On the Repair Time Scaling Wall for MANETs

Vinod Kulathumani[1], Mukundan Sridharan[2], Anish Arora[2], Bryan Lemon[1], and Kenneth Parker[2]

[1]West Virginia University, Morgantown WV 26506
[2]The Samraksh Company, Dublin OH 43017

***Abstract*—The inability of practical MANET deployments to scale beyond about 100 nodes has traditionally been blamed on insufficient network capacity for supporting routing related control traffic. However, this paper points out that network capacity is significantly under-utilized by standard MANET routing algorithms at observed scaling limits. Therefore, as opposed to identifying the scaling limit for MANET routing from a capacity stand-point, it is instead characterized as a function of the interaction between dynamics of path failure (caused due to mobility) and path repair. This leads to the discovery of the repair time scaling wall, which is used to explain observed scaling limits in MANETs. The factors behind the repair time scaling wall are identified and techniques to extend the scaling limits are described.**

***Index Terms*—MANET, path failure, repair time, network capacity, link estimation, local routing, neighborhood discovery**

I. INTRODUCTION

Despite several years of research in MANETs, it turns out that most practical deployments do not scale beyond about 100 nodes [1]. The most common reason attributed to this limitation is that of bounded wireless network capacity [17, 19]. As the number of nodes in a network grows, the overall capacity in the network only grows as $O(\sqrt{n})$ [2]. Several researchers have compared this growth rate with the required capacity to support routing related control traffic (i.e., the network layer overhead), and have argued that the wireless network capacity does not scale with the routing needs. But our studies show that at a scale of about 75-150 nodes, the amount of channel capacity used to support routing turns out to be under 2%, and yet the median path reachability (the number of connected paths) is less than 50%.

Specifically, we constructed and simulated the standard implementation of the Optimized Link State Routing (OLSR [3]) using network scenarios of size 75-150 nodes. To calculate the number of broken paths, we took periodic snapshots of the nodes routing tables every 100 milliseconds and traversed the path derived from the routing table to find the reachability between every pair of nodes. Further, in order to isolate the effect of Network Layer Overhead (NLO) on scaling, the simulations were done in the absence of any useful data traffic. At a scale of about 100 nodes, we observed that around 50-60% of the paths remained broken all the time. However, the network layer overhead was only accounting for about 1.5% of the network capacity. The details of this simulation are described in Appendix A, but in summary, our analysis shows that the *routing system failed well before the capacity limits were reached.*

The ***objective of this paper*** is to investigate alternate reasons for this failure. To do so, instead of formulating this problem from a capacity standpoint, we characterize routing path reachability as a function of the interaction between the dynamics of path failure and path repair. Informally speaking (formal definitions are provided in Section IV and V), the path connectivity interval refers to the average time that end to end paths in a network remain connected before mobility causes the paths to be disconnected. Note that breaking of a single link on a path is sufficient for the path to break. On the other hand, the path repair interval refers to the average repair time for end to end paths in a network. By comparing the path connectivity and repair intervals as a function of network size, we are able to identify factors other than channel capacity that limit the scaling of MANETs. Our analysis leads to several interesting findings which are summarized below.

a. The median path connectivity interval falls as $O(1/\sqrt{n})$, where *n* is the network size. On the other hand, the median path repair interval remains fairly constant irrespective of the network size and is roughly equal to the link failure estimation time, i.e., the time required to detect that a link no longer can be used for routing because a node has moved beyond its communication range. When the path connectivity interval falls below the path repair interval, we can expect a majority of the paths to be disconnected, and the scale at which this occurs is defined as the *repair time scaling wall* for the MANET.
b. Our analysis shows that the crucial controllable factor which impacts the repair time scaling wall is the link failure estimation delay. Using our analysis, one can

This work was supported in part by Defense Advanced Research Products Agency (DARPA)'s Fixed Wireless at a Distance program under contract FA8750-12-C-0278. The views expressed are those of the authors and do not reflect the official policy or position of the Department of Defense or the U.S. Government. Approval for public release was granted by DARPA on March 18, 2014.

determine the bound on link estimation delay for a given network size under a given mobility model so that the system stays within the repair time scaling wall.

c. By analyzing the distribution of path repair intervals, we determine that most of the broken paths can be fixed quite close to the failed links and thus local updates are mostly sufficient for fixing broken paths and restoring connectivity. *Therefore, contrary to previous knowledge, the channel capacity required for propagating link state updates throughout the network does not play the dominant role in limiting the scalability of MANETs.* Instead, it is the convergence of path connectivity interval towards the path repair interval as network size increases, which plays a dominant role in limiting the scalability of MANETs.

Outline of the paper: In Section II, we state related work. In Section III, we state the network model. In Section IV and Section V, we separately analyze the path connectivity and path repair dynamics in a fully connected mobile ad-hoc network. We utilize the results of these analyses in Section VI, to characterize the repair time scaling wall for MANETs. We conclude in Section VII.

II. RELATED WORK

Our interest in analyzing point to point routing protocols for MANET stems from the necessity of these protocols for information sharing in MANETs as the scale of the networks starts to increase. Previous studies [11] have compared the efficiency and impact of stateful communication strategies (such as point to point routing) and stateless strategies (such as flooding) for MANETs under different levels of network connectivity, density and mobility rates. The study in [11] points out that under high mobility and connectivity, flooding is the right choice. But the study in [11] does not consider network scale. While flooding based solutions may be acceptable for small scale networks, as the scale starts to increase to several hundreds and thousands of nodes, stateless protocols start becoming less of an alternative and point to point routing becomes necessary. Our point in this paper is that the scalability issues attributed to point to point routing are not caused by capacity constraints, but rather discovery latencies.

