How Would Ants Implement an Intelligent Route Control System?

Hamid Hajabdolali Bazzaz[1,2] and Ahmad Khonsari[1,2]

[1] ECE Department, University of Tehran, North Kargar Ave., Tehran, Iran
[2] School of Computer Science, IPM, Niavaran Sq., Niavaran, Tehran, Iran
h.hajabdolali@ece.ut.ac.ir, ak@ipm.ir

Abstract. Multihoming, the connection of a stub network through multiple Internet Service Providers (ISPs) to the Internet, has broadly been employed by enterprise networks as a sort of redundancy technique to augment the availability and reliability of their Internet access. Recently, with the emergence of Intelligent Route Control (IRC) products, IRC-capable multihomed networks dynamically select which ISPs' link to use for different destinations in their traffic in a smart way to bypass congested or long paths as well as Internet outages. This dynamic traffic switch between upstream ISPs is mostly driven by regular measurement of performance metrics such as delay, loss ratio, and available bandwidth of existing upstream paths. However, since IRC systems are commercial products, details of their technical implementation are not available yet. Having the incentive to delve into these systems deeply, in this paper, we employ traditional ant colony optimization (ACO) paradigm to study IRC systems in that domain. Specifically, we are interested in two major questions. Firstly, how much effectively does an ant based IRC system switch between upstream links in comparison to a commercial IRC system? Secondly, what are the realistic underlying performance metrics by which ants pick the path to a food source (destination network) in a multihomed colony? Through extensive simulations under different traffic load and link reliability scenarios, we observe that ants perform well in switching between available egress links. Moreover, delay of paths is not the only criterion by which ants select the path; instead, through their intuitive ACO paradigm, they tend to choose the path with a better performance in terms of both delay and loss ratio.

Keywords: Intelligent route control, ant colony optimization.

1 Introduction

Multihoming has been widely used as a redundancy technique to provide stub enterprise networks with a higher level of availability and reliability in their Internet access. In this practice, each stub network is connected to the Internet through several Internet Service Provider (ISP) upstream links rather than a single one. In the most straightforward deployment configuration, one of the ISPs is picked as the primary Internet provider while the others are just used as backup providers and only in the case of failures in the primary connection (Fig. 1).

K. Al-Begain, D. Fiems, and G. Horváth (Eds.): ASMTA 2009, LNCS 5513, pp. 29–41, 2009.

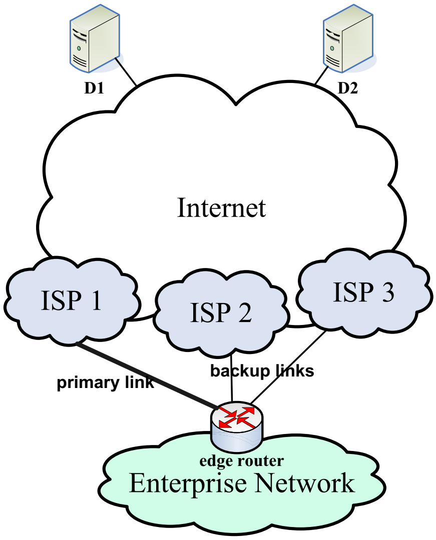

Fig. 1. A multihomed network with three upstream ISPs

Recently by the emergence of Intelligent Route Control (IRC) products [1]-[3], instead of using only one of the upstream links at each time, all egress links are utilized and the IRC edge router dynamically selects the best upstream ISP for every major destination of the network traffic. This is typically done driven by some performance criterion such as delay and loss ratio. As a result, the proliferation of IRC systems by stub networks is not only for bypassing Internet outage but also to have their traffic experience a low latency and packet loss.

There are just few recent studies on the potential performance/cost benefits of IRC systems [4]-[7]. However, since IRC systems are commercial products, only the basic implementation guidelines are discussed thus far. In this paper, we apply traditional ant colony optimization (ACO) paradigm to study these systems deeply. Because, as we observe throughout the paper, the way ants select their path to a food source using their heuristic pheromone-basis ACO paradigm is very similar to commercial IRC systems deployment. We translate the IRC problem into ant colony domain where ants are seeking the best egress path between their multihomed colony and source food destinations. Specifically, we are interested in two primary questions. Firstly, how much effectively does an ant based IRC deployment pick upstream links in comparison to a commercial IRC system? Secondly, what are the realistic underlying performance criterions by which ants select the path to a destination (food source) in a multihomed colony? We address these questions throughout this paper.

