Statistical Techniques for Evaluating Adaptive Routing Protocols in a Mobile Tactical Environment

Jonathan M. Nicholas
ITT Industries
Aerospace/Communications Division
Clifton, New Jersey

Abstract - **When studying adaptive routing protocols, we are interested in their ability to find correct routes in a dynamic environment, such as the one experienced in a mobile tactical environment. However, this means we are trying to measure performance in a transient situation, which by definition has changing performance and therefore does not easily provide a statistically reliable sample. This is complicated further by the fact that the dynamic environment may change in many different ways, and realistic assessments of what types of changes may occur might not be easily obtained. In this paper, we discuss techniques for obtaining statistically useful data to evaluate adaptive routing protocols under dynamic conditions.**

I. INTRODUCTION

Adaptive routing protocols are designed to accurately determine network topology. This knowledge is translated into routing tables, which are used to route information packets. If the routing table is inaccurate, packets will be misrouted, causing small delays at best, packet loss at worst. Ideally, an adaptive routing protocol should quickly identify changes in network topology to minimize the time during which information packets might be misrouted. Open Shortest Path First (OSPF) is a popular commercial example of an adaptive routing protocol.

The impact of having misrouted information packets is not always obvious. For instance, if a large number of routine messages are misrouted onto a channel reserved for high priority messages, the impact will be seen in the performance of those high priority messages. Furthermore, those high priority messages would have been routed correctly while the routine messages may have seen no change (or even an improvement) in performance.

Furthermore, the routing protocol itself has impact on network performance since it must use network resources to determine the topology. It is all too easy to construct an adaptive routing protocol which quickly and accurately determines the best routes, but leaves no bandwidth for user traffic.

Therefore, the performance of an adaptive routing protocol must be evaluated based on its impact on the network performance as well. Network performance is usually measured statistically. The reason for this is that many of the difficulties in communications are random in nature (i.e. channel noise). The only way to get a reliable evaluation of performance is to evaluate the statistical performance of a large number of messages.

However, adaptive routing protocols are specifically designed to deal with a situation when the communications environment is changing. The change could be for better or worse and if we take a large sample before the change and after the change we could make an evaluation of which environment was better. But what we are really interested in from the point of view of evaluating the adaptive routing protocol is the performance during the time routing information was inaccurate and how long misrouting occurred. If an adaptive routing protocol responds quickly (and most do), we may not observe enough messages to make an accurate evaluation. For example, if a decrease in completion rate is observed when a change occurs in the network, it could either be due to misrouting or some coincidental channel noise.

To complicate matters further, an adaptive routing protocol could quickly identify a new route to one destination and take longer to identify new correct routes to other destinations. In a large network, it may be difficult to identify exactly which packets were effected by misrouting, due to multiple paths and relatively large end to end delays.

Tests can easily be created to determine the speed with which a particular adaptive routing protocol makes corrections to the routing table for a particular change in topology. However, this sort of test may not be indicative of how the protocol will perform for a different change in topology. Exhaustive testing quickly becomes impossible as the network expands.

0-7803-5538-5/99/$10.00 (c) 1999 IEEE

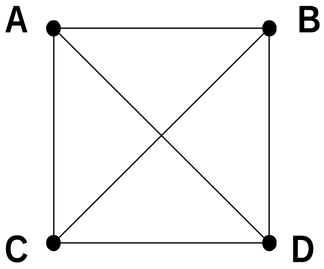

Fig. 1. An abstract 4 node network.

For example, the 4 node network shown in Figure 1 has 6 possible duplex links, which makes 64 possible topologies. Even if we exclude uninteresting topologies (i.e. one with no good links), the number of possible topology transitions is in the thousands. If we consider the possibility of further topology change when the adaptive routing protocol has only partially identified the previous topology change, the notion of exhaustive testing becomes untenable.

II. METHODOLOGIES

In order to properly assess performance of an adaptive routing protocol, we need to conduct tests on a statistically representative sample of topologies and topology transitions. We also need to isolate factors which influence performance independently of the routing protocol used such as RF noise and traffic load.

A. Network Modeling And Simulation

Modeling and simulation lends itself particularly well to this task. Outside factors can be easily held constant. Quick construction of test networks is a relatively simple process. Modeling and simulation gives better control over items of interest such as router failure, link failure and protocol timers. Finally, items which are not easily traceable in a field setting, such as packet flows and routing table changes, can be stepped through moment by moment in the course of a simulation.