There has been plenty of research on routing algorithms for mobile, ad-hoc networks in the past two decades. Some well-known examples are OLSR [3], TBRPF [4], STAR [5], ZRP [9], DSR [6], AODV [7], and DSDV [8]. At a high level, these algorithms can be classified in two ways: (i) proactive or reactive depending on whether they maintain up-to-date routing information at each node (proactive) or whether they determine routes on-demand when a packet is to be routed (reactive), and (ii) link-state or distance vector [11], depending on whether information about individual links is exchanged globally (link-state) or whether path information at each node is exchanged locally (distance-vector).

In general, the main emphasis in these studies has been on reducing the network level overhead required for propagating route information [19] so that they remain within the *capacity scaling limits* imposed in the seminal paper by Gupta and Kumar for wireless networks [12]. While such a strategy may yield satisfactory routing performance under low mobility and in small networks, practical deployments have not been successful beyond 75-100 nodes [1]. A case in point is the Optimized Link State Routing (OLSR) which introduces selection of multi-point relays (an optimally chosen subset of two-hop neighbors) for controlling the propagation of neighborhood information across the network. As a result, although the routing overhead is reduced, the quality of routing is poor, as quantified by our experimental results in the Appendix A of this paper. Moreover, at the point of poor routing performance, the capacity of the network is in fact severely underutilized. Therefore, our stand in this paper is that network layer overhead is not the right metric for characterizing routing performance and understanding scaling limits for MANETs. Instead, we formulate an alternate technique to quantify MANET scaling limits, wherein as opposed to analyzing routing protocols from a capacity standpoint, we analyze them by comparing the probability distributions of path connectivity and path repair intervals.

Our results in this paper show that it is not capacity but rather the *latency in discovering broken links and fixing broken paths that imposes the crucial scaling limit for MANETs*. We have shown this by characterizing, both analytically and experimentally, the dynamics of path connectivity and path repair in MANETs of different sizes. Our results point out that quick and efficient link estimation is critical for extending the scale of MANETs, *which previous studies have largely overlooked.* Another striking observation that we make in this paper is that the latency for fixing a majority of paths in the network is very close to the latency for discovering a broken link, implying that a majority of the paths can actually be fixed quite locally. This is significant because it shows that much of the routing overhead involved in propagating network wide information by traditional routing protocols is actually not needed.

Lastly, we would like to point out that in this paper we are interested in real-time data routing and hence the idea of introducing delay tolerance to increase the network capacity and recent results on capacity-delay tradeoffs for mobile ad-hoc networks [20, 21] are not directly

applicable here. Moreover, as soon as we weaken the notion of path connectivity by stretching temporally, our results do not apply.

III. MODEL

We consider a mobile network of N nodes deployed over a two dimensional region. The communication range of the nodes is constant irrespective of network size N. We assume a random walk mobility model [11] for the nodes. In this mobility model, at each interval a node picks a random direction uniformly in the range $[0, 2\pi]$ and moves with a constan speed randomly chosen in the range $[v_{min}, v_{max}]$ for a constant distance φ. At the end of each interval, a new direction and speed are calculated. This model is Brownian in its characteristics; the Brownian model can be described as a scaling limit of this motion model under small step sizes [12]. The random walk motion model results in node locations that are uniformly distributed across the network [13]. Therefore, we assume that over time the average number of neighbors per node is ρ and this number stays constant irrespective of the network size.

We assume that the network is never partitioned. Thus, there exists at least one path between every pair of nodes in the network, at all times. The MANET is assumed to be supported by an underlying routing algorithm that is responsible for discovering a path between each pair of nodes, which is optimal under a given metric. Let $\mathrm{P}_{s,d}(\mathrm{t}) = < s, g1, g2, \ldots, , gk , d >$ be the path between the pair of nodes <*s, d*> as determined by the routing algorithm at time *t*. In this path, g_1 represents the next hop towards d as determined by the routing table at node s, and g_{i+1} represents the next hop towards as d as determined by the routing table at node g_i. Since the network is mobile, neighboring nodes move in and out of each other's transmission range, thus adding and breaking links respectively. We assume that the rate at which links are added and deleted to be constants z and θ respectively. The path $\gamma(s, d)$ at time *t* is valid only if all the links along the paths exist at time *t*. Thus if any of the links are broken, the path is assumed to be disconnected. The underlying routing algorithm fixes such broken paths between node pairs, by first detecting broken paths and then restoring them via alternate routes. Thus, each end to end path in a MANET can be represented as a sequence of connected and disconnected intervals in time.

IV. PATH FAILURE DYNAMICS

A. Analytical characterization of path connectivity interval

Definition [Connectivity Interval of a given path]: Let $P_{x,y}(t)$ denote the path between a pair of nodes *x* and *y* in the network at a given time *t*. The *connectivity interval* for the path $P_{x,y}(t)$ is defined as the duration after which at least one of the links on that path breaks, causing the path to break.

Definition [Average Path Connectivity Interval]: The *average path connectivity interval* is defined as the average duration of the connected instance for each path in the network.

Definition [Average Path Failure Rate]: The *average path failure rate* is defined as the number of times that each end-to end path changes from a connected to disconnected state per unit time, averaged over the number of end-to-end paths in the network. Roughly speaking, the average path failure rate grows as the inverse of the average path connectivity interval.

Definition [Median Path Connectivity Interval]: The *median path connectivity interval* for a network is defined as the median of a distribution of path connectivity intervals in the network over a given duration. The median path connectivity interval is thus equal to the smallest time *T* such that 50% of the paths have a path connectivity interval smaller than *T*.

Theorem 3.1: The average length of each path in a network of n nodes is $\boldsymbol{O(\sqrt{n})}$.