The rest of the paper is structured as follows. Section 2 reviews the commercial IRC implementation guidelines from the literature. In Section 3, we translate the IRC problem into ant colony domain and argue about different design decisions we need to make. Performance, convergence time and criterion metric of the proposed ant based IRC system are evaluated in Section 4 and finally we conclude the paper in Section 5.

2 Commercial IRC Systems

There are few works on performance evaluation of IRC systems such as [4]-[11]. More recently, in [5][6] Akella *et al.* emulate an IRC system on Akamai content distribution network and argue that multihoming can bring about a considerable amount of benefits in terms of both availability and performance. As for the degree of multihoming, their results state that IRC has the potential to achieve an average performance benefits improvement of 25% or more for a 2-multihomed stub network. The authors also show that typically having up to four upstream providers is enough to gain the full benefit out of multihoming.

In previous works [12], an IRC system behavior is typically modeled as a periodic five-stage process, namely, idle, measurement, performance estimation, routing decision, and path switching. The whole period is T_R seconds. There is an optional idle stage in the beginning of every cycle in which the system doesn't probe the destinations to reduce the overhead of the routing probes. Then, in the measurement stage which is T_M seconds, different metrics of the existing upstream paths (e.g., delay, loss ratio and available bandwidth) are collected through sending N_p number of *probing* packets. After receiving the acknowledgements of these probes, the IRC system estimates a level of performance for each candidate path by a hybrid metric typically consisting of delay and loss ratio of the paths. In the routing decision stage, one of the paths is selected as the best upstream path based on the previous stage performance estimations. And finally, the IRC system routes the traffic through the chosen egress link at least during the entire coming routing period. Notice that performance estimation, routing decision, and path switching stages are almost immediate tasks because they only include simple calculation and forwarding table (memory) update. Our proposed deployment of the IRC system follows this basic 5-stage process. We discuss the details of the ant based approach in the following Section.

3 Ant Based IRC System

Ant Colony Optimization (ACO) initially was proposed by Marco Dorigo in 1992 in his PhD thesis [13]. The aim of the first algorithm was searching for an optimal path in a graph based on the behavior of ants seeking a path between their colony and a food source [14]-[16]. Intuitively, ants do so by wandering around randomly on different paths while laying down *pheromone* trail, a chemical substance they use to form an indirect communication with each other. As soon as an ant finds a path to a source food, it comes back to the colony while still putting down pheromone on its way back to the home colony. Ants greedily believe that paths which are traveled more (by their companions) and, thus have more pheromone are more likely to be leading to source food. Thus, while an ant meanders randomly to find a short path to a food source, the chance it follows a specific existing path is somehow related to the amount of pheromone its senses on that specific path. Over time, however, the pheromone trail evaporates [17]. The more time it takes for an ant to travel down a specific path and back to the home again, the more time the pheromones have to evaporate. Yet, a short path gets marched over faster, and thus its pheromone density remains high even though its pheromones evaporate. As a result, the problem of finding the best path (e.g., shortest one) which the ants are trying to solve naturally is indeed heuristically solved by this simple approach they follow. Thus far imitating ants' behavior, diverse numerical problems are solved by employing ACO. See [15] for a complete survey.

In this section, after this brief overview on ACO, we formulate the IRC problem as an ACO problem. Specifically, we are to denote the notations of ACO we employ in a typical IRC system, namely:

(3.1) What are the correspondent elements of IRC in our ant based model?
(3.2) How much pheromone do ants lie down on the path they travel? Is this amount constant or it may change dynamically?
(3.3) How long does it take for pheromones to evaporate? In other words, what is the pheromone evaporation function?
(3.4) How exactly do ants select the path they march on?

In the coming subsections we address these questions.

