of using the generator.

Other applications of ant colony optimization include bus driver scheduling, graph coloring and partitioning, load balancing in telecommunications networks, and routing in packet-switched networks. Additional sources of information on ACO are Marco Dorigo's Web site (<http://iridia.ulb.ac.be/~mdorigo/ACO/ACO.html>) and the book by Bonabeau, Dorigo, and Theraulaz [4].

Some of the research described above has tested the performance of the ant-based heuristics and compared the results to Genetic Algorithm, simulated annealing, and other search approaches. This information is summarized in Table 2. Neural networks and Tabu Search, which are also well-established techniques used for problem solving [11], are included in the table. For several test cases of the traveling salesman problem, ant methods seemed to outperform Genetic Algorithm and simulated annealing approaches, but not Tabu Search. Vehicle routing results, on average, were better for the ant heuristic when compared to simulated annealing and a neural network approach for a set of fourteen sample problems, but did not surpass three Tabu Search algorithms. Ants worked better for the single machine total tardiness problem on 125 benchmark problems when compared to two different simulated annealing approaches. Slightly lower costs were found using an ant approach, versus a Genetic Algorithm, on a small test case for the power economic dispatch problem. Overall, ant techniques seem to provide some performance advantages over the natural science techniques described earlier, but do not seem to have an advantage over Tabu Search. These results, although mixed, are still

encouraging. However, it must be remembered that this statement is based on a small sample of documented problems.

Problem	Genetic Algorithms	Simulated Annealing	Tabu Search	Neural Networks
Traveling Salesman	Better	Better	Same	---
Vehicle Routing	---*	Better	Worse	Better
Single Machine Tardiness	---	Better	---	---
Power Economic Dispatch	Better	---	---	---

*Not tested. Table summarized from data in [4], [2], [5], [12]

Other Insect Behaviors

There are many other collective insect behaviors which fall under the umbrella of swarm intelligence and can be applied to problem solving. Certain species of ants collectively sort their young according to various stages of development. Even though different types of young may begin randomly scattered throughout a given area in the nest, eggs and the smallest of larvae will end up in the center, with the sizes increasing and the largest larvae at the edges. An explanation for this is that the ants will tend to group things according to sameness. Items that look out of place within a group will be picked up and moved. Items will be dropped next to those that look similar. Researchers have already applied this behavior to problems dealing with the sorting and grouping of database information [4]. An example of this is a bank trying to determine which people will receive approval for loans based on various application data. Other problems that might benefit from this approach include stock analysis, product line design, analysis of online auction purchase patterns, and the dynamic placement of Web advertisements based on user behavior within a site.

In another type of collective behavior, ants will work together to retrieve prey or a piece of food that is too large for one individual to move – a process known as cooperative transport. While pushing the object, each ant will periodically change position and alignment until the object moves toward the nest. This behavior has been implemented in the field of robotics, where a group of independently acting devices performs a task such as moving a box to the center of an arena [4]. Other problems that could be solved by modeling this type of behavior include assembly line design and balancing, especially in flexible manufacturing environments. Workers would team up for certain extraordinary tasks to maximize the overall efficiency of the line. Additional “division of labor” problems might be addressed in a similar manner.

Honeybees will perform tasks based on age. Older bees may forage for food, while younger bees will stay at the hive and nurse the young. The allocation of these tasks, however, may change when demand dictates. For example, when food is scarce, younger bees will also forage for food. This concept has been applied to flexible manufacturing processes [1]. For example, an expensive piece of assembly equipment is dedicated to a specific task, but will change tasks if it senses that it is advantageous to do so.

Similarly, ants perform tasks according to specific “job classifications”. These are primarily forager, patroller, nest maintainer, and midden worker [10]. Patrollers scout the area surrounding the nest for signs of danger, and determine the general direction that foragers will look for food. Nest maintenance workers keep the entrance to the nest clear and remove soil from inside the nest. Other workers sort and maintain the midden (refuse) pile. Sometimes ants will switch tasks when an urgent situation arises, but for

the most part worker classifications are fixed. Perhaps heuristics can be developed that incorporate the notion of different types of ants, and ant activities, occurring simultaneously to more effectively solve a problem.

The pheromone marking of trails discussed previously is not limited to ants. Instead of performing a “dance” to communicate the location of food, certain stingless bees will deposit scent marks between a nest and a food source [9]. Many times the placement of these marks is irregular, and the pheromones evaporate within a matter of minutes. Sometimes this is not sufficient to guide other bees to the food, so a scout will initially lead them to the source. Bees, however, can also mark unproductive food sources with pheromones as a signal to other bees not to spend too much time there trying to extract nourishment. Maybe particular problems can benefit from techniques that make use of the idea of marking both good and bad alternatives.

Certain species of caterpillars forage for food as well, but do not bring the food back to the nesting site [6]. Caterpillars will usually leave two types of trails, prefeeding and postfeeding. Prefeeding (or exploratory) trails mark the path the caterpillar has been following, and may or may not lead others to food. A postfeeding trail is usually meant to act as a recruitment trail that will lead others to the food source. In a twist to this, however, one species of caterpillar sometimes leaves a postfeeding trail to a new resting site, undermining the use of these trails as pathways to food. But perhaps there is something of value in this behavior as well. This approach might be used to devise schemes that do not get trapped at local optima, and search groups of good solutions more efficiently.