Proof: We assume that the spatial distribution of nodes is uniform across the network with a constant density of ρ. Consider any node in the network. The number of nodes within a circle of diameter d around the node is $O(\rho d^2)$, where $1 \leq d \leq \sqrt{n}$. Therefore, the number of nodes between a distance of d and $(d-1)$ from point p is $O(d)$. Therefore the number of paths between lengths d and $(d-1)$ in the network is $O(nd)$. The total length of all n^2 paths in the network can be found as follows:

$$\mathrm{O}(\sum_{d=0}^{\sqrt{n}} \mathrm{n}d^2) = \mathrm{O}(\mathrm{n}.(\sqrt{\mathrm{n}})^3) = \mathrm{O}(\mathrm{n}^{\frac{5}{2}})$$

The average length of each path is thus equal to $O(\sqrt{n})$.

Theorem 3.2: The average path connectivity interval for a MANET decays as $O(1/\sqrt{n})$, where n is the number of nodes in the MANET.

Proof: The link failure rate is a constant of θ per unit time. Thus the average interval between two successive times that a given link breaks is $1/\theta$. Now consider a set of *p* links whose failures are uniformly distributed in time with a constant rate of θ per unit time. In this case, the average time between two successive link failures is $1/\theta p$. We note from the previous theorem that the average length of a path grows is $O(\sqrt{n})$ and that it is

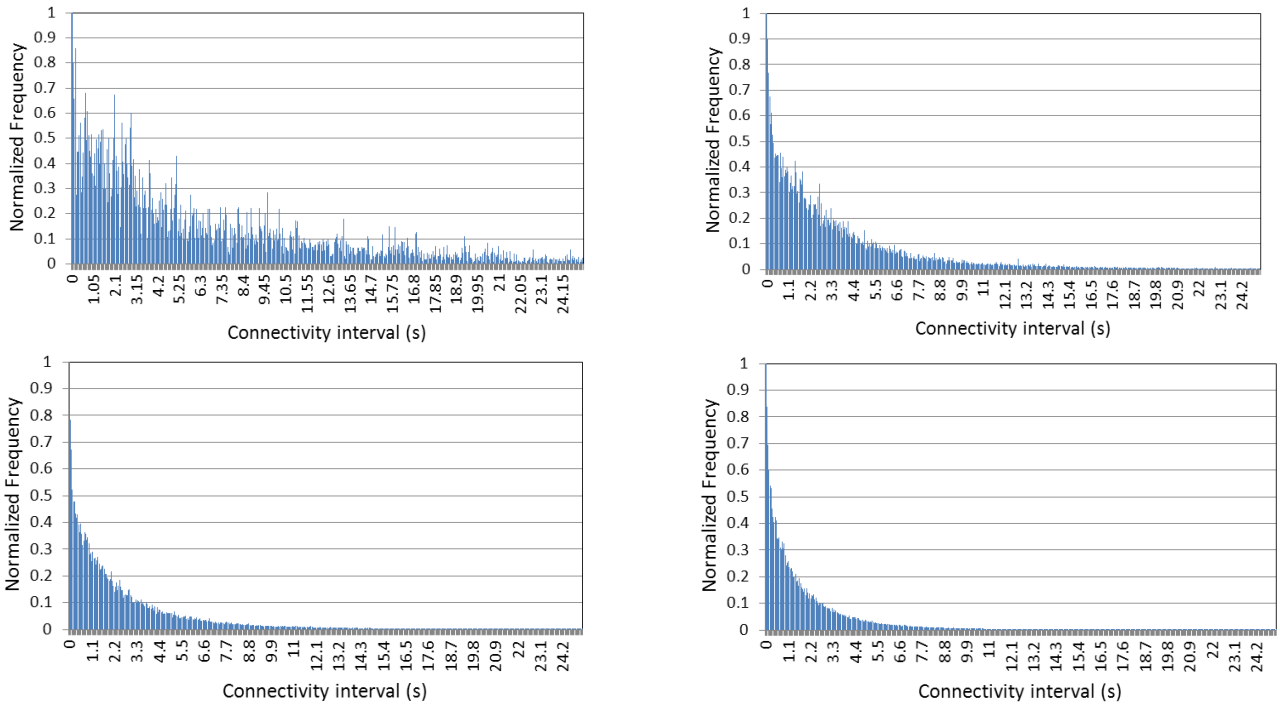

Figure 1: **Histogram of Path Stability Intervals with 50, 150, 300 and 500 nodes. Note that the median and mean path stability intervals progressively shift to the left.**

sufficient for one link on a path to break for the path to be disconnected. Thus, the average time that a path remains connected is $1/\theta\sqrt{n}$.

B. Experimental Characterization of Path Stability Interval

1) Network setup:

We characterized the path connectivity intervals for different network sizes using simulations in ns-3. We considered a random walk 2-d mobility model where each node moves inside a fixed rectangular area with a speed chosen uniformly within the range 2-4 m/s and changes directions after moving in a randomly chosen direction for 30m. The mobility model is such that the number of link changes per second per node is approximately 0.25 irrespective of the network size. The deployment area relative to the communication range is such that the average number of neighbors for each node is approximately 8. Occasionally, the network may be disconnected because of mobility.

2) Routing and Link Estimation:

We assume that there is a correction process, i.e., a routing algorithm executing in the system. Specifically, we used an event-based link state routing algorithm (LSR) that generates a link state update per link estimation event (i.e., either the discovery of a new neighbor or the discovery of the loss of a neighbor). A beacon based algorithm is used for link estimation, i.e. each node beacons a heartbeat message at a steady rate of *H* Hz. The radio range is kept constant to enable knowledge of ground truth. 3 missed heartbeats are used to signal the loss of a neighbor.