3.1 Ant Colony Model

We consider a simple multihomed network with an IRC capable edge router which major traffic is destined to several probably big networks which are given a priori. In our model, the source IRC-equipped stub network represents the ant colony while each of its major destinations is like a food source to which ants are trying to find the best path (Fig. 2). In the ant based IRC system, ants play the role of the routing probes the edge router periodically transmits through its upstream links in order to obtain performance metrics of each of its paths to a specific destination. Both systems (ant based IRC and commercial one) have the same objective; calculating the best path to

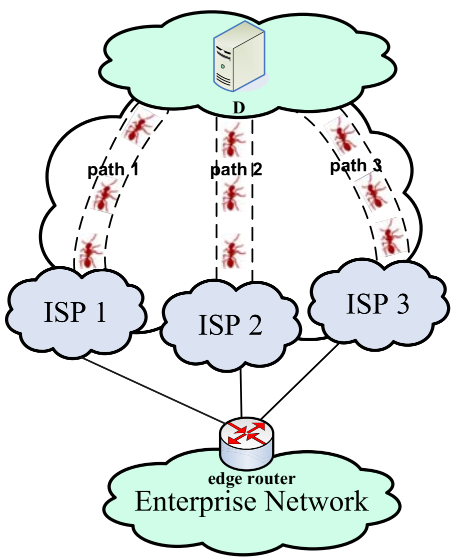

Fig. 2. An ant colony IRC system

the major destinations of the stub network. There are actually two choices for the rate by which new ants march from the source to each destination. First, we may consider this rate as the traffic rate between the stub network and that destination. The alternative is to map this rate only to that of probing packets. Indeed, there are no significant differences between these two schemes, but as the principle purpose of the IRC system is finding out the best path per destination through probing, we take the second option. As a result, in our model the rate of data packets between the source and destination is not of importance given the fact that solely probing ants are trying to calculate the best path. This roughly means that we have two sorts of ants: the elite type (routing probes packets) which generates pheromone while wandering and the normal ants (data packets) which solely traverse the best path picked by their elite companions. Unless otherwise specified, we are referring to the elite ants simply by ants throughout the paper.

3.2 Pheromone Creation Formulation

As the ants march on different upstream links, they lay down pheromone trails on the path they travel. Eventually, this would result in the augmentation of the pheromone value associated with good paths (i.e., shorter paths) to encourage more ants to pick them, and decrease in that of bad paths (i.e., longer paths) through pheromone evaporation. Here, we formulate the pheromone update scheme we use throughout this paper by adapting the *aging* phenomena in the ACO [14]. Aging phenomena simply states that an older ant may produce less pheromone as a result of its decrepitude. It is one of the methods used to control the amount of pheromone of the whole system in the ACO. Thus, by adopting this fact, we assume that (i) the initial amount of pheromone an ant lies down on the paths it travels on is a constant value K_0; (ii) ant's pheromone generation rate decreases exponentially while it becomes older as the result of the aging phenomena. Specifically, the amount of pheromone an ant puts on a path at the age of t (seconds) is $K_0 e^{-t}$. Let $\tau_{s,i}(.)$ represent the amount of pheromone at egress point of path i of the enterprise source node s. Thus, if at time t_0 an ant a with the age of $age(a)$ marches on the egress point of path i of the IRC equipped enterprise network s, we have:

$$\tau_{s,i}(t_0 + \varepsilon) = \tau_{s,i}(t_0) + K_0 e^{-age(a)} \qquad (1)$$

in which $0 < \varepsilon << 1$ is a very small number.

3.3 Evaporation Rate Formulation

As explained earlier, to reduce the impact of past experience, an approach called evaporation [17] is typically applied in ACO. Evaporation avoids pheromone concentration in optimal paths from being excessively high and preventing ants from exploring other (new or better) alternatives [14]. Furthermore, with employing evaporation it is possible to switch between paths in case of dynamic changes in the optimal solution of the problem. Indeed, this is exactly what happens in the ant based IRC and we discuss this more later. Here, we formulate the pheromone evaporation paradigm we use in our ant based IRC. We assume that the initial K_0 units of pheromones an ant put down on the path are evaporated T_E seconds later. This roughly results in a pheromone evaporation rate of $\frac{K_0}{T_E}$.