All of these advantages mean that when an adaptive routing protocol is observed in a simulation to give poor performance, the reasons for the poor performance can be identified. When these reasons are identified, they can be used to improve the routing protocol, network architecture, or any other factor which is identified as critical to performance.

B. Representative Topologies

Unfortunately, planned topologies frequently are highly inhomogeneous (i.e. bridges or gateways which are choke points in a network), which means that some changes in topology may have greater impact than others. Ideally, we wish to isolate those inhomogeneities to maximize statistical reliability.

Fortunately, (for assessment purposes) getting an accurate projection of what is a realistic set of representative topologies is a nontrivial task. Frequently, the planned topology is only known in generalities with details subject to unexpected changes.

As a result, when confronted with a range of topologies to choose from, researchers can usually choose a topology which has a fair degree of homogeneity. Of course, any remaining inhomogeneities must be handled in a manner appropriate to the magnitude of the inhomogeneity.

Small scale inhomogeneities (i.e. choke points) can be considered noise factors suitable for isolation. That is, if a topology has only one gateway between two areas in a network, data will only be taken while that gateway is a stable part of the network topology.

Larger scale inhomogeneities (i.e. two areas using different transmission media) can be treated statistically and separate data taken for messages sent within each homogeneous area and between homogeneous areas. This may reveal that different routing protocols work better in different areas thereby pointing to a compromise solution.

C. Changing Topology

Given a chosen topology, we can then use knowledge of failure and recovery rates to create a set of representative topologies and topology transitions. That is, for a homogeneous network, the proportion of links or nodes which have failed at any given time should be relatively constant. How often the topology changes can also be calculated based on failure and recovery rates as links and nodes fail and recover.

Usually, a realistic assessment is nearly impossible, especially in a mobile tactical environment. Therefore, a range of values should be tested to assess a variety of situations.

Ideally, failure and recovery rates should be uniform as well. However, inhomogeneities may arise such as forward battlefield units having higher failure rates than rear supply units. These inhomogeneities should be handled in a manner functionally identical to the way topological inhomogeneities are handled.

Small scale inhomogeneities should be treated as noise factors which should be fixed or carefully monitored and statistically isolated from other measurements. Large scale inhomogeneities should be treated as two (or more) separate, but correlated, statistical experiments.

D. Network Loads

In many ways, network loading presents the same problems for evaluating adaptive routing protocols as does the issue of network topology. The load is often produced primarily by one source and realistic values for the projected traffic load are difficult to obtain.

As a result, traffic loads can usually be treated in the same manner as topologies. A range of traffic loads should be tested to ensure that the adaptive routing protocol is robust. The researcher should attempt to keep inhomogeneity to a minimum. This primarily means that traffic should be as evenly distributed as possible among the network members and all messages should be similar in size and type.

When inhomogeneities are unavoidable, small scale inhomogeneities, such as one main traffic source, should be treated as fixed noise factors. Large scale inhomogeneities, such as different message types or sizes, should be treated as separate, but correlated, statistical data sets.

III. EXAMPLE - INTRANET RELAY

In a wireless network, links between stations can become unusable due to a variety of factors such as range, intervening terrain, jamming etc. To deal with this problem, an intranet relay scheme can be used to allow two members of a network with no connection to communicate though a third party which is connected to both.

Obviously, in a mobile tactical environment, the status of wireless links between any two stations is subject to change. Furthermore, wireless links frequently have limited bandwidth to exchange topology information and links may not be duplex.

For this example, we will be evaluating the intranet relay protocol described in MIL-STD-188-220B over a SINCGARS radio network. We wish to determine optimum topology update timer, which is the minimum time between transmission of successive topology updates. We will be testing values of 60, 90 and 180 seconds. As a baseline, we will also be evaluating a scenario with no relay. We will be conducting our study using an OPNET model of an 8 member SINCGARS network.

In our final analysis, we decided that message completion rate is the only performance measure we are interested in. Message delay is also of potential interest, but differences are expected to be relatively small and therefore will not be evaluated for this example.

A. Topology

In this case, the driver of topology, and topology change, is the status of the wireless links between stations. If we consider all links to have an equal probability of being connected or disconnected at any one time, we have created a homogeneous topology.

