

Swarm Intelligence and Problem Solving in Telecommunications

Abstract

This paper describes how biologically-inspired agents can be used to solve control problems in Telecommunications. These agents, inspired by the foraging behaviour 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 telecommunication networks.

Introduction

The notion of complex collective behaviour emerging from the behaviour of many simple agents and their interactions is central to the ideas of Artificial Life. Nature provides us with many examples of social systems where individuals possess simple capabilities when compared to their collective behaviours which are much more complex. Such systems span several 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 behaviour and promises to generate approaches to distributed systems management 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 amount of data that must be transferred and processed as the network increases in size. 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 communications capacity. 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 behaviours of ants [Beckers et al 92] has led to new approaches to control in telecommunications networks.

Swarm Intelligence and the Ant Colony

Swarm Intelligence [Beni and Wang 89] is a property of systems of unintelligent agents exhibiting collectively intelligent behaviour. 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 behaviour 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 behaviour arises as a result of the actions of societies of such agents.

Individual ants are behaviourally 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 behaviour have been documented [Frank 89; Hölldobler and Wilson 94] including:

1. Forming bridges;
2. nest building and maintenance;
3. cooperating in carrying large items;
4. finding the shortest routes from the nest to a food source;
5. regulating nest temperature within a one degree celcius range;
6. preferentially exploiting the richest source of food available.

In the above examples two forms of stigmergy have been observed. Sematectonic stigmergy involves a change in the physical characteristics of the environment. 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 behaviour 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 signalling 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 behaviour which emerges is a form of autocatalytic behaviour 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.

Routing

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 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 bandwidth.

Ideally, we would like to have the network adapt routing patterns to take advantage of free resources and move existing traffic if possible. To date, two applications of the ant metaphor for routing have been documented [White 96], [Schoonderwoerd 96]. This paper describes the approach of [White 96].

Ant Routing

The "Routing by ants" systems consists of three agent types; explorers, allocators and deallocators. Explorer agents exhibit the foraging behaviour of ants and preferentially follow trails of pheromones laid down by previous explorers. Allocator agents traverse the path determined by explorer agents and allocate the bandwidth on the links used in the path. Similarly, when the path is no longer required deallocator agents traverse the path and deallocate the bandwidth used on the links.

The system works in the following way. A call request arrives at a given node. The call is either a point to point

(P2P) or point to multi-point (P2M) request. For P2P requests a new species of ant (agent) is created and sent out into the network. For a P2M request n agents of a new species are created and sent out into the network. These explorer agents execute the following algorithm:

1. Initialize
 - set $t := 0$
 - For every edge (i,j) set an initial value $T_{ij}(t)$ for trail intensity. Place m ants on the source node. [Generate new explorers at freq. e_{fj}]
2. Set $s := 1$ { tabu list index}
 - for $k := 1$ to m do
 - Place starting node of the k th ant in $\text{tabu}_k(s)$.
3. Repeat until dest'n reached:
 - Set $s := s + 1$
 - for $k := 1$ to m do
 - Choose node j to move to with prob $p_{ij}^k(t)$
 - Move the k th ant to node j .
 - Update explorer route cost: $r_k = r_k + c(i,j)$
 - if $(r_k > r_{\max})$
 - kill explorer r_k
 - Insert node j in $\text{tabu}_k(s)$.
 - At destination go to 4.
4. While $s > 1$
 - traverse edge (i,j)
 - $T(i,j) = T(i,j) + p_e$
 - $s := s - 1$
5. At source node do:
 - if $(\text{path}_e = \text{pathBuffer} * d)$
 - create and send allocator agent
 - if $t > T_{\max}$
 - create and send allocator agent

Explorer agents are created at a given frequency e_f and continue to be created and explore during the lifetime of the call. In this way it is possible to have recovery from node or link failure and (potentially) have the system re-route calls in order to overcome temporary congestion situations.