Now, by considering aforementioned pheromone generation and evaporation formulas, we can analytically calculate the pheromone update function $\tau_{s,i}(.)$, recursively by adding up the pheromone value changes during interval $T = [t, t + \Delta t]$ on each path. Specifically, we have:

$$\tau_{s,i}(t+\Delta t) = [\tau_{s,i}(t)(1-\Delta t\frac{K_0}{T_E})]^+ + \sum_{a \in i \,\mathrm{in}\, T} [K_0 e^{-age(a)}(1-(\Delta age(a))\frac{K_0}{T_E})]^+ \quad (2)$$

where $[z]^+ = \max\{0, z\}$. In this formula, the first term corresponds to the residual amount of the initial value of pheromone on the path at time t after decreasing evaporated amount out of it. The second term is the residual amount of pheromone added to the path as the result of ants walking on it during the interval T. Note that the summarization is over all the ants marching on the egress point of path i in the time period T. $age(a)$ represents the age of ant a exactly at the moment it is on the path and $\Delta age(a)$ denotes its age increase from the initial time it is on the egress point till $t + \Delta t$.

3.4 Ants Routing Strategy

In this subsection we bring up the routing strategy ants utilize to reach source foods and then travel back to their colony. We distinguish between three cases as follow:

1) Egress upstream link: Ants select the egress upstream link attached to their source colony in a probabilistic nature based on the amount of pheromone they sense at each egress point at time they want to wander out of their colony. Specifically, if we denote the probability that an ant selects the egress upstream link i attached to the source colony s at time t by $p_{s,i}(t)$, we have:

$$p_{s,i}(t) = \frac{\tau_{s,i}(t)}{\sum_i \tau_{s,i}(t)}. \quad (3)$$

2) Path from a specific egress link to a destination food source: Once an ant has chosen the egress upstream link, it has to find out the path to the source food through that link. There are two choices for this path selection strategy. Firstly, to do so similar to egress link selection in a probabilistic nature through some ant based optimization schemes (e.g., following the pheromone trails). Secondly, to utilize the current available forwarding tables of each intermediate node to reach the destination. We stress that in an ant based IRC system, the objective is solely picking the egress link to a specific destination rather than the whole path which is already plain with the aid of underlying routing algorithms. Thus we select the second alternative. One may note that the first choice would work as well with some additional overhead in the convergence time of the ACO. In other words, here, we

assume that ants have already solved the traditional problem of finding the best path from each node in the given network to the source foods and our focus is on IRC nature of the problem.

3) Path from the destination to the source colony: Once an ant reaches a source food, it can simply travel back on the path it has already found between the source and destination to get back to its home colony. As a matter of fact, this travel back methodology has been traditionally used in ACO [13].

4 Experimental Results

This section presents a performance evaluation of our proposed ant colony based IRC system. Specifically, we are interested in two aforementioned questions. Firstly, how much effectively does the ant based system switch between paths in comparison to a commercial IRC system? Secondly, what are the underlying performance criterions by which ants switch between paths?

4.1 Simulation Setup

Our simulations are performed with a customized simulator, implemented in C++, which precisely model the network topology and traffic at packet level. Unless otherwise specified, we set simulation parameters like this: $K_0 = 1$, $T_E = 3$, $T_R = 2$, $T_M = 0.8$, and N_p, the number of probing packets in each measurement period, is selected 30. We have conducted extensive simulations to study effects of each of these key parameters on the results and then selected aforementioned values for each parameter. In subsection 4.2, we review how we derived these values for the simulation parameters.

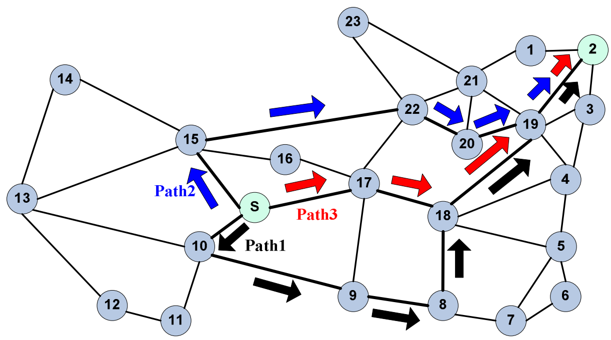

Fig. 3. The simulated network

The test network used in our experiments consists of a 23-node 38-link network [19] depicts in Fig. 3 All links are 1 MB/s with the same routing metric cost. The background traffic consists of a number of random flows between nodes in the topology. The node pairs, start time, duration, and rate of these flows are selected randomly according to a uniform distribution. In the simulated network, source node S is multihomed by three different upstream links, namely to nodes 10, 15, and 17. We pick node 2 as the major destination (food source) of this node due to the appropriate path diversity of selected upstream nodes to this destination. In particular, given the same cost for every link in the topology, any shortest-path routing algorithm (such as OSPF) picks the following paths as the shortest path between aforementioned upstream nodes to node 2:

