IDEA GROUP PUBLISHING
701 E. Chocolate Avenue, Suite 200, Hershey PA 17033, USA
Tel: 717/533-8845; Fax 717/533-8661; URL-http://www.idea-group.com

ITP4548

Rates of Change in Ad-hoc Networks

Alec Yasinsac[1]
Florida State University
Tallahassee, Florida 32306-4530 USA
Tel: 850.644.6407 ~ Email:yasinsac@cs.fsu.edu

ABSTRACT

Ad hoc networking techniques allow low power devices to communicate among themselves utilizing one another as communications relays. Often the resulting networks are highly dynamic, with nodes entering and leaving the network, often for short duration membership.

In this paper, we systematically address issues associated with changes that occur in ad hoc networks. We consider the functionality impact of change and address bounds on optimization that exist when change rates are high.

1. DYNAMICS OF AD HOC NETWORKS

Networks come in all shapes and sizes, with a wide variety of characteristics. We are concerned with networks that have no permanent structure, essentially, all nodes are not only mobile, but they characteristically regularly move about. These networks are comprised of nodes with limited transmission ranges and depend on other nodes to relay traffic in order to expand their broadcast domain. We commonly term these *ad hoc networks* because networks form, change, and dissolve in an ad hoc way, often and quickly, and as a matter of routine. The networks that they form are often highly dynamic.

This paper addresses questions about functional limitations that high and fluctuating rates of change cause in ad hoc networks. In the rest of this section, we systematically set up the discussion by defining key terms and follow with an argument about the important metrics and the bounds that apply given assumptions about these metrics.

1.1. Nodes, Links, Networks and Notation

Ad hoc networks are collections of nodes that intercommunicate by relaying messages across peer to peer links. We label our nodes in capital letters, while links are pairs using the lower case letters that correspond to the nodes that the link connects. Thus, a link between nodes A and C is represented as (a, c), or equivalently as (c, a).

A network consists of a collection, or set, of interconnected nodes. If we label networks with upper case letters from the end of the alphabet, we can say that nodes A and B are elements of network X:

$$\{A, B\} \subseteq X$$

and that if link (a, b) exists, it is also an element of X.

$$(a, b) \in X$$

We define a path as a set of interconnected links that connect two nodes. Paths are represented as ordered tuples, with the number of entries dependent on the number of links that must be crossed. Thus, a path from A to B that must go through C and D (in that order) would be labeled (a, c, d, b), or equivalently (b, d, c, a). These relationships are casually illustrated in Figure 1.

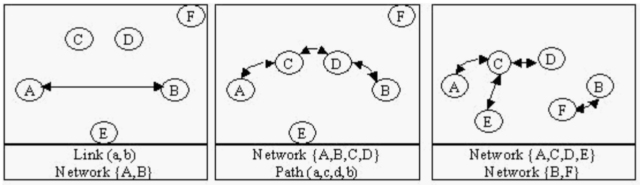

Link (a,b)
Network {A,B}

Network {A,B,C,D}
Path (a,c,d,b)

Network {A,C,D,E}
Network {B,F}

Figure 1

1.2. Network Structure Rate of Change

The *ad hocness* that is a primary characteristic of the networks we consider, results in dynamic networks. As the rate of change increases, the nature of these networks becomes progressively more complex. For example, when a link forms, it may join a node to a network, establish a cycle in an existing network, or merge two networks. Conversely, dissolving a single link can have the opposite three effects.

We consider the specific types of network structure change in order to better understand the nature of networks with high rates of change. We note that there is presently no existing set of measures to reflect these notions.

1.3. Discrete Structures in High Rate of Change Networks

One way to think about network changes is to consider the network structure during static periods, as addressed in [1]. If we define change to occur instantaneously, then we can theoretically identify the network structure at any instant. Practically, network structure change does not occur instantaneously, but rather injects a "change interval" where the system neither has the previous structure, nor the next structure. Still, if the change interval is sufficiently small, we can act as though changes are instantaneous with little impact on our results.

