Connection Management by Ants: An Application of Mobile Agents in Network Management

Tony White, Bernard Pagurek, Franz Oppacher[1]
Systems and Computer Engineering, Carleton University,
1125 Colonel By Drive, Ottawa, Ontario, Canada K1S 5B6
email: {tony, bernie}@sce.carleton.ca

Abstract

This paper describes how biologically-inspired agents can be used to solve control and management problems in Telecommunications. These agents, inspired by the foraging behavior of ants, exhibit the desirable characteristics of simplicity of action and interaction. The collection of agents, or swarm system, deals only with local knowledge and exhibits a form of distributed control with agent communication effected through the environment. In this paper we explore the application of ant-like agents to the problem of routing in circuit switched telecommunication networks.

Introduction

The notion of complex collective behavior emerging from the behavior of many simple agents and their interactions is central to the ideas of Artificial Life [Langton 87]. There are many examples in Nature of social systems where individuals possess simple capabilities which, when compared to their collective behaviors, are much more complex. Such systems span many levels of evolutionary complexity, from simple bacteria [Shapiro 88], to ants [Goss et al, 90], [Franks 89], caterpillars [Fitzgerald and Peterson 88] and beyond.

The continuing investigation and research of naturally-occurring social systems offers the prospect of creating artificial systems that are controlled by emergent behavior and promises to generate engineering solutions to distributed systems management problems found, for example, in telecommunications networks.

Controlling distributed systems such as those found in telecommunications networks by means of a single central controller, or requiring each controlling entity to have a global view of the system, has many disadvantages. In the case of the single controller, a considerable quantity of information must be communicated from the network to the controller, necessitating the sending of data from all parts of the network to the centralized control point. These systems scale badly due to the rapid increase in the quantity of data that must be transferred and processed to the central point as the network increases in size. Such systems invariably have to deal with data that is time delayed, i.e. stale. Providing a single point of control also provides for a single point of failure, a highly undesirable characteristic of any system. In the case where multiple global views are constructed and maintained, the problem of synchronization of such views can lead to instability and can lead to excessive use of network resources. The optimal design of a centralized controller is often difficult to achieve in that design decisions must be made based upon a static (and idealized) view of the way in which demands on resources in the network are likely to change. Decentralized control mechanisms need not suffer from the above problems and potentially can take advantage of local knowledge for improved use of network resources.

In this paper we describe the essential principles of Swarm Intelligence (SI) and how an understanding of the foraging behaviors of ants [Beckers et al 92] has led to new approaches to control in telecommunications networks.

This paper consists of seven subsequent sections. In the next section, a brief overview of Swarm Intelligence and Ant Colony search is presented. There then follows a brief description of the important attributes of the routing problem from the perspective of this paper. The next section describes the Routing By Ants system and provides an overview of the algorithm used. Experimental setup and results are then described and the final sections provide conclusions and future work that is planned with this system and derivatives of it.

---

[1] Email: oppacher@scs.carleton.ca

Swarm Intelligence and the Ant Colony

Swarm Intelligence [Beni and Wang 89] is a property of systems of unintelligent agents of limited individual capabilities exhibiting collectively intelligent behavior. An agent in this definition represents an entity capable of sensing its environment and undertaking simple processing of environmental observations in order to perform an action chosen from those available to it. These actions include modification of the environment in which the agent operates. Intelligent behavior frequently arises through indirect communication between the agents, this being the principle of stigmergy [Grasse' 59]. It should be stressed, however, that the individual agents have no explicit problem solving knowledge and intelligent behavior arises as a result of the actions of societies of such agents.

Individual ants are behaviorally simple insects with limited memory and exhibiting activity that has a random component. However, collectively ants manage to perform several complicated tasks with a high degree of consistency. Examples of sophisticated, collective problem solving behavior have been documented [Frank 89; Hölldobler and Wilson 94] including:

- Forming bridges
- Nest building and maintenance
- Cooperating in carrying large items
- Finding the shortest routes from the nest to a food source
- Regulating nest temperature within a one degree Celsius range
- Preferentially exploiting the richest source of food available.

In the examples listed above, two forms of stigmergy have been observed. Sematectonic stigmergy involves a change in the physical characteristics of the environment. Ant nest building is an example of this form of communication in that an ant observes a structure developing and adds its ball of mud to the top of it. The second form of stigmergy is sign-based. Here something is deposited in the environment that makes no direct contribution to the task being undertaken but is used to influence the subsequent behavior that is task related.