3) Path connectivity results

We collect routing statistics at regular intervals of 50ms during each trace of the simulation. This is done as

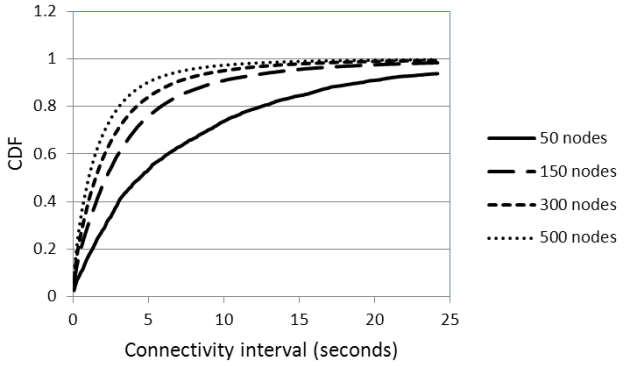

Figure 2: Cumulative distribution function of path connectivity intervals, for different network sizes

follows. Between each pair of nodes <*s, d*> in the network, $P_{s,d}(t) = < s, g1, g2, \ldots, , gk, d >$ is the path as determined by the routing algorithm at time *t*. In this path, g_{i+1} represents the next hop towards as *d* as determined by the routing table at node g_i. We mark the path between <*s, d*> as disconnected at time t, if any of the links along the path is not valid at time t. Link validity is determined

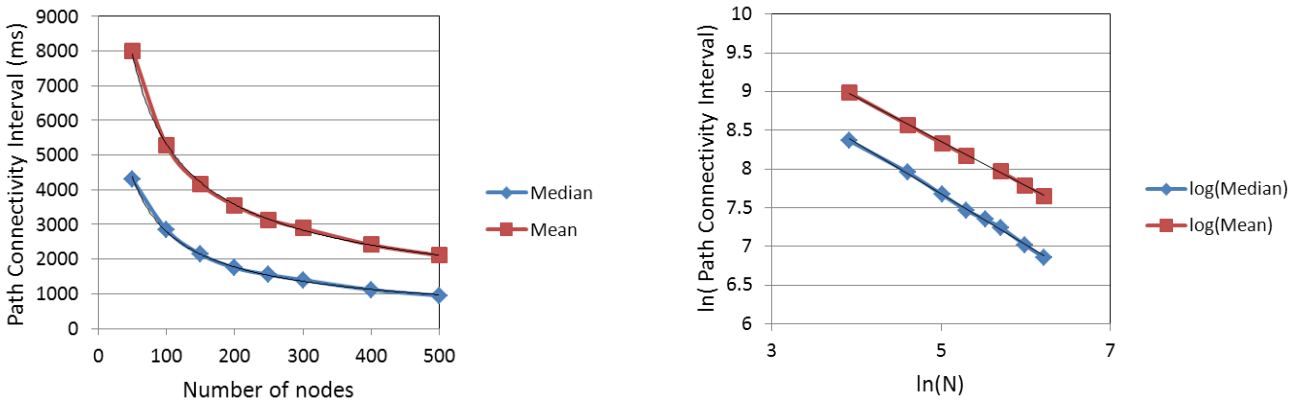

Figure 3: (a) Median and mean path connectivity interval as a function of network size. (b) log-log plot with a slope of approximately -0.5 highlight the $O(1/\sqrt{n})$ asymptotics

by comparing against the ground truth location data and noting that a link between two nodes is invalid if the distance between them is greater than the radio range.

For each pair of nodes in the network, we note down the durations for which the path between them remains connected over the course of the simulation. Note that during the course of each simulation, paths will be continually disconnecting and getting repaired. Thus, the path connectivity interval for each instance of a path being connected is noted down. This data is used to plot the histogram of path connectivity intervals (shown in Fig. 1). The data in Figure 1 shows the histogram of path connectivity intervals at networks sizes ranging from 50 to 500 nodes. The y-axis is normalized by the total number of instances at each network size. We observe from this figure that the mean and median of the histograms shift progressively to the left. This indicates that the path connectivity intervals decrease with the network size. To better understand the histograms, in Figure 2 we plot the cumulative distribution function for the path connectivity intervals, computed over the duration of each simulation. This figure shows that the tail of the distribution becomes smaller as network size increases.

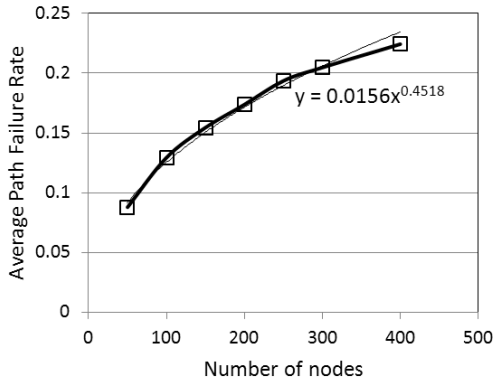

Figure 4: Average path failure rate as a function of network size.

Then we characterize the decay rate for the median and mean path connectivity intervals. Figure 3 (on the next page) shows that the median and mean path connectivity interval for the network decay as $O(1/\sqrt{n})$. Figure 3b is a log-log plot of the path connectivity intervals. The slopes of the lines are approximately -0.5, highlighting the $O(1/\sqrt{n})$ asymptotics. The medians are lower than the respective means, indicating the large range for path connectivity intervals in the network. The connectivity intervals for some paths are relatively high.

In Figure 4, we show the average failure rate per path in the network. Specifically, we note the number of times that each path moves from the connected to disconnected state over the duration of the simulation, and average this over the number of paths in the network. Figure 4 shows that the average path failure rate grows as $O(\sqrt{n})$. All these results match our analysis in Section III.

Finally, we note that during the simulations the network might be occasionally partitioned, i.e. there may be no paths in the ground truth between nodes at certain times. For computing the data shown in Figure 1, Figure 2 and Figure 3, we ignore such instances. Thus, the path connectivity intervals are only determined over pairs of nodes for which some path exists in the ground truth data.