Most attempts to manage ad hoc networks are based on two assumptions regarding the rate of change:

(1) The change interval is insignificant and
(2) There are long network structure static intervals that have a computationally significant interval between relevant changes in the network structure.

The impact of the former depends on the accuracy of the latter. Functions on ad hoc networks assume that the network is static for relevant changes for a period longer than is required to complete the function. For example, a node count function may not succeed if it cannot expect that the connected nodes will not change before the start function completes.

In networks with low rates of change and long static intervals, the change interval is less significant, since network changes occur regularly, and quickly, in ad hoc networks. Adding or deleting a link in a network routinely takes a few seconds at the most. If the network structure is routinely static for hours at a time, the few seconds it takes to make a structural change is insignificant relative to the network structure static intervals. However, if the network structure intervals are short, the few seconds that it takes to make changes have a larger impact.

1.4. Sparsely Populated Ad hoc Networks

Considering the rate of change in network structure necessarily requires scope. We now consider some subtleties of how changes effect sparse versus dense networks. We suggest that large networks will have more changes than will smaller networks. Thus, larger networks are more difficult to manage because of the larger number of changes that occur, and correspondingly, the static intervals are shorter. Conversely, each link in a smaller network tends to be more important to the traffic in that network, so a larger percentage of changes are significant relative to more network functions.

Some networks will have few nodes and few links, while others will have numerous nodes, but are sparsely connected, while still others will have few nodes that are highly connected. We posit that the rate that changes take place has different impacts in each situation. In a sparsely populated network, dissolution of a node is more likely to split the network into two disconnected networks than loss of a single node in a densely populated network. Similarly, loss of a single link is more likely to separate a node from the network if the network is sparsely connected (few links per node) than a more densely connected network.

For the rest of this paper, we employ the somewhat uncomfortable use of the phrase “more dynamic” to address this rate of change. A network that is more dynamic than another has a higher rate of network structural change and a longer static interval on average. We now offer a series of metrics that categorize the dynamic nature of ad hoc networks in the next section.

2. RATE OF CHANGE METRICS

We now define measures of the dynamic nature of ad hoc networks. We introduce building blocks that we use later to make our argument about bounds on routing efficiency. We partition our metrics into architectural and application oriented metrics.

2.1. Architectural Rate of Change Metrics

2.1.1. Link Lifetime

The first metric that we introduce is the network average link lifetime. Intuitively, a network with shorter average link lifetime is more dynamic than a network with longer average link lifetime. This metric is easy to compute: simply sum the duration of the existence of each link that has existed in the network, and divide that by the number of links that have existed. We give the Link Lifetime Average for network X with n total links as:

$$\text{Avg\_LLt(X)} = \frac{\sum_{i=1}^{n} LLi}{n}$$

Practically, it is more difficult to compute the average link lifetime of an ad hoc network, since acquiring complete information is unlikely. Rather than computing the average link lifetime of a network, we begin our arguments by assuming a value for this metric and reason about the resulting impact on network functionality. It is strait forward to model link lifetime using statistical methods. By fixing link lifetime and varying the distributions and impacts of differing assumptions, we can observe the results as the link lifetime increases and decreases.

2.1.2. Node Lifetime

A second metric is node lifetime. Each time a node enters or leaves the network, there is a connectivity impact that may be greater than having a single link change. The computation for the average node lifetime is similar to the average link lifetime, where m is the number of nodes that have existed in X:

$$\text{Avg\_NLt(X)} = \frac{\sum_{i=1}^{m} NL_i}{n}$$

2.1.3. Number of Links per Node

We introduce a metric that connects links and nodes: the average number of links per node. This metric characterizes the redundancy and connectivity of the target network. When considered as a factor of change, it also allows reasoning about how functionality changes as connectivity changes. Simply stated , the number of links per node is represented as the number of links (n) divided by the number of nodes n in the network.