Path1: <S, 10, 9, 8, 18, 19, 2>
Path2: <S, 15, 22, 20, 19, 2>
Path3: <S, 17, 18, 19, 2>

Here, we have chosen the node with lower id in case of any ties. For instance in the Path1, between two available paths with the same cost from intermediate node 22 to node 2 (<22, 20, 19, 2> and <22, 21, 1, 2>), the path through node 20 is selected rather than that of node 21.

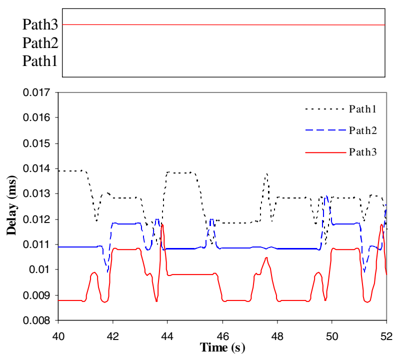

Fig. 4. Comparison of delay of paths at different time (bottom graph) along with chosen path (top graph) in the commercial IRC in scenario 1

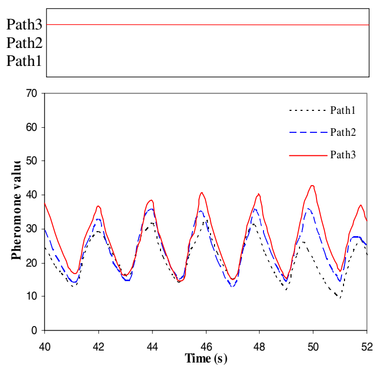

Fig. 5. Comparison of pheromone of paths at different time (bottom graph) along with chosen path (top graph) in the ant based IRC in scenario 1

We explore the behavior of the ant based system and compare its effectiveness with a commercial delay based IRC system, i.e., an IRC system which path switching decisions are made solely on the basis of delay metric of the paths, in three different scenarios. In the first scenario, the underlying network traffic load on the paths is low and the paths are not congested as well as reliable and stationary. In other words, there are no temporal congestion, or delay oscillation and bursty packet losses. In the second scenario, we put a temporal congestion event on the Path1 which brings about a longer delay on this path. We then study the reaction of the systems to this delay increase. In the final scenario, we examine the impacts of unreliable links, e.g., lossy ones in a wireless like environment, on the efficiency of the systems.

4.2 Performance Evaluation under Stationary Conditions

Assuming the same propagation delay for all the links in the simulated network, Path3 has the lowest minimum round trip time (RTT) due to having minimum number of intermediate nodes between source node S and destination node 2; after that comes Path2 and finally Path1 which has the longest native delay. In this section, we examine the stationary scenario in which the underlying network paths are under low load and they are stable in terms of delay, packet loss and congestion.

Fig. 4 depicts the delay of the paths during a part of simulation time interval. The top graph in this figure shows corresponding chosen path of the delay based IRC system in each interval. As the switching decisions are made exactly every T_R seconds (2 sec in this case) based on the average measured delay of the latest measurement interval T_M (0.8 sec in this case), at each point the nearest measured delay is shown in the graph. This explains, the reason the curves are straight for an interval of T_R - T_M (1.2 seconds in this case) every T_R seconds. According to the stable traffic load of this scenario, Path3 always has the lowest delay on average; Path2 comes next while Path1 has the longest delay. Thus, the delay based IRC system has always selected egress link of Path3 (the top graph) in this scenario.

On the other hand, Fig. 5 shows the ant based IRC system path switching events along with correspondent switching metric, i.e., amount of pheromone on each egress link. As it is seen in this figure, the system works effectively and similar to the commercial delay based system in terms of selecting the best upstream link during every iteration.

