Directional Routing for Wireless Mesh Networks: A Performance Evaluation

Bow-Nan Cheng
ECSE Department
Rensselaer Polytechnic Institute
chengb@rpi.edu

Murat Yuksel
CSE Department
University of Nevada - Reno
yuksem@cse.unr.edu

Shivkumar Kalyanaraman
ECSE Department
Rensselaer Polytechnic Institute
shivkuma@ecse.rpi.edu

***Abstract*— Routing in multi-hop wireless networks involves the *indirection* from a *persistent name (or ID)* to a *locator*. One of the biggest issues in routing is providing adequate connectivity while scaling the network. Recently, [1] has attempted to mitigate this issue by using directional communication methods to find intersections between source-rendezvous and rendezvous-destination paths, providing effective routing in unstructured, flat networks. [1] showed that by "drawing" two lines orthogonal to each other at each node, it is possible to provide over 98% connectivity while maintaining only order O($N^{3/2}$) states. It is interesting, however to investigate what happens when additional lines are "drawn" and how that affects connectivity, path length and state complexity. In this paper, we examine how transmitting along one, two, three, and four lines affects routing and provide both analytical bounds for connectivity as well as packetized simulations on how these methods stack up in a more realistic environment. We show that by sending packets out in more directions, increased connectivity and smaller average path length results only up to a point. The trade-off, however, is added state information maintained at each node. We also show that in mobile environments, adding additional lines increases the chances for successful packet delivery only marginally.**[1]

I. Introduction

Routing in wireless ad-hoc networks have had to grapple with the twin requirements of connectivity and scalability. Early MANET protocols such as DSR [8], DSDV [6], AODV [7], among others, explored proactive and reactive routing methods which either flooded information during route dissemination or during route discovery respectively. While effective in providing high connectivity, as networks grow, however, flooding poses an obvious scalability problem. In response, several topology-based routing protocols such as OLSR [9], Hierarchical Routing [10], among others, have implemented limited flooding techniques to disseminate route information. Additionally, position-based routing paradigms such as GPSR [3] were also proposed to reduce the state complexity and control-traffic overhead by leveraging the Euclidean properties of a coordinate space embedding. These schemes require nodes to be assigned a coordinate in the system, and still require a mapping from nodeID to coordinate location.

[1]This material is based upon work supported by the National Science Foundation under Grant Nos. IGERT 0333314, ITR 0313095, and STI 0230787. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

A recent trend in wireless communications has been the desire to leverage directional forms of communications (e.g. directional smart antennas [12] [11], FSO transceivers [14]) for more efficient medium usage [11] [12] [13], routing [1], [2] and scalability. With the advent of fixed directional communications methods such as tesselated free-space-optical spheres and chained directional antennas which are highly directional in nature, it has become increasingly important to study how to incorporate directionality into routing schemes.

Recently, [1] has attempted to mitigate the issues of connectivity and scalability by using directional communication methods to find intersections between source-rendezvous and rendezvous-destination paths, providing effective routing in unstructured, fixed, flat mesh networks. [1] showed that by "drawing" two lines orthogonal to each other at each node, it is possible to provide over 98% connectivity while maintaining only order O($N^{3/2}$) states. It is interesting, however to investigate what happens when additional lines are drawn and how that affects connectivity, path length and state complexity. In this paper, we examine how communicating along one, two, three, and four lines affect routing and provide both analytical bounds for connectivity as well as packetized simulations on how these methods stack up in a more realistic environment.

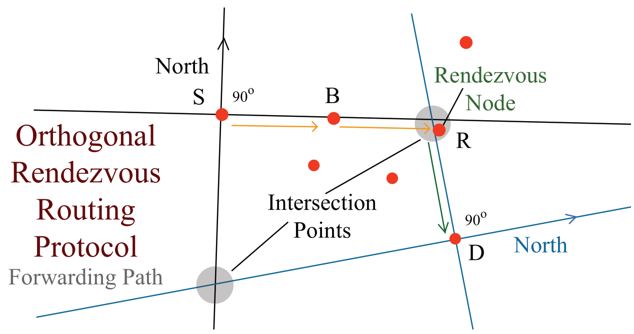

Fig. 1. ORRP Basic Example: Source sends packets to Rendezvous node which in turn forwards to Destination

Specifically, we will show that:

- Using the Multiplier Angle Method (MAM) heuristic suggested in [1], even only one line provides a high degree of connectivity in symmetric topologies.
- Addition of lines yields significantly diminishing returns from a connectivity-state maintenance perspective.

- Addition of lines yields better paths from source to destination.
- Although not the focus of the paper, as mobility is added into the equation, addition of lines yields only *marginally* better delivery successes.

