Matthew Guidry

Artificial Intelligence in Networking: Ant Colony Optimization

Abstract

Ever since the internet became a must have in today's technological world people have been looking for faster and faster ways to connect one machine to another. Many eloquent techniques have been proposed for this problem, some that are highly effective in individual cases. However, when it comes to the internet as a whole and connecting the mass of people to their individual destinations, an efficient algorithm is more difficult to produce.

One proposed technique is to use Artificial Intelligence at the global router level to help different computers use the most efficient route between each other. One direction that researchers have gone to pursue this is to study the behaviors of ants for their techniques to find the fastest path between two points. They also analyze their problem solving techniques when the path that they are following has been broken. Computer Scientists began researching the behavior of ants in the early 1990's to find new routing algorithms. The result of these studies called is Ant Colony Optimization (ACO) and in the case of well implemented ACO techniques, optimal performance is comparative to existing top-performing routing algorithms. [2] This research has yielded ways to minimize the number of nodes that are taken to get to the destination, techniques for quickly resolving an efficient path, and ways to avoid having loops within a routing scheme.

Problems with Routing in the Internet

Ants have the ability to create long bi-directional paths between their homes to their goals in optimal time. These paths may develop obstructions which cannot be crossed and the path may need to be changed at any point in time. It is this ability of the ants which interests researchers the most. By mimicking these abilities with routing algorithms researchers believe they can help the internet as a whole flow faster. At any given moment there are routers and other routing devices which are taken down for an assortment of reasons. These devices could have been vital nodes for a path given by another router. It is important that the routers which were previous in the path be able to find a new efficient path quickly. Ant Colony Optimization is aimed at doing just that.

Packet Switching and Circuit Switching

There are two different techniques currently used by computers to send information across the global internet, circuit switching and packet switching. Circuit switching is often compared to a telephone call because it follows the same basic principles. The phone is picked up, a call is made to a receiver which is notified, the receiver receives the call and can only talk to one client at a time, and finally when the call is over both the caller and the receiver hang up the phone and end the call. This is the same basic procedure as circuit switching, once a connection has been made between two points that connection is held open for just those points until the call is over. All of the packets of data follow the same path to their destination and they arrive in the same order as they are sent. For the router to be able to set up this kind of connection it must have the knowledge of the entire topology of the network. This requires a lot of information to be stored on the actual router level. The router must work with this data to formulate the best path to send the packets on to the receiver.

Packet switching, however, is much less organized. The packets which are sent from one computer to another do not necessarily follow the same path or arrive in the same order. At each different hop that a packet encounters in the network, that router decides which path to send it to next. These two techniques force different requirements to be in place at the router level. Packet switching requires much less state to be stored at the router level than circuit switching, routers only need to know who their neighbors are and how much cost is associated with getting each of them. The packets get routed at each hop and it is not needed for the sending router to know how packets will get to their destination. This scheme does not guarantee that the packets will be received in the same order they were send, and the receiver must be able to handle this occurrence.

After a transmission has been started, more efficient paths may become usable. It may be because routers along the path become bogged down, it may be because some router is taken down, or a number of other reasons that can change the efficiency of a path. This is described by "good news travels slow" which is discussed later in the paper. By replacing these algorithms with ACO, which implements Artificial Intelligence, packets can be transmitted more efficiently.

Developing Ants to Travel the Network

In the real world ants will search aimlessly until they find food, once they do they will return to the colony the fastest way they know how and they will mark their path with a pheromone trail. Over time this pheromone trail will begin to evaporate if it is not followed by other ants which keep it strong by relaying the pheromone. Furthermore as ants begin to follow the trail to the food source, if they discover better short cuts by following another route, they will lay their own pheromone trail of their routes. Over time the following ants will take these paths and make that scent stronger, and the old trail's scent will evaporate. In this way pheromone evaporation allows ants to enhance their paths and converge to find the most optimal route to the goal. [5]

Researchers tried to mimic this artificially with their own version of the pheromone trail. Intelligence in an ant algorithm is achieved by using this "artificial pheromone trail", this trail is essential to the metaheuristics. This trail is updated by probabilities which are increased whenever a successful path is achieved. Whenever a packet comes to a point where it can proceed down multiple paths it picks the next step based upon shared learning of previous packets and a greedy heuristic designed for routing.[3] This trail is developed and maintained by two different types of ants, the regular ants and the uniform ants.

In ACO it is the regular ants which transport packets from their origin to destination in the most efficient manner. At every node in the path these ants look at the probabilistic routing tables, which have been set up initially by the uniform ants, and determine which next hop to take as it progresses to the destination. When a route is taken more often and is proved to be the most efficient route, its probability of being taken again will be increased. This will make all the regular ants converge on to that route which is then taken by almost all of the packets. When this state is achieved the ants are said to be stable. This is how the technique finds the fastest path through the topology. The regular ants do not use as much intelligence as the uniform ants, which are used to discover the fastest routes through the network.