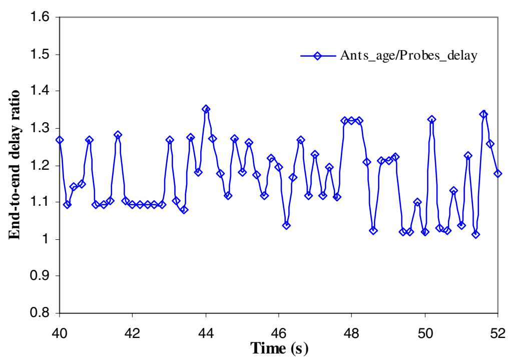

Fig. 6. Comparison of ants and routing probes end-to-end delay at different time for Path1 in scenario 1

Another observation to make is regarding the ant age in the ant based IRC system and its correspondent in the commercial IRC system, i.e., packets end-to-end delay. Fig. 6 shows the ratio of this end-to-end delay measurement of the ant based system to that of the commercial IRC system for Path1. As it is seen in this graph, ants always have greater end-to-end delay than IRC probing packets. The reason behind this is that although both ants and routing probes follow the same path to the destination, yet on traveling back to the source network, routing probes come back on the shortest path between the destination and source network (Path3 in this case) while ants travel back on exactly the same path (Path1 in this case) they have found from their source colony to the destination source food. This clearly led to a higher end-to-end delay for the ants due to the larger delay ants experience on their path from destination to their home.

A discussion on the selection of the system key parameters is due here. First, note that K_0, the initial amount of pheromone ants lay down while marching around, doesn't have a critical impact on the path selection. Indeed, increasing/decreasing this parameter would increase/decrease the total amount of pheromone per iteration yet switching decisions are not altered. However, evaporation duration time, T_E, has a significant effect on the convergence duration of the ant based IRC system. As an example, in Fig. 5, by choosing T_E = 3, we observe that at the end of each measurement period, the ant are able (e.g., have enough time) to find the best path correctly. However, if we reduce T_E, to 2 seconds or so, the convergence duration may be longer and ants wouldn't find the best path until the second iteration due to the rapid rate of pheromone evaporation which results in a lower difference between the amount of pheromone on each path. The lower difference of pheromone values in turn makes finding the best path a problematic issue for the ants as a result of probabilistic nature of ants' movement. On the other hand, a bigger value of T_E would make the initial convergence duration lower while bringing about a longer convergence time for the consequence path switching events in the case of sudden congestion on the selected path. Furthermore, it is evident that $T_E \geq T_R$ is a necessary constraint. With these in mind, T_E = 3 is the best selection as expected and seen in our results. We would suggest $T_E \approx 1.5 \times T_R$ in a more generalized scheme.

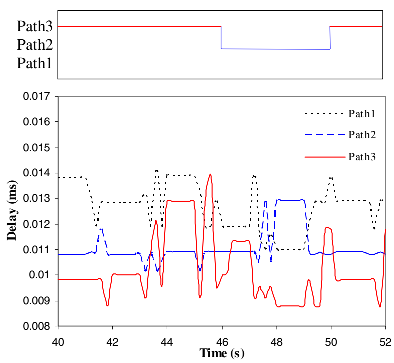

Fig. 7. Comparison of delay of paths at different time (bottom graph) along with chosen path (top graph) in the commercial IRC in scenario 2

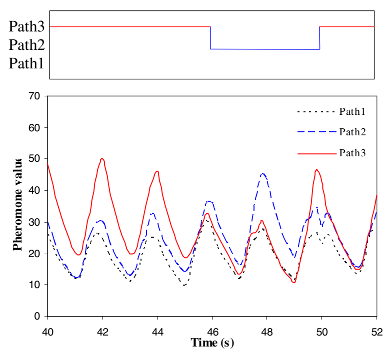

Fig. 8. Comparison of pheromone of paths at different time (bottom graph) along with chosen path (top graph) in the ant based IRC in scenario 1