The rest of the paper is organized as follows: Section II gives a brief introduction of Orthogonal Rendezvous Routing Protocol (ORRP) as well as extensions to the protocol to accommodate routing along additional lines. Section III provides some analysis to find connectivity upper bounds and path stretch without perimeter routing. Section IV provide performance evaluations in packetized simulations for each case and finally, section V concludes the paper.

II. Orthogonal Rendezvous Routing Protocol Extensions

The basic concept behind ORRP is simple: knowing that in 2-D Euclidian space, a pair of orthogonal lines centered at different points will intersect at two points at minimum, rendezvous points can be formed to forward packets as shown in Figure 1. To achieve this, ORRP relies on both a proactive element which makes up the "rendezvous-to-destination" path and a reactive element which builds a "source-to-rendezvous" route on demand. Nodes periodically send ORRP announcement packets in orthogonal directions and at each node along the orthogonal route, the node stores the route to the source of the ORRP announcement and the node it received the announcement from (previous hop). When a source node wishes to send to some destination node that it does not know the path for, it sends out a route request packet (RREQ) in its orthogonal directions and each subsequent node forwards in the opposite direction from which it receives the packet. Once a node containing a path toward the destination receives an RREQ, it sends a route reply packet (RREP) in the reverse direction back to the sender and data transmission begins.

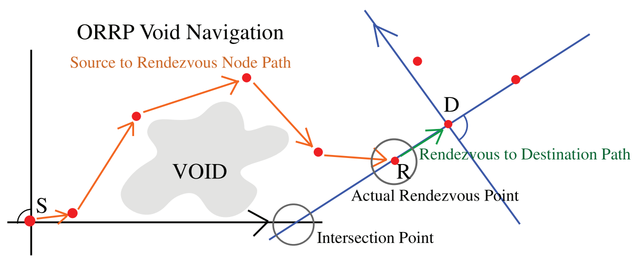

Fig. 2. Traversing voids in sparse networks with differing intersection points

To handle perimeter, void, and path deviation issues, ORRP implemented a Multiplier Angle Method (MAM) heuristic to navigate around voids, perimeters, and maintain relatively straight-line paths for announcement and RREQ packets as shown in figure 2. [1] showed that ORRP (2 lines) achieves connectivity with high probability even in sparse networks with voids, scales well without imposing GHT-like graph structures [19] (eg: trees, rings, torus etc), maintains a total state information of O($N^{3/2}$), evenly distributed for N-node networks, and does not resort to flooding either in route discovery or dissemination. The price paid by ORRP is suboptimality in terms of path stretch compared to shortest path, but [1] showed that the path stretch is small for generalized networks.

Because MAM allows for even the possibility of sending along one line to also achieve high connectivity (intersections outside of topology region would then be met along the perimeter), it is interesting to explore the tradeoff between the amount of state maintenance required to achieve similar reach statistics. In the same way, we are interested to see if addition of lines garners significant increases in reachability and better path selection. Extension of ORRP, therefore, is rather straight forward: instead of sending out interfaces that are orthogonal to each other (90^o from each other) as in ORRP, we send out announcement and RREQ packets out interfaces 180^o from each other for the "1 line" case, 60^o from each other for the "3 line" case, and 45^o from each other for the "4 line". All these cases are compared to the base orthogonal case.

TABLE I
Comparison of Reach Probability vs. Number of Lines

	1 Line (180°)	2 Lines (90°)	3 Lines (60°)
Circle (Radius 10m)	58.33%	99.75%	100%
Square (10mx10m)	56.51%	98.30%	99.99%
Rectangle (25mx4m)	34.55%	57%	57.61%

III. Analysis: Reachability and Path Stretch

Given a Euclidian area over which nodes are scattered, assuming no deviation correction with MAM, a source-destination pair cannot reach each other if all rendezvous points are outside the boundaries of the area. The general idea behind obtaining the reachability upper bound is to find intersections between lines drawn between the source and destination. In cases where all the intersections lie outside of the rectangular area for a particular source and destination oriented in a certain way, our analysis assumes that there is no path from source to destination. Notice that this analysis assumes that probe packets do not travel along perimeters of the Euclidian area under consideration and therefore inspects a worst-case upper bound on reachability.

Like in [1], our analysis begins with randomly selecting two source and destination pairs along with random orientations. We then formulate the equations of the lines generated by these two nodes and randomly selected orientations and find their intersection points. The equations of the lines will be different depending on whether we are looking at 1, 2, or 3 lines. If at least one of these intersection points lies in the boundaries of the topology, then we consider that particular source-destination pair as reachable. By iterating through all possible orientations for each possible source-destination pairs, we find a percentage of the total combinations that provide reachability vs. the total paths chosen. Because different Euclidian area shapes will no doubt yield different reachability requirements, we calculated the reachability probability for various area shapes by using Matlab in a grid network. Table I shows the reach probability vs. the number of lines used for calculations.