$$\text{Avg\_LpN(X)} = \frac{n}{m}$$

2.1.4. Percentage of change per unit time

We now return to our earlier example of the total and percent of network changes and use this metric as the springboard into talking about application metrics. We define this metric as the number of changes divided by the desired number of intervals of the selected time units. If we elect hours as our time of choice, we sum the number of changes to links and nodes and divide by the number of hours that elapsed in the desired measurement interval.

$$\text{Change(X)} = \frac{\Delta m + \Delta n}{\# hrs}$$

We generate our recommended enhancement to this metric by including the total number of nodes in the computation.

$$\text{Percent\_Change(X)} = \frac{\Delta m + \Delta n}{\# hrs * (m + n)}$$

All of the metrics that we have defined so far are related. If the link and node lifetimes for network X are longer than those for network Y, the percent change of X will necessarily be larger than the percent change of Y. We can also observe limits between these metrics. For example, the number of links can change without the number of nodes changing, since a node may have several links in the network. Conversely, if there are changes in the number of nodes, there must also be changes in the links, since a node is only a member of the network if it has a link in the network.

2.2. Application-Oriented Rate of Change Metrics

We now move on to application-oriented metrics. These metrics reveal the properties that allow us to recognize boundaries on ad hoc network functionality.

2.2.1. Path Length

Nodes communicate through a network over a series of links that together constitute a path. Consider an ad hoc network of n nodes where we desire to identify a path between nodes A and B. Notationally, a path is an ordered set of nodes (represented in lower case) that begins at the source and terminates at the destination. For example, if node C lies between A and B and if A can send messages to B, but they must be relayed by C, we represent the path between A and B as {a, c, b}, or equivalently {b, c, a}, and we term C an intermediate node between A and B.

2.2.2. Path Lifetime

As with links, paths come and go in ad hoc networks. There must be at least one path between any two nodes in a network, however, as links dissolve, paths may also dissolve. We now consider the average path length within a network. In order to facilitate discussion, we assume that we can enumerate the total number of paths in a network, call it q. We then define the average path length as the sum of the number of links in each path divided by the sum of the number of paths in each link.

$$\text{Avg\_PL(X)} = \frac{\sum_{i=1}^{q} PL_i}{q}$$

Another metric for functions that are concerned with paths is path lifetime. This metric can be convenient to perform functions between nodes, e.g. forming a circuit or authenticated route. The path lifetime provides a guideline on how much time the function can take and yet expect not to run out of time, as we described earlier.

We can generate a metric for average network path lifetime similar to that for link lifetime. Path lifetimes will be significantly shorter than link lifetimes in ad hoc networks, since dissolution of any link in the path also dissolves the path.

$$Avg\_PLt(X) = Avg\_LLt(X) / Avg\_PL(X)$$

Similarly, we can compute the path lifetime based on the node lifetime, by recognizing that any path between A and B that traverses i links, must also traverse i-1 intermediate nodes. Then the definition of the average path lifetime for network X may be stated, relative to nodes rather than links, as:

$$Avg\_PLt2(X) = Avg\_NLt(X) / (Avg\_PL(X) - 1)$$

2.3. A Few Illustrations

We argue that the metrics above are practical, that is, that ad hoc networks where these metrics are meaningful exist. To illustrate their utility, we identify four potential ad hoc network categories that correspond to different rates of change.