V. PATH REPAIR DYNAMICS

Definition [Repair Interval of a given path]: The repair interval of a path between a pair of nodes *x* and *y*, is the time taken to restore the connection between *x* and *y*, starting from a disconnected state.

Definition [Average Path Repair Interval]: The average path repair interval is the average duration of the disconnected instance for each path in the network.

Definition [Median Path Repair interval]: The median path repair interval for a network of size *N* is defined as the median of a distribution of path repair intervals for all

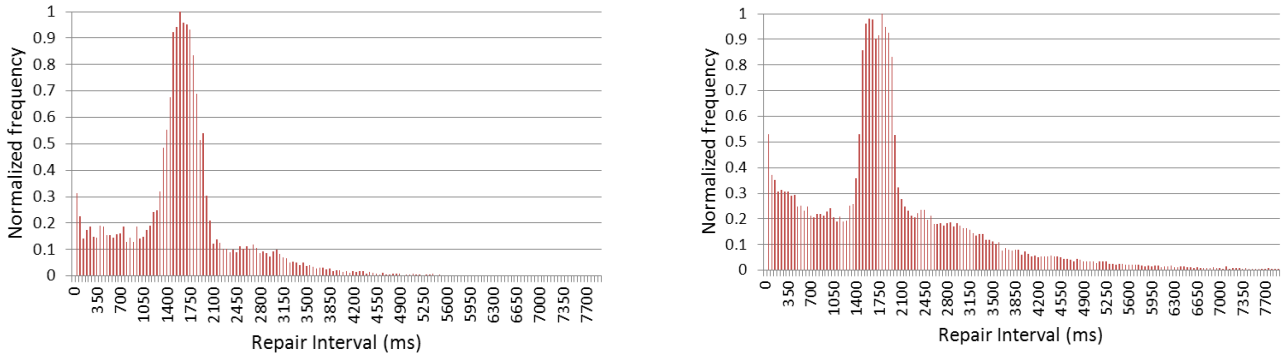

Figure 5: (a) Histogram of path repair interval with 150 nodes (b) Histogram of path repair intervals with 500 nodes

paths in the network over a given

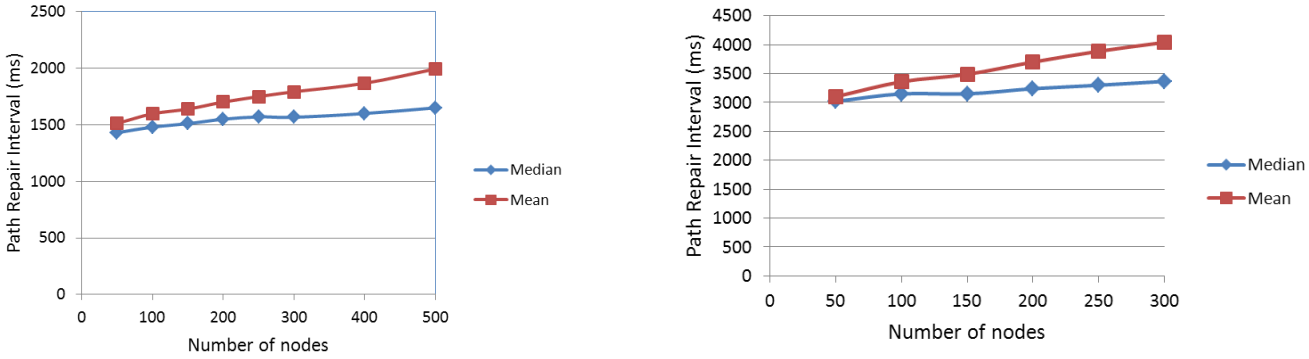

Figure 6: (a) Median and mean path repair interval with a beacon interval of 500 ms (b) Median and mean path repair interval with a beacon interval of 1000 ms. The plots show that the median repair interval stays approximately equal to (slightly higher than) three times the heartbeat interval, i.e. roughly equal to the time that it takes to discover a failed link.

duration. The median Path Repair Interval is thus equal to the smallest time *T* such that 50% of the paths have a Path Repair Interval smaller than *T*.

Once a link state change has caused a path to disconnect, the repair process now consists of two components: (i) discovery of the failed link and (ii) propagation of the discovery to other nodes in the network so as to restore paths that were broken as a result of the link state change. In general the updated link state information only needs to flood a small region completely surrounding the destination node in order to fix all broken routes. Once that has occurred packets originating from outside the region will follow the old routes until they intersect the region of updated routes, at which point they will follow the updated routes to the correct destination. This is what we call as the repair interval. It is important to note *that restoring connectivity does not mean that the path is now optimal and it does not mean that the path has stabilized*. As the information about link changes propagate in the network, better paths may continue to be established between the same pair of nodes.

A. Experimental Characterization of Path Repair Interval

As described in Section III.B, we collect routing statistics at regular intervals of 50ms during each trace of the simulation. For each pair of nodes in the network, we note down the durations for which the path between them remains disconnected (i.e., under repair) over the course of the simulation. Note that during the course of each simulation, paths will be continually disconnecting and getting repaired. The path repair interval for each instance of a path being connected is noted down. This data is used to plot the histogram of path repair intervals (shown in Fig. 5). The y-axis is normalized by the total number of instances at each network size. As stated in Section III, we ignore node pairs for which no path exists in the ground truth data.