Sign-based stigmergy is highly developed in ants. Ants use highly volatile chemicals called pheromones (a hormone) to provide a sophisticated signaling system. Ants foraging for food lay down quantities of pheromone marking the path that it follows with a trail of the substance. An isolated ant moves essentially at random but an ant encountering a previously laid trail will detect it and decide to follow it with a high probability and thereby reinforce it with a further quantity of pheromone. The collective behavior which emerges is a form of autocatalytic behavior where the more the ants follow the trail the more likely they are to do so. The process is characterized by a positive feedback loop, where the probability that an ant chooses any given path increases with the number of ants choosing the path at previous times.

The use of ant foraging behavior as a metaphor for a problem-solving technique is generally attributed to Dorigo [Dorigo et al 91]. However, since his early work on the Travelling Salesman Problem (TSP) and Asymmetric TSP, the technique has been applied to several other problem domains. These include the Quadratic Assignment Problem (QAP) [Maniezzo et al 94, Taillard et al 97], graph coloring [Costa and Hertz 97], vehicle routing [Bullnheimer et al 97] and, as we shall see in the next section, communications network routing.

Routing

Many strategies have been proposed and researched for network routing. For the purposes of this paper, two characteristics of such strategies are considered important.

Firstly, networks are either packet or circuit switched. A packet switched network is one in which routing decisions are made on a packet-by-packet basis. In this case, no fixed connections are made and resources are not reserved. In a circuit switched network a connection is maintained for the duration of a session between two or more entities. Resources are allocated to the circuit for the duration of the session.

Secondly, network routing can be either static or adaptive. Static routing usually employs shortest path algorithms such as provided by Dijkstra's algorithm [Dijkstra 59] in order to compute routing tables that are subsequently downloaded to the network. Adaptive routing uses link cost metrics, which are functions of the utilization of network resources in order to force changes in routing during periods of network congestion.

Considerable information on routing can be found in [Schwarz 89] and [Tanenbaum 88].

The motivation for exploiting the ant metaphor for routing in telecommunications networks arises from the fact that routing systems frequently depend upon global information for their efficient operation. Ant systems do not need such global information, relying instead upon pheromone traces, or rather their digital equivalent, that are laid down in the network as the ant, or agent, moves through the network. Global information is frequently out of date and transmission of the information required from one node to all others consumes considerable network resources. Ideally, we would like to have the network adapt routing patterns to take advantage of free resources and move traffic if possible. This is particularly desirable in broadband networks where traffic patterns change rapidly and maintaining a global view of available network resources is almost impossible.

To date, three applications of the ant metaphor in the domain of routing have been documented [White 97], [Schoonderwoerd et al 97] and [Di Caro and Dorigo 97]. Schoonderwoerd's work embraces routing in the circuit switched networks while Di Caro and Dorigo deal with packet switched networks. Di Caro and Dorigo, in particular, provide compelling experimental evidence, based upon on simulation, as to the utility of ant search in the network routing problem domain by comparing ant-based routing with the current and proposed routing schemes used in NSFNET. This paper describes the approach of [White 97]. This paper does not propose that this ant-based system be employed in networks as described but presents a model through which the interplay of model parameters can be explored, thereby increasing the utility and generality of the results obtained. For example, the use of a cost function for the network links rather than round trip delay measurements (as are used in Di Caro and Dorigo's experiments) allows more general statements to be made regarding the experimental results. Also, this system is more likely to be of use in a network planning or management context such as is required in a SONET transmission network rather than the control domain as suggested by Di Caro and Dorigo.

Ant Routing

There are three agent types in the Routing By Ants (RBA) system. These are explorers, allocators and deallocators. Each agent type possesses a small memory to store the route being traversed and the constraints on routing. Explorer agents exhibit the foraging behavior of ants and preferentially follow trails of pheromones laid down by previous explorers. Constraints on routing allow for the inclusion or exclusion of specific nodes on the route, or the exclusion or inclusion of certain links. The Allocator agents traverse the path determined by explorer agents and allocate network resources on the nodes and links used in the path. Similarly, when the path is no longer required, a deallocator agent traverses the path and deallocates the network resources used on the nodes and links.