1. Low Rate of Change Ad hoc Network. The least dynamic (or most stable) category includes ad hoc networks where links exist on the average some number of hours, up to several days. Changes at this rate are relatively easy to handle and do not consume a significant percentage of network resources. An example of such a network is an office environment where employees carry their laptop computers, connected by roaming wireless communications, home and to work with them. While the change rates may peak in the morning and again in the afternoon, the average link lifetimes will likely be hours.
2. Medium Rate of Change Ad hoc Network. We consider networks with average link lifetime of ten minutes to a few hours as a medium rate of change network. The changes at these intervals do not consume even a local majority of the network resources, but the resources consumed are statistically significant. An example of such a network is a delivery service network, where communication between carriers is via short wave radio. Each vehicle may operate primarily within its own area with links to adjoining areas that are interrupted intermittently.
3. High Rate of Change Ad hoc Network. High rate of change networks are characterized by link lifetimes between a few seconds and a few minutes. Managing change in these networks can take a majority of the available resources. An example of such a network is a wireless network between hand held devices in a crowd where individuals move about independently and communicate via low-power, broadcast medium.
4. Very High Rate of Change Ad hoc Network. These networks are characterized by average link lifetimes of just a few seconds. The primary concern for any function on these networks is resource allocation, and their utility is suspect with current technology. An example of such an ad hoc network environment is that of airplanes in a combat or other high-speed environment.

3. IMPACT AREAS

Thus far, we have presented the foundation for reasoning about the nature of rate of change limitations. We now address the impact that rate of change has on applications.

There has been a significant amount of work done on ad hoc routing [2, 3, 4, 5, 6] most geared toward optimizing either the number of messages or time required to acquire an effective route, where a route is available (we do not consider "wait and see" routing protocols where route requests are held and re-forwarded when new links appear). Flood routing provides a ceiling in both areas.

More formally, for any ad hoc network comprised of n nodes, the largest number of messages that are required to derive a route is 2n. If we choose to optimize the number of messages in a new routing protocol, any routing protocol that systematically produces a route with fewer than 2n messages is superior to flooding. In a given environment, if there is no protocol that can systematically produce an effective route with fewer than 2n messages, then flooding is the optimal routing algorithm.

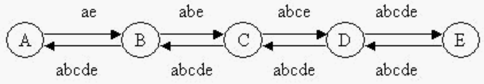

Figure 2

The metrics described above help us to reason about this problem. In this section, we argue that optimization is not possible for some functions in Highly Dynamic networks and use of predetermined routes or circuits may not be possible in networks that are Very Highly Dynamic.

3.1. Bounds on Routing Protocols in Ad hoc Networks

As we described earlier, applications that require circuits are particularly vulnerable to the network dynamic nature. On demand routing protocols generally produce such a circuit. We use the Secure Routing Protocol (SRP) [6] to illustrate how rate of network change can limit network functionality.

SRP is a leapfrog protocol that begins with a route request. Each node that receives the route request appends their address and retransmits if the request is new, and discards otherwise. The route request protocol continues until all nodes in the network receive the route request. If the destination node receives the route request, it prepares a route reply packet directed to the reverse path of the first received route reply. When the originating node receives the route reply, it utilizes the established circuit to communicate with the destination node. Figure 2 illustrates the messages in SRP.

The goal of SRP is to establish a secure route, (circuit) between two hosts on an ad hoc network. SRP establishes a route with only n + *l* messages, where n is the number of nodes in the network and *l* is the path length, a substantial reduction in the number of messages over flooding.

The time required to complete SRP is twice the sum of the times required to move between nodes on the resulting path. Our first observation regarding the impact of the dynamic network nature is that SRP cannot be effective unless the average path lifetime is at least twice as long as the average time required to complete SRP. Otherwise, we should expect that the path identified in SRP would be invalid by the time the protocol completes.

While SRP offers an improvement in the number of messages over flooding, in terms of time, SRP is no better than flooding. Flooding can establish a circuit in the time that it takes to traverse the path from the source to the destination and back; the same amount of time as SRP.

This leads to our first rule regarding bounds on functionality of highly dynamic ad hoc networks, where T(f) is the time required to complete function f.

Rule 1. For any function f that must access a circuit on network X will not be effective unless T(f) < avg_PLt(X)/2, or equivalently if

$$T(f) < Avg\_LLt(X) / 2*(Avg\_PL(X)).$$

Consider some subtleties of this observation. First, we do not claim that functions that violate this rule will never work. Certainly, for shorter circuits or with low probability on longer circuits, functions that violate this rule may occasionally work. However, we cannot *expect* the function to complete its task if Rule 1 is not met.