As for T_R and T_M parameters, the bigger the T_M is, the faster ants tend to find the best path. Thus, this actually imposes a trade-off between the amount of overhead we put into the network for probing the destinations through each available path and the speed of convergence which satisfy us. Furthermore, for faster response in case of network congestion, the idle period T_R -T_M within each iteration may be minimized or even avoided. And, finally, it is important to note that the number of routing probe packets, N_p, is of importance in the effectiveness of the system. Choosing a low value for N_p, affects the correctness of the measurements and convergence time of the ant based system while a large value for N_p imposes too much useless overhead. In addition, an approximate upper bound on N_p for each path can be driven to assure that the probes are acknowledged by the destinations *on time*, i.e., before the next switching decision event. If we represent N_p on egress link l by $N_{p,l}$, we have:

$$\forall_l \frac{T_M}{RTT_l} \succ N_{p,l} \tag{4}$$

where, RTT_l denotes the average round trip time on the path correspondent to l. Considering the path with the maximum RTT, we can calculate the upper bound on N_p. For instance, in our experiments, $T_M = 0.8$ and $\max\{RTT\} \approx 0.2\,ms$ results in an upper bound of 40 for N_p. However, we notice that $N_p = 30$ is big enough in terms of convergence speed.

4.3 Performance Evaluation under Congested Links

In this scenario, to study the effectiveness of the ant based IRC system in relation to that of a delay based one in case of path switching events, we introduce a congestion event at Path3 by CBR cross traffic between nodes 17 and 5, i.e., <17, 18, 5>, during the time interval [44, 48].

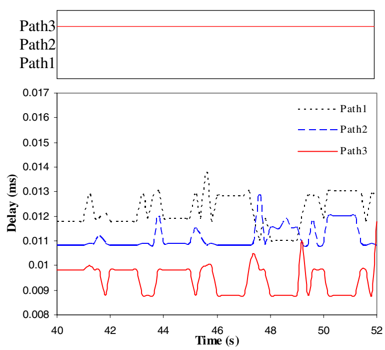

Fig. 9. Comparison of delay of paths at different time (bottom graph) along with chosen path (top graph) in the commercial IRC in scenario 3

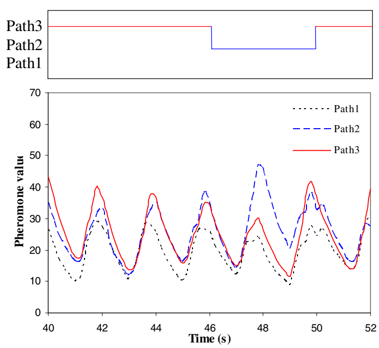

Fig. 10. Comparison of pheromone of paths at different time (bottom graph) along with chosen path (top graph) in the ant based IRC in scenario 3

Fig. 7 and Fig. 8 depict the path switching events of the IRC systems along with their switching metric. We observe that both systems recognize the congestion event on Path3 during interval [44, 46] and on the next switching event in 46 both switch to the best alterative after that path, Path2. Besides, once Path3 has recovered from temporal congestion in 48, both systems have switched back to that path on the next switching event at time 50.

4.4 Performance Evaluation under Lossy Links

In order to analyze effects of lossy links in a wireless like environment, we assume that link <17, 18> losses packets with probability 40% during time interval [44, 48] (See Fig. 11). Fig. 9 and Fig. 10 show the path switching events of the ant based and delay based IRC systems along with their switching metric. We make a few observations. First, the ant based system is still working effectively. Surprisingly, in time period [44, 46] ants have recognized the poor reliability of the link(s) on Path3 and on the next switching event in time 48, they have switched to the Path2 as it is the best alternative considering low delay after Path3. However, the delay based system hasn't figure out this lossy behavior as it only considers the delay metric of the probes which arrive to it for making its decisions. Furthermore, again ants have switched back to Path3, once the lossy behavior of this path is finished.

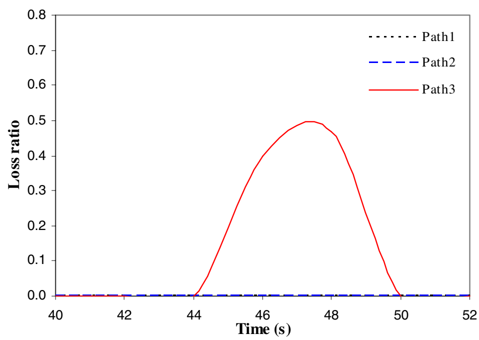

Fig. 11. Comparison of loss ratio of paths at different time in scenario 3

In sum, the results of experimental evaluation of this section evidently suggest that (i) ants efficiently choose the best egress links in comparison to a commercial IRC system. Indeed, all of the switching events of a delay based IRC system are detected and followed by the ants as well. (ii) The path ants heuristically select has low latency and also low packet drops. In other words, ants choose paths with a hybrid metric of delay and loss ratio.