The system works in the following way. A connection request is generated at a given node, the source. The connection request is either a point to point (P2P) or point to multi-point (P2MP) request. For P2P requests a new species of ant (agent) is created and m are sent out into the network. For a P2MP request with *n* destinations, *nm* agents of a new species are created and sent out into the network, with *m* being sent to each of the *n* destinations. These explorer agents execute the following pseudo-code algorithm:

1. Initialize the route finding simulation
   Set t:= 0
   For every edge (i,j) set an initial value $T_{ijk}(t)$
   of zero for trail intensity. Place m ants on the source node. {Create new explorers agents at a frequency e_f }
2. Set s:= 1 { tabu list index }
   for k:= 1 to m do
      Place starting node of the kth ant in $Tabu_k[s]$.
3. Repeat until destination reached:
   Set s := s + 1
   For k:=1 to m do
   Choose node j to move to with probability p_{ij}^k (t)
   Move the kth ant to node j.
   Update explorer route cost: $r_k = r_k + C(i,j)$
   If ($r_k > r_{max}$) then
      Kill $explorer_k$
      Insert node j in $Tabu_k[s]$.
      At destination go to 4.
4. While s > 1
   Traverse edge $Tabu_k[s]$.
   $T_{ijk}(t) = T_{ijk}(t) + ph_k$
   s := s - 1
5. At source node do:
   If the path in $Tabu_k$ is the same as p% of paths in

PathBuffer then Create and send an allocator agent
   If t > Tmax then
      Create and send an allocator agent

In the above algorithm, the following symbols are used:
- $T_{ij}(t)$ is the quantity of pheromone present on the link between the ith and jth nodes,
- C(i,j) is the cost associated with the link between the ith and jth nodes.

- r_k is the cost of the route for the kth explorer agent.
- $Tabu_k$ is the list of edges traversed.
- Tmax is the maximum time that is allowed for a path to emerge.
- PathBuffer is the array of paths obtained by the (up to m) explorer agents.
- r_{max} is the maximum allowed cost of a route.
- ph_k is the quantity of pheromone laid by the kth explorer agent.
- p_{ij}^k (t) is the probability that the kth agent will choose the edge from the ith to the jth node as its next hop.

The probability with which an explorer agent (k) chooses a node j to move to when currently at the ith node at time t is given by:

$$p_{ij}^k(t) = [T_{ijk}(t)]^{\alpha}[C(i,j)]^{-ß} / N_k$$

$$N_k = \Sigma_{j \text{ in } (S(i) - Tabu_k)} [T_{ijk}(t)]^{\alpha}[C(i,j)]^{-ß}$$

where α and β are control constants and determine the sensitivity of the search to pheromone concentration and link cost respectively. N_k is simply a normalization factor that makes $p_{ij}^k(t)$ a true probability. S(i) is the set of integers, {l} such that there exists a link between the ith and lth nodes.

Explorer agents are created at a given frequency e_f and continue to be created and explore the network during the lifetime of the connection. In this way it is possible to have recovery from node or link failure and (potentially) have the system re-route connections in order to overcome temporary congestion situations. An alternate solution to the node or link failure is to compute multiple node and link diverse paths from source to destination. This is currently being researched and will be reported in a future publication.

Pheromone levels also decrease with time by the process of evaporation. Left alone to decay, and with no reinforcement by other ants, the pheromone levels drop at a rate given by r. In our system the evaporation was simulated but in a real engineering solution this process might be effected by mobile evaporation agents traversing the network.

When explorer agents reach their destination they backtrack along the route chosen and drop pheromone in order to mark the path. Explorer agents that are searching for P2MP paths use the same pheromone and so tend to reinforce partial paths consequently identifying multi-cast nodes within the network. Constraint handling in the RBA system is important because of the inability of certain nodes in the network to perform multi-casting.

Upon arrival back at the source node, a decision is made whether or not to send an allocator agent. The decision is made based upon *m* previous allocator agents' paths. If *p* percent of the agents follow the same path, the path is said to have emerged. An allocator agent is then created and enters the network and allocates network resources along the route. In the case of P2MP allocator agents, the decision to create and send one is made based upon whether the spanning trees chosen are the same. It should be noted that P2MP explorer agents are searching for the lowest cost, spanning tree within the network for the source and destinations chosen and this problem is known to be NP hard. Another incidental but convenient property of the P2MP path search is that new connections can be added dynamically as remote sessions come online, as would typically be the case in a distance learning application. Potentially the entire spanning tree found by the P2MP agents might then change as a more efficient multi-cast solution is found.

Allocator agents traverse the paths indicated by the highest concentrations of the pheromones dropped by their associated explorer agents. P2MP allocator agents behave slightly differently when compared to their P2P counterparts in that they consume resources only *once* on a node or link during the allocation process.

It is possible that network resources have already been allocated by the time the allocator agent is sent. In this situation, the allocator agent backtracks to the source node rolling back resource allocation and decreasing pheromone levels such that a later explorer agent will not tend to follow this path. A decision to re-send an allocator agent is made after a back-off period has been observed. During the back off period explorer ants continue to search for alternative routes.