We observe from Figure 5 that the median path repair intervals stay approximately constant. This is highlighted more clearly in Figure 6, where the median and mean repair intervals are plotted as function of the network size. This data is shown at two different heartbeat intervals (500ms and 200ms). We make 3 observations.

a. First, we see that the repair intervals follow a uniform distribution until the points close to the median. Then

we notice that most of the paths have repair interval close to the median and the repair intervals follow a power law distribution after this point and quickly taper off.

b. Second, we observe that the median and mean exhibit a small increase with the network size, but this increase is more pronounced for the mean. This is because as network size increases, the worst case repair intervals tend to increase. For some paths, the discovery of a failed link needs to be propagated throughout the network before the path can be fixed.

c. The median repair interval stays approximately equal to (slightly higher than) three times the heartbeat interval, i.e. roughly equal to the time that it takes to discover a failed link. Thus for a majority of the paths, the repair interval is close to the link failure estimation time. This indicates that a majority of the paths are fixed close to the location where a link failure is discovered,

VI. JOINT ANALYSIS OF PATH FAILURE AND PATH REPAIR

time scaling wall as the point at which the median path stability interval falls below the median path repair interval.

We consider this particular definition of the scaling wall to be significant because we expect that at this point a majority of data packets being routed through the network are expected to almost always encounter a broken path somewhere along their path. In systems where intermediate nodes drop data packets upon reaching a dead-end, this would imply poor throughput. In systems where intermediate nodes buffer dropped packets and re-establish a route, this would imply higher latency and a higher buffering overhead.

For a given mobility model and repair interval, our analysis can be used to determine the repair time scaling wall. This is shown in Figure 7. The path connectivity interval and path repair interval are compared with two different link estimation parameters, one with a heartbeat rate of 500ms and another with a heartbeat rate of 1000ms. The repair time scaling wall is indicated for both these scenarios. With a heartbeat frequency of 500ms, the expected scaling limit is about 300nodes. With a heartbeat frequency of 1s (i.e., a link failure estimation time of approximately 3s), the expected scaling limit is about 75-100 nodes.

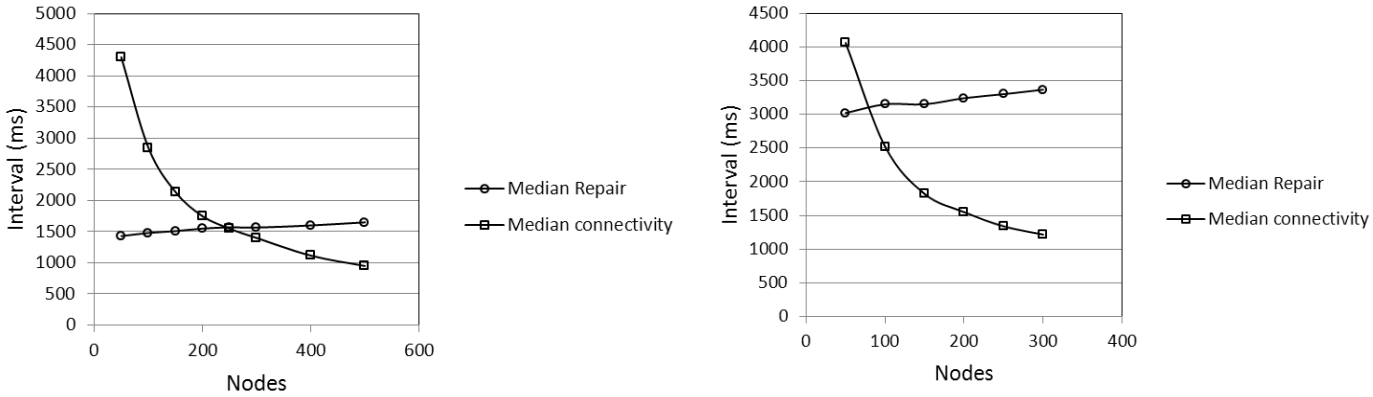

Figure 7: (a) Repair time scaling wall: heartbeat beaconing interval 500ms (b) Repair time scaling wall: heartbeat beaconing interval 1000ms

By putting together the analysis of path stability and path repair dynamics, we are able to make some interesting observations on the scaling limits for MANETs.

A. Existence of repair time scaling wall

The median path connectivity interval in a network of N nodes decays as $O(1/\sqrt{n})$. This is mainly because, as the network size increases, the average path length increases as $O(\sqrt{n})$ and given that links fail at a uniform rate, the expected time before a path hits a broken link decreases as $O(1/\sqrt{n})$. On the other hand, we observe that the median path repair interval stays fairly constant. In general, we expect good connectivity in the network when paths are repaired much faster than the rate at which they are broken. In this paper, we have defined ***repair***

B. Factors behind the scaling wall

The crucial *controllable* factor that impacts the stabilization scaling wall is the link failure estimation time. The faster the failures can be detected the lower will be the expected repair interval – thus increasing the scalability of the system. This highlights the importance of faster link estimation in a MANET.

By decreasing the link failure estimation time, the scaling wall can be pushed farther. In Figure 8, we show the scaling wall with a heartbeat frequency of 200ms (5Hz). The link failure estimation time (and hence the median repair interval) is approximately around 600ms. This pushes the scaling wall to approximately 1500

nodes. Note that the path connectivity intervals only depend on the network size, and are independent of the heartbeat frequency or the link estimation time.

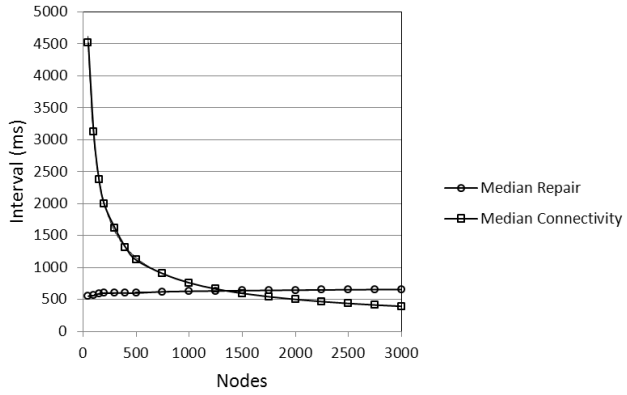

Figure 8: Repair time scaling wall with a heartbeat interval of 200ms, i.e. an average link failure estimation time of 600ms.