It can be seen that the addition of more lines yields significant gains from the one to two line case but only slight gain

TABLE II

COMPARISON OF PATH STRETCH VS. NUMBER OF LINES

	1 Line (180°)	2 Lines (90°)	3 Lines (60°)
Circle (Radius 10m)	3.854	1.15	1.031
Square (10mx10m)	4.004	1.255	1.039
Rectangle (25mx4m)	4.73	3.24	1.906
Grid (No bounds)	1.323	1.123	1.050

afterwards. Particular interest is given to the rectangular case where even with three lines, the raw reach probability is very low. We suspect the reason for this is the slim shape yielding to much more path intersections outside of the topology area. [1] showed that most of the unreach happens at the topology perimeters and even with additional lines, these perimeter nodes need a very high degree of angular match between lines before a path can be made. The result is that by adding only 30° more to match on, the angle of incidence is still too high to find an intersection within the area.

A similar analysis is done to find path stretch. If a source and destination pair has a line intersection within the topology boundaries, the shortest total distance (from source to intersection point and intersection point to destination) is selected as the path. This distance is divided by the distance between the source and destination to obtain a path stretch. In cases where there is no intersection inside the topology boundaries, we simply add the distance of the perimeter as that is the maximum path we can obtain with MAM. Table II gives the Matlab calculated path stretch for 1, 2, and 3 lines.

Table I and II show the reachability and path stretch simulation results for 1-3 lines all equidistantly separated from each other. While for reach probability, the affect from one to two lines is dramatic, it can be seen that very little gain is achieved by adding additional lines. In the case of path stretch, however, the addition of additional directions to send announcement and RREQ packets result in much better path selection as more packet interceptions occur. We suspect that in sparser networks or networks with voids, the gains would be negligible as control packets would take similar paths with MAM. It is important to note that with MAM, almost all the corner case reach issues can be resolved with only 2 lines.

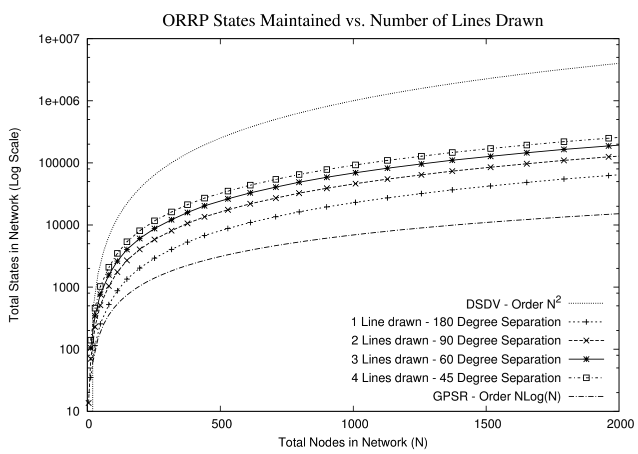

Fig. 3. Total states maintained in network with respect to the number of transmission lines used. As number of lines increase, the number of states maintained throughout network increases.

Figure 3 demonstrates the potential increase in state maintenance needed with the addition of transmission lines. While increasing steadily, it is still much less than order N^2.

IV. PERFORMANCE EVALUATION

In this section, we will evaluate the metrics of *reach probability*, *average path length*, *total state maintenance*, and *aggregate throughput* under conditions of varying *network densities*, *network topologies*, *void conditions*, and *basic random waypoint mobility*. Unless otherwise noted, all simulations were performed using Network Simulator [16] with default simulation parameters listed in table III Interfaces were setup so that they are all aligned equally spaced radially from a single point (the node) with the transmission and receive angle for each interface equal. Adding all the transmission angles together provided for omnidirectional coverage. For example, a node with 24 interfaces would have a transmit/receive and interface separation angle of 15°. In the same way, a node with 4 interfaces would have a transmit/receive and interface separation angle of 90°. Unless otherwise noted, all nodes are outfitted with 24 interfaces and simulation results averaged over 30 runs each under random node orientation.