Low values of α indicate that the search process is insensitive to pheromone concentration, whereas low values of β indicate that link cost is unimportant. The balancing of these two parameters *strongly* affects the efficiency and stability of the search process.

Experimental Setup

Several networks were used during the evaluation of the RBA system. A small example is shown in figure 1.

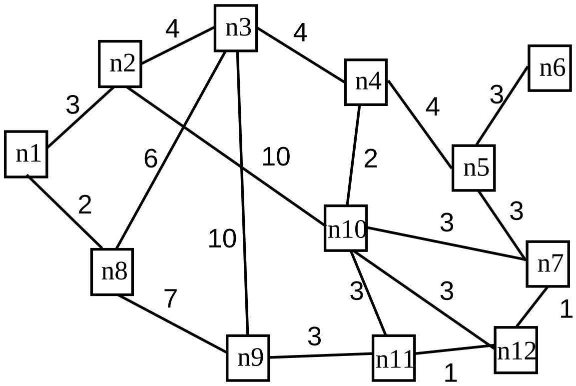

Figure 1: A Simple Example Network

The labels on the links represent available capacity for connections. Larger networks of up to 200 nodes were used in several experiments, these networks being examples of large transmission networks currently being planned.

In order to explore the sensitivity of the search algorithm, routes for all possible source-destination connection requests were computed for a range of α and β values. Routes were computed 100 times in order to determine the variation of search times. Values of α,β in the range [1,10] were used to investigate the sensitivity of the search to these parameters. Table 1 shows the values of principal parameters for the system.

Parameter	**Value**
Agent creation frequency	Every 10 cycles
Quantity of pheromone dropped	10 units
Emergence criterion	90% follow a path
Number of agents created	15
Path buffer size	40 agents
Pheromone evaporation rate	1.0 units/cycle
Maximum search time	300 cycles

Table 1: Experimental Parameters

Results and Discussion

When the algorithm starts, the cost part of the probability function is important, and the greedy heuristic comes into play. The actual value of α is unimportant, as only a small amount of pheromone (approaching zero) is present on the links, therefore the only factor influencing choice of links is actual cost on that link. As routes are found, pheromone is laid on the links that form that path. The amount of pheromone laid is inversely proportional to the total cost of the route, and acts as a global measure of ‘goodness’. This brings the reinforcement part of the probability function into play. The sensitivity to pheromone, α, influences the choice links, and links with more pheromone are more likely to be chosen.

α,β	2,1	2,2	4,2	8,2
Minimum	75	75	130	130
Maximum	400	220	195	190
Mean	220	175	159	150
Std Dev.	45	25	14	10

Table 2: Results

In the first column of the results table high standard deviations on search times implies that the confidence in a path is not high, and the system continues to use the greedy heuristic to explore other paths. The second and third columns of the results table show similar results for minimum, maximum and mean, but with standard deviation decreasing. The system finds results quickly, and reinforces good solutions. However, confidence is at a level such that other solutions are not rejected, and exploration continues. These seem to be the best sets of results. The final column in the results table show a fast, almost deterministic algorithm. This makes it undesirable for dynamic routing.

It is important that other solutions are explored and evaluated, as this is a stochastic approach. Reducing the standard deviation shows that the algorithm is more deterministic, which is not desirable. On the other hand, too high a standard deviation is also undesirable, as good solutions are not sufficiently reinforced.

Observations and simple analysis show that with high values of β, the system becomes ‘locked into’ the solution found first. This leads to problems later when bandwidth is removed from a link on that path. The system reorganizes, but when the bandwidth is reinstated, the previous solution, which is typically the best solution, is not found as the high sensitivity to pheromone is forcing the choice of paths where ants have previously been.

The times for route emergence varied considerably with the values chosen for α and β. Values of α=2 and β=2 were chosen as a reasonable compromise between exploitation of an emerging path and exploration of other paths in the network.

As a result of the sensitivity of the search process to parameter settings, the search process should probably be made adaptive. The values α and β should not fixed for the all agents but allowed to vary based upon the effectiveness of the search resulting from them. This enhancement to the basic system is currently undergoing investigation and will be reported in a subsequent publication.

Summary

This paper has described a search process that solves the routing problem for networks containing both point to point and point to multi-point connection requests. The process requires three agent types and is dynamic in nature, thereby allowing the potential for re-routing in situations where local congestion occurs. An interesting property of the process is the potential for reconfiguration of multi-cast path solutions as new sessions come online or existing sessions terminate. Results have shown that shortest path routes can be quickly computed and that response to failure events in the network is rapid.