Specifically, if a station has a 5 in 7 chance of being connected to any other station, then that station will be connected to 5 out of the 7 other stations, on average, at any given time, for an average connectivity of 71%. However, for a dynamic scenario, the specific stations a given station is connected to will change even though the number of connections remains the same. For this example we will use an average connectivity of 60%.

Since changes in connectivity are driven largely by the movements of mobile units, we will base our assessment of connectivity on the notion of an individual moving within the network. Ideally, this motion is continuous, but for our purposes discrete changes, when network connectivity changes occur, will be sufficient. Again, since all members of the network are identical, the requirement for homogeneity is satisfied.

In this example, we will assume that at certain times a randomly selected member of the network moves. At that time, all of the moving member’s links will be reevaluated based on the desired average connectivity. For each other station in the net, the simulation will randomly determine whether it is connected to the moving station or not. The frequency of these changes will be varied for different simulations to obtain results over a range of conditions.

B. Network Load

For this situation, we have determined that we are interested in two types of traffic: unicast and broadcast. To satisfy homogeneity, each station will send the same amount of traffic. Unicast and broadcast traffic will be evaluated separately. For further statistical reliability, we will be sending many small messages rather than a few large messages. In a complete study, we would use a range of traffic loads to ensure robust performance, but since this is just an example, one relatively light load will be used.

C. Results and Analysis

The results of the simulations are plotted in Figures 2 and 3 below. It is clear that the shorter topology update timer values performed better for both broadcast and unicast traffic.

One unexpected result is the rise in completion rate for unicast messages at shorter topology change intervals. This is due to the retransmission attempts at the link layer. In essence, the topology is changing so quickly that before a station determines that one of its neighbors is not in range, the neighbor returns. Since the link layer attempts to retransmit a message twice at 30 second intervals (plus access delays) if an acknowledgment is not received, that particular neighbor returns in time to receive one of the retransmissions.

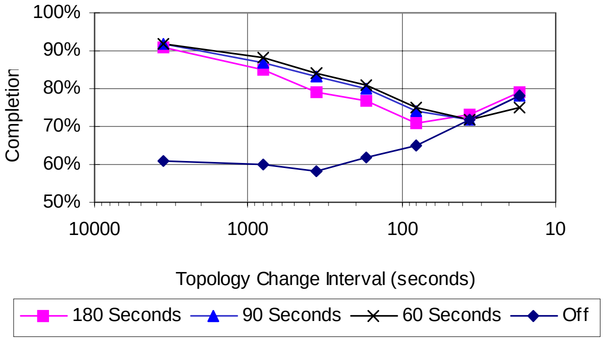

Fig. 2. Unicast completion rate as a function of topology change and topology update message rates.

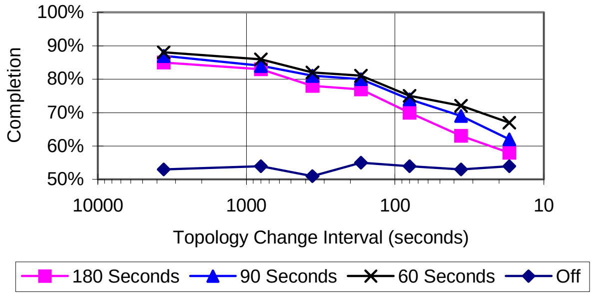

Fig. 3. Broadcast completion rate as a function of topology change and topology update message rates.

In contrast, the link layer does not attempt retransmissions for broadcast messages, so the broadcast message completion declines as the topology change interval decreases.

Finally, we note that the lower completion rate of broadcast messages in the relay off case reflects the effects of on air collisions with no retransmissions. Detailed analysis (not presented here) shows that approximately 10% of all transmissions are lost due to collisions in this scenario. Thus, unicast messages (with retransmission) show a 60% completion rate (the predetermined average connectivity), while broadcast messages show a 54% completion rate.

IV. CONCLUSIONS

Evaluating how effectively adaptive routing protocols maintain efficient network performance is a difficult task. However, by carefully constructing tests which look to maximize homogeneity and evaluating the primary network performance measures, we can come to definitive conclusions concerning the relative performance of adaptive routing protocols and possibly devise improvements to the routing protocol.

REFERENCES

[1] MIL-STD-188-220B

0-7803-5538-5/99/$10.00 (c) 1999 IEEE