Secondly, the average link lifetime is a critical element of this computation. In networks in category 4 (very high rate of change), where link lifetimes are only a few seconds, it is likely impractical to expect to be able to utilize circuits at all. Even category 3 networks may be constrained if reliability is essential, or if transmission times are long because of high traffic load or other reasons. Intuition has sensed these observations in the past, but Rule 1 formulates a mechanism to systematically reason about these limiting factors.

3.1. Tuning Factors for Effective Functions in Ad hoc Networks

Another important question is "Can we use Rule 1 to derive a rule that guarantees that such a function will complete"? Since our approach is loosely probabilistic, we prefer to deal with terms such as "likely" and "expected" rather than "guarantee". However, if we accept a slightly loosened form of guarantee, "statistically insignificant" and set that threshold arbitrarily to be less than one percent, we can derive some helpful results.

Rule 2. Any function f that must access a circuit on network X will not be time constrained if

T(f) < avg_PLt(X)/200, or equivalently if
T(f) < (Avg_LLt(X)) / 200*(Avg_PL(X)).

Our arbitrary selection of one percent as our statistically insignificant threshold is fine for illustration, but likely not practical. Although the exact figure will be highly context driven, most network functions cannot endure a one percent failure rate. Fortunately, we can easily tune this threshold and restate Rule 2 as:

Rule 2': Any function f that must access a circuit on network X will not be time constrained relative to the threshold factor (tf) if:

T(f) < avg_PLt(X)/(2* tf), or equivalently if
T(f) < (Avg_LLt(X)) / 2*(tf*(Avg_PL(X))).

4. CONCLUSION

In this paper we have shown how the varying rates of change in ad hoc networks affect the their functionality. We categorized these rates and established metrics to allow systematic analysis of their impact. We went on to address specific functional bounds that may occur for highly dynamic ad hoc networks, showing how one secure routing algorithm cannot be effective in very highly dynamic networks. Using our threshold factor, we show how to use our metrics to gauge functionality in any ad hoc network.

Our work and examples in this paper are limited by space to focus on applications that employ circuits in ad hoc networks. However, these metrics and techniques are applicable to a wide variety of functions and environments and can be a productive mechanism for designing and analyzing applications in ad hoc networks.

ENDNOTE

[1] This material is based upon work supported in part by the U.S. Army Research Laboratory and the U.S. Army Research Office under grant number DAAD19-02-1-0235.

5. BIBLIOGRAPHY

[1] Prosenjit Bose, Pat Morin, Ivan Stojmenovic and Jorge Urrutia, "Routing with Guaranteed Delivery in Ad Hoc Wireless Networks", Wireless Networks, Vol. 7, pp. 609-16, 2001, Kluwer.

[2] Stephen Carter and Alec Yasinsac, "Secure Position Aided Ad hoc Routing", to appear in Proceedings of the Third International Conference on Computer and Communication Networks, IEEE Computer Society Press, November, 2002

[3] C. Perkins and E. Royer, "Ad hoc On-Demand Distance Vector Routing," Proceedings of the 2nd IEEE Workshop on Mobile Computing Systems and Applications, February 1999, pp. 90-100.

[4] Y. Ko, and N. Vaidya, "Location Aided Routing in Mobile Ad Hoc Networks," The Fourth Annual ACM/IEEE International Conference on Mobile Computing and Networking, October 1998.

[5] S. Yi, P. Naldurg and R. Kravets, "Security-Aware Ad-Hoc Routing for Wireless Networks," UIUCDCS-R-2001-2241 Technical Report, August 2001.

[6] Papadimitratos and Z.J. Haas, "Secure Routing for Mobile Ad hoc Networks," SCS Communication Networks and Distributed Systems Modeling and Simulation Conference, January 2002.