Future Work

The system is currently being extended to allow for interactions between pheromone species. In the multi-pheromone ant colony system (mPAC), a chemistry $\mathfrak{C}$, is defined for the system:

$$c_i: s_1 S_1 + \ldots + s_m S_m \rightarrow s_1{}' S_1{}' + \ldots + s_m{}' S_m{}'$$

The chemical reactions can be defined link by link, globally defined and multiple reactions are possible $\mathfrak{C} = \{c_i\}$. Each reaction has an associated reaction rate (defined by Arrhenius' equation), i.e. temperature dependent. It is envisaged that the addition of such chemistry will provide for a mechanism to define behavior-based network management and control systems.

References

Beckers R., Deneuborg J.L and Goss S. (1992), Trails and U-turns in the Selection of a Path of the Ant Lasius Niger. In *J. theor. Biol*. Vol. 159, pp. 397-415.

Bullnheimer B., R.F. Hartl and C. Strauss (1997). Applying the Ant System to the Vehicle Routing Problem. *2nd Metaheuristics International Conference (MIC-97)*, Sophia-Antipolis, France.

Beni G., and Wang J., (1989) Swarm Intelligence in Cellular Robotic Systems, *Proceedings of the NATO Advanced Workshop on Robots and Biological Systems*, Il Ciocco, Tuscany, Italy.

Di Caro G. and Dorigo M. (1997) AntNet: A Mobile Agents Approach to Adaptive Routing. Tech. Rep. IRIDIA/97-12, Université Libre de Bruxelles, Belgium, 1997.

Costa D. and A. Hertz (1997), Ants Can Colour Graphs. *Journal of the Operational Research Society*, 48, 295-305.

Dorigo M., V. Maniezzo & A. Colorni, (1991) The Ant System: An Autocatalytic Optimizing Process. Technical Report No. 91-016, Politecnico di Milano, Italy.

Dijkstra E.W. (1959) A Note on Two Problems in Connexion with Graphs In *Numerische Mathmatik* vol. 1.

Franks N.R. (1989), Army Ants: A Collective Intelligence, *Scientific American*, Vol 77.

Fitzgerald T.D., Peterson S.C. (1988), Cooperative foraging and communication in caterpillars, *Bioscience*, 38, pp. 20-25.

Goss S., Beckers R., Deneubourg J.L., Aron S., Pasteels J.M. (1990), How Trail Laying and Trail Following Can Solve Foraging Problems for Ant Colonies, in Hughes R.N. (ed.) *NATO ASI Series, Vol. G 20, Behavioural Mechanisms of Food Selection*, Springer Verlag, Berlin.

Grassé P.P. (1959), La reconstruction du nid et les coordinations inter-individuelles chez Bellicoitermes natalenis et Cubitermes sp. La theorie de la stigmergie: Essai d'interpretation des termites constructeurs. In *Insect Societies*, Vol. 6, pp. 41-83.

Hölldobler B. and Wilson E.O. (1994), *Journey to the Ants*. Bellknap Press/Harvard University Press, 1994.

Langton, C.G. (1987) Artificial Life, *Proceedings of an Interdisciplinary Workshop on the Synthesis and Simulation of Living Things*, Los Alamos, New Mexico, Addison Wiley.

Maniezzo V., A. Colorni and M. Dorigo (1994), The Ant System Applied to the Quadratic Assignment Problem. Tech. Rep. IRIDIA/94-28, Université Libre de Bruxelles, Belgium.

Shapiro J. A. (1988), Bacteria as multi cellular organisms, *Scientific American*, pp. 82-89.

Schoonderwoerd R., O. Holland and J. Bruten (1997). Ant-like Agents for Load Balancing in Telecommunications Networks. *Proceedings of Agents '97*, Marina del Rey, CA, ACM Press pp. 209-216, 1997.

Schwarz M., and Stern T.E. (1989), *Routing Protocols, Computer Network Architectures and Protocols*. 2nd Edition. Sunshine, Carl A. (ed.). Plenum Press New York/London.

Taillard E. and L. M. Gambardella (1997), An Ant Approach for Structured Quadratic Assignment

Problems. *2nd Metaheuristics International Conference (MIC-97)*, Sophia-Antipolis, France.

Tanenbaum A.S. (1988) *Computer Networks*, Prentice Hall.

White T. (1997), Routing and Swarm Intelligence, Technical Report SCE-97-15, Systems and Computer Engineering, Carleton University, September 1997.