Uniform ants are dispatched to travel the network and seed the probabilistic routing tables so that the worker ants can use their heuristics to choose the best path. These ants do not necessarily need a destination because they are just searching to find the fastest paths to different nodes. This is also important because the origin node may not have knowledge of the entire network layout. Uniform ants are also immune to the probabilities which have been previously set at each router to guide the regular ants; this makes them immune to the oscillation problems in the network. These ants use backwards reporting and once they reach a certain destination will report back to the previous node the cost that was encountered from that hop, the routers then update their probabilistic routing tables. These ants also use their own heuristics to determine which route to take next, which ensures that not all of the uniform ants take same path at a given node. Every path has the same probability of being taken. They also have an attribute associated with them called time to live, at every node that the worker ant passes a variable is increased by one and once this hits a certain limit the ant dies. This ensures that the ant will not be crawling around the network forever.

Bad News Travels Fast, Good News Travels Slow

An inherent trait in networking which affects ACO and other routing algorithms is that bad news travels fast and good news travels slow. If a router suddenly goes down the ants are trained to deal with this situation and do so quickly. As the reports of ants reaching their destination stop coming back to the router, the next ants will choose the second best path and the best route will be altered quickly. However, if for some reason the best route which the ants are all choose slows down; for instance if that router has a high load transferring through it. The algorithm does not go to the second best route as quickly.

The effectiveness with which the algorithm deals with the possible changes in the topology or link cost changes is called its adaptiveness. If one of these changes occur on a path once the regular ants converge upon it as the best path, their policy will prevent them from changing to a new path as quickly in accordance with their probabilistic routing table. This transition is sped up by the learning rate of the algorithm and the worker ants which are not affected by the routing tables. Once the second path has been traveled by a few regular ants and there are no ants which reinforce the effectiveness of the old best path, the new path quickly replaces the old one. The probabilistic routing table is now being changed by every successful trip that is made to the second best route.

Strengths of Ant Colony Optimizations

One important advantage from using ACO is that these ants have a constant size in bits. This size is small enough to be piggybacked on top of other packets that need to be sent along this path. This means that these messages cost little to no extra overhead, unlike circuit switching and packet switching. For circuit switching every time there is an update in the network, it has to be propagated to every router so that they can keep an accurate memory of what the topology is. In the instance of packet switching periodically routers must send messages to all of their neighbors and vice-versa so that each will know the costs associated with sending packets to their neighbors. This optimization could prove to be very useful for networks with routers having bandwidth issues. For every algorithm the fact that the network topology is ever changing is one of the biggest hindrances. The changes in the topology must be accounted for in one way or another.

Uses for Ant Colony Optimization

Ant Colony Optimization is not just useful for computer networks as discussed in this paper. Forms of this algorithm have also been applied for the Traveling Salesman Problem. This algorithm is also being researched at MIT in an effort to steer robotic cars through a busy city. At the root level, determining the best way to get a person from their current location to their destination has a lot of the same difficulties as routine packets. ACO has also been applied to areas in assignment, scheduling, subset, machine learning, and bioinformatics problems. [4]

Problems with Ant Colony Optimizations

There are a few reasons why this revolutionary new technique for handling complex trees has not been more featured throughout the internet. The main reason is the cost of change. The internet has come out with countless upgrades including already upgrading to IPV6 which has a long way to go before it has been fully integrated. It is very difficult to upgrade routers across the board to handle new routing techniques. This is a major hassle for ACO to be widely integrated across the internet at major routers, however, individual networks can do this.

Conclusion

Ant Colony Optimization has been researched for and applied to some core areas of Computer Science. One main area of interest is networking; because of ant's natural ability to go from their starting point to their destination with the most limited amount of cost. We are also interested in their quick problem solving techniques when their path has been blocked for one reason or another. Ant Colony Optimization is one of the latest possible solutions to be applied to networking and because of its potential effectiveness it has raised interest in other related fields as well. Although network administrators are reluctant to upgrade from their system which is already in place, this eloquent new approach has gotten a lot of attention and could definitely be applied in future releases.

Another area of research within Ant Colony Optimization which could be expanded upon is the tendency of ants to stay within safer areas until they reach a certain point at which their target is a straight shot. If developers could figure a way to apply a sort of security heuristic this could be applied to routers which have had a tendency to go down in the past, or routers which repeatedly become bogged down once they become an optimal next hop in the topology. If this safety heuristic was added it may even enhance the effectiveness of this algorithm against man in the middle attacks where a certain node is suspected or actually discovered to be inspecting all of the packets which travel through it.

References

[1] Ants and reinforcement learning: A case study in routing in dynamic networks (1997) by Devika Subramanian,Peter Druschel,Johnny Chen Proceedings of the Fifteenth International Joint Conf. on Arti Intelligence

[2] Webpage: " http://www.codeproject.com/KB/recipes/Ant_Colony_Optimisation.aspx", Lawrence Botley, 2008

[3] Webpage: " http://www.sciencedirect.com/science?_ob=ArticleURL&_udi=B6VCT-4S5FJCY-4&_user=2139768&_rdoc=1&_fmt=&_orig=search&_sort=d&view=c&_acct=C000054272&_version=1&_urlVersion=0&_userid=2139768&md5=f683a68f58b4d6e5eeaf682b76f7d6ed ", Sara Morin, Caroline Gagné, and Marc Gravel, 2008

[4] "Ant Colony Optimization ", Marco Dorigo and Thomas Stützle

[5] Webpage: "http://en.wikipedia.org/wiki/Ant_colony_optimization "