5 Conclusion

The emergence of Intelligent Route Control (IRC) products has enabled the augmentation of the reliability, availability and also the end-to-end performance of multihomed stub networks' connection to the Internet. Leveraging ideas from ant colony optimization (ACO), this paper presents an ant based deployment of IRC systems to see how ants would intuitively implement such systems. Using extensive simulation experiences under different network traffic load and link reliability scenarios, we obtain two important results. First, ants perform quite efficiently in selecting the best upstream path in comparison to commercial IRC systems. Second, we observe that the path ants tend to pick at each time through their heuristic pheromone-basis ACO paradigm is actually the most optimized upstream path in terms of a hybrid metric of both delay and loss ratio.

References

1. InterNAP, Premise-Based Route Optimization
2. Route Science, Adaptive Networking Software
3. Cisco Systems, Optimized Edge Routing (OER)
4. Akella, A., Seshan, S., Shaikh, A.: Multihoming Performance Benefits: An Experimental Evaluation of Practical Enterprise Strategies. In: Proc. USENIX, Boston, MA (June 2004)
5. Akella, A., et al.: On the Performance Benefits of Multihoming Route Control. IEEE/ACM Transactions on Networking (2008)
6. Akella, A., et al.: A Measurement-Based Analysis of Multihoming. In: Proc. ACM SIGCOMM, Karlsruhe, Germany (August 2003)
7. Goldenberg, D., Qiu, L., Xie, H., Yang, Y., Zhang, Y.: Optimizing Cost and Performance for Multihoming. In: Proceedings of ACM SIGCOMM (2004)
8. Dai, R., Stahl, D.O., Whinston, A.B.: The economics of smart routing and QoS. In: Proceedings of the Fifth International Workshop on Networked Group Communications (2003)
9. Schreck, G., Rustein, C., Porth, M.: The end of the private WAN. Forrester Brief (March 2002)
10. Sevcik, P., Bartlett, J.: Improving user experience with route control. Technical Report NetForecast Report 5062. NetForecast (2002)
11. Guo, F., Chen, J., Li, W., Chiueh, W.: Experiences in Building A Multihoming Load Balancing System. In: Proceedings of IEEE INFOCOM (2004)
12. Gao, R., Dovrolis, C., Zegura, E.W.: Avoiding oscillations due to intelligent route control systems. In: Proc. IEEE INFOCOM (April 2006)
13. Dorigo, M.: Optimization, Learning and Natural Algorithms. PhD thesis. Politecnico di Milano, Italy (1992)
14. Sim, K.M., Sun, K.M.: Ant colony optimization for routing and load-balancing: survey and new directions. IEEE Transactions on Systems, Man and Cybernetics 33(5), 560–572 (2003)
15. Bonabeau, E., Dorigo, M., Theraulaz, G.: Inspiration for optimization from social insect behavior. Nature 406, 39–42 (2000)
16. Dorigo, M., Caro, G.D.: The ant colony optimization metaheuristic. In: Corne, D., Dorigo, M., Glover, F. (eds.) New Ideas in Optimization, pp. 11–32. McGraw-Hill, New York (1999)
17. Dorigo, M., Blum, C.: Ant colony optimization theory: A survey. Theoretical Computer Science 344(2-3), 243–278 (2005)
18. Schoonderwoerd, R., Holl, O., Bruten, J., Rothkrantz, L.: Ant-based Load Balancing in Telecommunications Networks. Tech. Rep., Hewlett Packard Lab. Bristol, U.K., HPL-96-35 (1996)
19. Zarifzadeh, S., Yazdani, N.: Joint Resource Conserving and Load Distributing Approach for Routing of Survivable Connections. Elsevier Computer Communications 31(14), 3384–3393 (2008)