Our analysis of the path connectivity intervals can be used to determine bounds on the path repair interval, such that the system does not cross the repair time scaling wall for any given network size. For instance, we can observe that for a network size of 2500 nodes, the path repair interval should be lower than 400ms. Such a determination can be used to guide the design of the link estimation strategy and parameters such as beaconing interval, duty cycling etc.

While the link estimation delay can be decreased by faster beaconing, we note that the scheduling of these beacons will impose a lower bound on the achievable link estimation. In the duty cycled, almost always-off scenario, the lower bound will be significantly higher.

C. Capacity required for link state updates – not the critical scaling factor

The distribution of path repair intervals shows that a majority of the paths have repair time almost equal to the link failure estimation time and the repair intervals fall off as a power law distribution after the median. This indicates that very few paths require multi-hop information propagation before connectivity is restored. In general, local updates are mostly sufficient for fixing broken paths and restoring connectivity. Thus, our analysis shows that the channel capacity required for propagating link state updates does not play the dominant role in limiting the scalability of MANETs. Instead, it is the convergence of path connectivity intervals with the link estimation latency at higher network sizes, which imposes a dominant scaling factor.

Link estimation by itself only requires local information exchange and is not capacity intensive. We have shown elsewhere [15] that if the density of a network remains constant (irrespective of network size), then the required capacity for neighborhood state exchange scales as $O\left(\log(n).\sqrt{n}\right)$ bit-meters per second which exceeds the capacity growth rate for a network only by $O(\log(n).)$ Now, as the network size increases, the required network density to maintain connectivity w.h.p grows as $O\left(\sqrt{\log(n)}\right)$. Under this scenario, the required capacity for neighborhood state exchange scales as $O\left(\sqrt{\mathrm{n}}.\log^{1.5}(n)\right)$ bit-meters per second which still only marginally exceeds the capacity limits by a factor of $O(\log^{1.5}(n))$. Therefore, the channel capacity does not impose the critical scaling constraint for route reachability. On the other hand, the average path stability interval in a network of N nodes decays as $O\left(1/\sqrt{n}\right)$, and is likely to hit the scaling limits imposed by the repair latency sooner than the capacity imposed scaling limit.

VII. CONCLUSIONS AND FUTURE WORK

In this paper, we have identified an important factor that limits scalability of routing systems in MANETs, namely the repair time scaling wall. The repair time scaling wall occurs because as network size increases, the average duration that a path remains connected decreases, while the average duration to repair a path remains fairly constant. When the average path connectivity interval falls below the path repair interval, the scaling wall is reached. We have shown that the path repair interval is roughly equal to the link failure estimation delay in the system. Thus, faster link estimation is critical for extending the path stability scaling wall. This leads us to explore efficient techniques for link estimation and to identify the limits on link estimation intervals, which is a subject of our ongoing work.

Our analysis also shows that much of the channel capacity utilized by standard versions of LSR in propagating link state updates throughout the network do not yield significant improvements in route reachability. This leads us to explore alternate routing protocols, where link state updates are only conditionally forwarded based on their expected impact on the path changes. This is also a subject of our ongoing work. While it is true that such a protocol may not always correct all the paths, need to know LSR can be supplemented by low frequency link state updates that propagate throughout the network.

In this paper, we have compared the path connectivity intervals with the path repair intervals. However, we note that even when a path stays connected, the paths may fluctuate. Some of these could be genuine fluctuations in search of better (optimal) paths while others may be a result of false link estimation events generated by the link estimation service. Such unnecessary fluctuations may have a cascading impact on the system performance because each link state event results in control traffic for route correction. Therefore, the *stability of connected*

paths is an important metric that needs to be studied further.

REFERENCES

[1] Request for Information, Novel methods for information sharing in large scale mobile ad-hoc networks, Defense Advanced Research Projects Agency (DARPA), DARPA-SN-13-35, April 2013

[2] P. Gupta and P. R. Kumar, "The capacity of wireless networks", IEEE Transactions on Information Theory, 46, 2, 2006, pp. 388-404.

[3] P. Jacquet, P. Muhlethaler, T. Clausen, A. Laouiti, A. Qayyum, and L. Viennot, "Optimized link state routing protocol for ad hoc networks", In Proceedings of IEEE INMIC, Pakistan, 2001.

[4] B. Bellur and R. Ogier, "A reliable, efficient, topology broadcast algorithm for dynamic networks", In Proceedings of IEEE INFOCOM, 1999

[5] J. J. Garcia-Luna-Aceves and M. Spohn, "Source tree routing in wireless networks", In proceedings of IEEE ICNP, 1999

[6] D. B. Johnson and D. Maltz, "Dynamic source routing in mobile ad-hoc networks", In Mobile Computing, Edited by T. Imielinski and H. Korth, Kluwer Academic publishers, 1995

[7] C. Perkins, "Ad-hoc, On demand distance vector routing", In proceedings of IEEE MILCOM, 1997

[8] S. Basagni, I. Chlamtac, V. R. Syrotiuk and B. A. Woodward, "A distance routing effect algorithm for mobility", In proceedings of IEEE / ACM Mobicom, 1998

[9] Z. Haas and M. Pearlman, "The performance of query control schemes for the zone routing protocol", In proceedings of ACM SIGCOMM, 1998

[10] P. Jacquett and L. Viennot, "Overhead in mobile ad-hoc network protocols", INRIA research report, France, June 2000