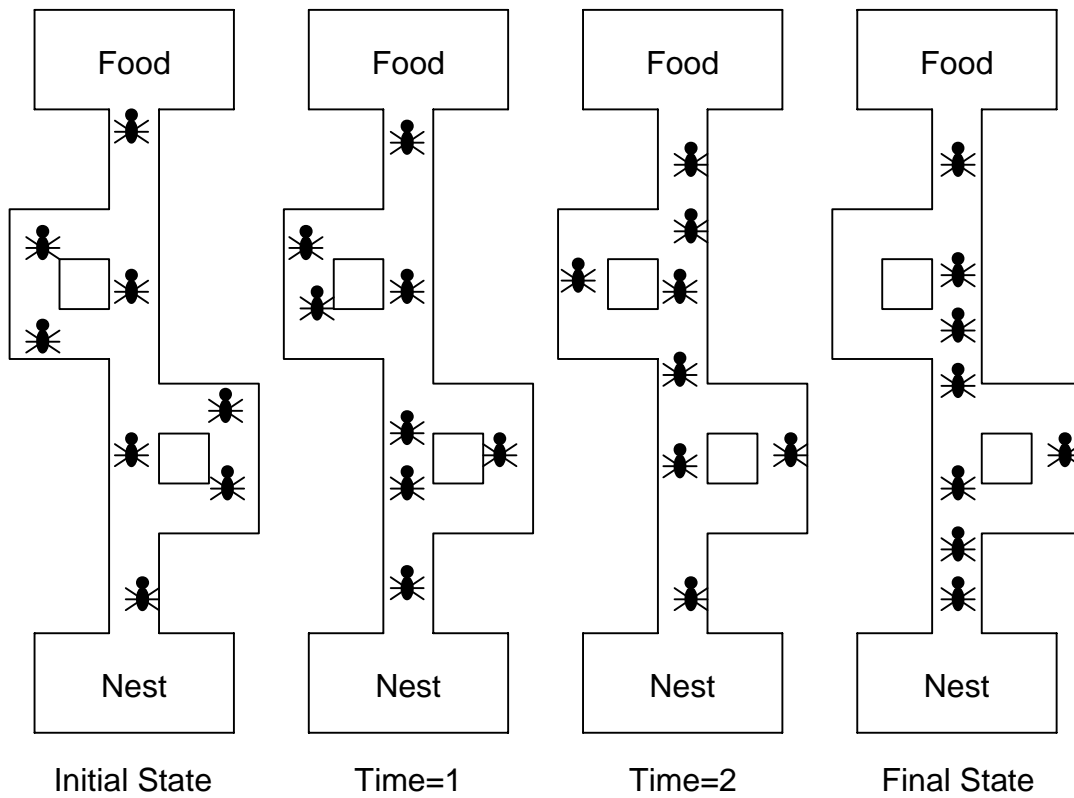
Swarming into Action

Scientists are still learning about the collective social behaviors of insects, and researchers are just beginning to apply what is known about swarm intelligence to heuristic design. The studies that have been done to date, especially with ant colony techniques, show the potential of this approach to effectively find good solutions to many types of practical problems which would otherwise remain difficult to solve. Many other collective insect behaviors remain to be investigated. The world is just beginning to see the great power that can be collectively harnessed from the simple behaviors of these insects.

REFERENCES

1. Anderson, C., and Bartholdi, III, J. J. Centralized Versus Decentralized Control in Manufacturing: Lessons From Social Insects. In *Complexity and Complex Systems in Industry* (McCarthy, I. P., and Rakotobe-Joel, T., eds.). University of Warwick, UK, 2000, 92-105.
2. Bauer, A., Bullnheimer, B., Hartl, R. F., and Strauss, C. An Ant Colony Optimization Approach for the Single Machine Total Tardiness Problem. *Proceedings of the 1999 Congress on Evolutionary Computation*, 1999, 1445-1450.
3. Bland, J. A. Space-Planning by Ant Colony Optimisation. *International Journal of Computer Applications in Technology*, 12, 6 (June, 1999), 320-328.
4. Bonabeau, E., Dorigo, M., and Theraulaz, G. *Swarm Intelligence: From Natural to Artificial Systems*. Oxford University Press, New York, 1999.
5. Bullnheimer, B., Hartl, R. F., and Strauss, C. An Improved Ant System Algorithm for the Vehicle Routing Problem. To appear in *Annals of Operations Research: Nonlinear Economic Dynamics and Control* (Dawid, Feichtinger, and Hartl, eds.). 1999.
6. Costa, J. T., and Pierce, N. E. Social Evolution in the Lepidoptera: Ecological Context and Communication In Larval Societies. In *The Evolution of Social Behavior in Insects and Arachnids* (Choe and Crespi, eds.). Cambridge University Press, Cambridge, UK, 1997, 407-442.

7. Dorigo, M., Di Caro, G. and Gambardella, L. M. Ant Algorithms for Discrete Optimization. *Artificial Life*, 5 (1999), 137-172.
8. Dorigo, M., Maniezzo, V., and Coloni, A. The Ant System: Optimization by a Colony of Cooperating Agents. *IEEE Transactions on Systems, Man, and Cybernetics – Part B*, 26, 1 (January, 1996), 29-41.
9. Free, J. B. *Pheromones of Social Bees*. Comstock Publishing Associates, Ithaca, New York, 1987.
10. Gordon, D. M. *Ants at Work: How an Insect Society is Organized*. The Free Press, New York, 1999.
11. Pham, D. T., and Karaboga, D. *Intelligent Optimisation Techniques*. Springer-Verlag, London, 2000.
12. Song, Y. H., Chou, C. S. V., and Min, Y. Large-Scale Economic Dispatch by Artificial Ant Colony Search Algorithms. *Electric Machines and Power Systems*, 27, 7 (July 1999), 679-690.

Figure 1 – Ants Presented with Alternate Pathways¹

¹ Adapted from Goss, S., Aron, S., Deneubourg, J. L. and Pasteels, J. M. Self-organized short-cuts in the Argentine ant. *Naturwissenschaften*, 76 (1989), 579-581.