When explorer agents reach their destination they backtrack along the route chosen and drop pheromone in order to mark the path. 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 $p\%$ of the agents follow the same path, the path is said to have emerged and an allocator agent is created and enters the network

in order to allocate bandwidth. In the case of P2M allocator agents, the decision is made based upon whether the spanning trees chosen are the same. A convenient property of the P2M 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 P2M agents might then change as a more efficient multi-cast solution is found.

Allocator agents traverse the path indicated by the highest concentrations of the pheromones dropped by their associated explorer agents. It is possible that network bandwidth has already been allocated by the time the allocator agent is sent and in this case the allocator agent backtracks to the source node rolling back resource allocation and decreases pheromone levels such that a later . A decision to re-send an allocator agent is made at a later time after a back-off period has been observed. During the back-off period explorer ants continue to search for routes.

The probability with which an explorer agent (k) chooses a node j to move to is given by:

$$p_{ij}^k(t) = \frac{[T_{ij}(t)]^\alpha [C(i,j)]^\beta}{N_k}$$

$$N_k = \sum_{k \text{ in } (S\text{-Tabu}(k))} [T_{ij}(t)]^\alpha [C(i,j)]^\beta$$

where α and β are control constants and determine the sensitivity of the search to pheromone concentration and link cost respectively. N_k is a normalization term that

makes $p_{ij}^k(t)$ a true probability. 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 of the search process.

As a result of the sensitivity of the search process to parameter settings, the search process was made adaptive. The values α and β were not fixed for the all ants but allowed to vary based upon the effectiveness of the search resulting from them. In order to achieve this each ant was given an associated fitness value – the cost of the path found – and whenever new ants were created α and β were determined by proportional selection of two parent ants based upon fitness followed by crossover and mutation.




figure 1. An example of crossover

The figure above shows the genotype of two ant parents that encode α and β . A single crossover point is chosen and two offspring are generated. One offspring is discarded and the other undergoes mutation as shown below.




figure 2. An example of mutation


As can be seen from the two figures, the resulting offspring is rarely the same as either of the parent. Using this approach significantly improved the search process as the agents learned efficient search parameters.

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 call 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 \mathcal{C} , is defined for the system:



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

References

[Beckers et al 92] Beckers R., Deneuborg J.L and Goss S., *Trails and U-turns in the Selection of a Path of the Ant Lasius Niger*. In J. theor. Biol. Vol. 159, pp. 397-415.

[Beni and Wang 89] Beni G., and Wang J., *Swarm Intelligence in Cellular Robotic Systems*, Proceedings of the NATO Advanced Workshop on Robots and Biological Systems, Il Ciocco, Tuscany, Italy, June 26-30, 1989.

[Franks 89] Franks N.R., *Army Ants: A Collective Intelligence*, Scientific American, Vol 77, March-April, 1989.

[Fitzgerald and Peterson 88] Fitzgerald T.D., Peterson S.C., *Cooperative foraging and communication in caterpillars*, Bioscience, 38, 1988, pp. 20-25.

[Goss et al 90] Goss S., Beckers R., Deneubourg J.L., Aron S., Pasteels J.M., *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, 1990.

[Grassé 59] Grassé P.P., *La reconstruction du nid et les coordinations inter-individuelles chez Bellicositermes natalensis et Cubitermes sp. La theorie de la stigmergie: Essai d'interpretation des termites constructeurs*. In *Insect Societies*, Vol. 6, pp. 41-83.

[Hölldobler and Wilson 94] Hölldobler B. and Wilson E.O., *Journey to the Ants*. Bellknap Press/Harvard University Press, 1994.

[Shapiro 88] Shapiro J. A., *Bacteria as multi cellular organisms*, Scientific American, 1988 pp. 82-89.

[White 96] White A.R.P., *Routing with Ants*, Nortel internal report, 1996.

[Schoonderwoerd 96] Schoonderwoerd, R., Holland, O.E., Bruten, J.L., & Rothkrantz, L.J.M. (1996) *Ant-based load balancing in telecommunications networks*, HP Technical Report HPL-96-76.