[11] A. S. Tanenbaum, "Computer Networks", Prentice-Hall, 1996

[12] G. F. Lawler and V. Limic. 2010, "*Random Walk: A Modern Introduction*". Cambridge University Press, New York, NY, USA

[13] J.-Y. Le Boudec and M. Vojnovic, "Perfect simulation and stationarity of a class of mobility models," In Proceedings of IEEE INFOCOM, 2005, pp. 2743–2754

[14] P. Nain, D. Towsley, B. Liu, and Z. Liu, "Properties of random direction models," In Proceedings of IEEE INFOCOM, 2005, pp. 1897–1907

[15] K. Parker, M. Sridharan, A. Arora, V. Kulathumani, J. brown, P. R. Kumar, R. Correa, H. Greenburg, T, Herman, D. Patru, G. Thaker, T. Winters, "Scaling Mobile Ad-hoc Networks: Alternate semantics for local routing combined with leveraging small amounts of global capacity", Final report, Fixed Wireless at a Distance program, Defense Advanced Research Projects Agency (DARPA), 2013

[16] J. Redi and R. Ramanathan, "The DARPA WNaN network architecture," In Proceedings of Military Communications Conference (MILCOM), 2011.

[17] R. Ramanathan, R. Allan, P. Basu, J. Feinberg, G. Jakllari, V. Kawadia, S. Loos, J. Redi, C. Santivanez and J. Freebersyse, "Scalability of mobile ad hoc networks: theory vs practice," In Proceedings of Military Communications Conference (MILCOM), 2010

[18] V. Manfredi, M. Crovella, and J. Kurose, "Understanding stateful vs stateless communication strategies for ad hoc networks", In *Proceedings of the 17th annual international conference on Mobile computing and networking* (MobiCom '11), 2011.

[19] A.A. Abouzeid and N. Bisnik, "Geographic Protocol Information and Capacity Deficit in Mobile Wireless Ad Hoc Networks," *IEEE Transactions on Information Theory*, 57(8): 5133-50, August 2011.

[20] M. Grossglauser and D. N. C. Tse, "Mobility increases the capacity of ad hoc wireless networks. *IEEE/ACM Trans. Netw.* 10(4), pp.477-486, 2002.

[21] G. Sharma, R. Mazumdar, and N. B. Shroff, "Delay and capacity trade-offs in mobile ad hoc networks: a global perspective", *IEEE/ACM Trans. Netw.* 15, 5, 981-992, 2007

VIII. Appendix A

In this section, we describe our experiments on the standard implementation of the Optimized Link State Routing (OLSR) to study its documented poor performance [4] in network sizes of 75-150 nodes. We used the default parameters of OLSR with a *hello* interval of 2 seconds, and *topology control* interval of 5 seconds.

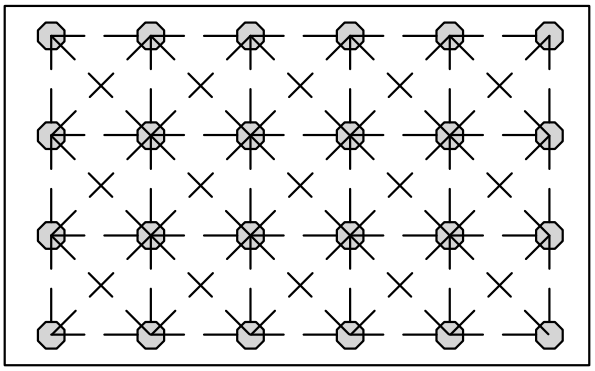

Figure 9: **The grid network used for initial investigation**.

Specifically, we constructed a grid network (**Figure 9** shows this topology) with varying sizes and introduced random link changes at a rate corresponding to low-to-medium mobility. Although a grid network is not a very realistic deployment scenario, it allows us to control very precisely the ground truth; that is, the number of neighbors and the rate of change of links, which would be somewhat more difficult to control precisely in a scenario with random mobility. Further, in order to isolate the effect of Network Layer Overhead (NLO) on scaling, the simulations were done in the absence of any useful data traffic. The expected result of these simulations was that when the network size reached around ~100 nodes, the system would start to fail and the NLO at that point would have reached between 15%-30% of the network capacity, thereby leaving very little capacity for useful data traffic. Thus, *our definition of the capacity wall for these simulations was NLO reaching 15%-20% of channel capacity.*

Counter to our expectations, for a simulation of 100 nodes the OLSR-NLO did not consume around 15% of the capacity. In fact, the NLO was accounting for around 1% of the physical capacity. But further investigation revealed that the routing performance in the network was indeed bad. The real problem turned out to be the percentage of broken paths. OLSR in fact was minimizing NLO at the expense of routing performance. To calculate the number of broken paths, we took periodic snapshots of the nodes routing tables every 100 milliseconds and traversed the path derived from the routing table to find

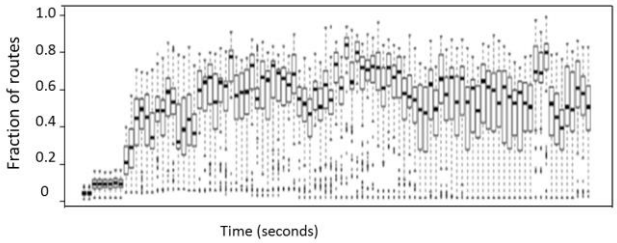

Figure 10: **Fraction of routes that were reachable during 100 seconds of simulation using OLSR with 100 nodes**.

the reachability of the nodes between every pair of nodes. Figure 10 shows the plot of the percentage of broken paths with time for OLSR. From the figure the reader can see that around 50-60% of the paths remain broken almost throughout the entire period of the simulation. These results show that the routing system failed well before the capacity limits were reached.