TABLE III

DEFAULT SIMULATION PARAMETER

Parameter	Values
Transmission Radius	60m
Number of Interfaces	24
TTL for Control Pkts	10
Topology Boundaries	300m x 300m
Announcement Interval	2.0s
Route Timeout	10s
Simulation Time	50s
Mobility	None

A. *Affect of Additional Lines on Various Topologies*

Section III showed that under differing topologies without any angle correction, connectivity and path stretch is drastically affected by number of lines used for transmissions. It is interesting, therefore, to see how the analysis matches up with packetized simulations with angle correction. We suspected that even with one line, MAM should be able to deal with the majority of perimeter nodes and therefore provide fairly high reachability in symmetric topologies. In asymmetric topologies, however, as the "incident angle" a packet hits a perimeter node becomes steeper and steeper, it becomes more difficult to do angle correction since we set a hard limiter to not forward more than 90° to avoid loops so we suspect in these topologies, additional lines will affect reach probability more drastically.

In the same way, because additional lines provide additional paths to choose from, we expect that as the number of lines increase, the average path length from source to destination will decrease. Table IV outlines the simulation parameters that differ from the default and figure 4 and figure 5 show our results

As illustrated in figure 4, for square topologies, there is a large gain in reach probability going from one line to two

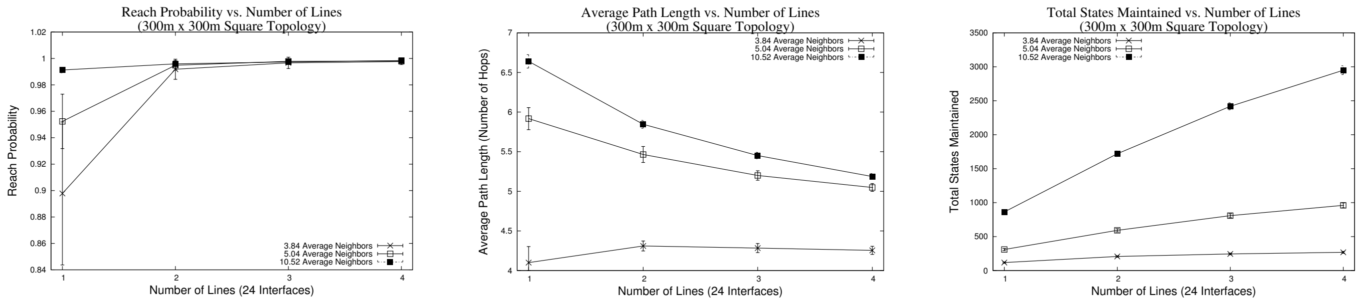

Fig. 4. Reach probability, total states maintained, and average path length vs. number of lines used for transmissions for dense and sparse with no voids present. As expected, as number of lines increased, the reach probability and total states maintained increased while average path length decreased.

TABLE IV

SIMULATION PARAMETERS: ADDL. LINES ON VARIOUS TOPOLOGIES

Parameter	Values	
TTL for Control Pkts	10	15, 20
Topology Boundaries	300m x 300m	1000m x 200m
Number of Nodes	25, 50, 100	75, 100
Average Number of Neighbors	3.84, 5.04, 10.52	3.6, 5.48

lines but the gain thereafter is small even for varying network densities. Average path length, as well, seems to trail off after transmitting orthogonally with two lines. This is expected as even in our analysis, path stretch was close to shortest path even for two lines. In contrast to this, states maintained at each node increased seemingly linearly with increased number of lines. This is expected as more states need to be maintained along linearly increasing number of lines of transmission.

We saw very similar results for rectangular topologies except that the jump from two to three lines provided a larger jump in reach probability. Even with just one line, MAM was able to ensure roughly 67% packet delivery success as compared to the 34.55% shown in our analysis. By increasing the number of lines, additional paths were available despite the rather "thin" topology. Figure 5 showed that the average path length curve mimicked the reach probability curve. At first this seems counter intuitive since one would expect that with additional lines and thus, additional paths to choose from, the average path length would be less as lines are increased. However, it is important to note that our simulations only calculate average path length based on *successful* transmissions. Thus, nodes at the edges of the rectangular topology, which would most likely incur the highest number of hops to reach, would be left out if no path is found. This is therefore consistent with our hypothesis and as expected, total states maintained in the network grew fairly linearly with increased number of lines.

B. Affect of Number of Lines on Network Voids

It is interesting to see how the number of lines of transmission affect reachability and path length in networks with large voids. We hypothesized that while reach would increase with increased number of lines, average path length would remain fairly constant. This is due to few paths to choose from to navigate around voids and therefore, as long as there is a path, most likely, that path would be the one chosen. Our simulation parameters are listed in table V.

TABLE V

SIMULATION PARAMETERS: ADDL. LINES ON NETWORKS WITH VOIDS

Parameter	Values
Number of Nodes	25, 50
Average Number of Neighbors	3.92, 6.2

Figure 6 shows our results for various lines on networks with voids. As expected, the increase from one to two lines yielded a fairly large connectivity gain as well as increased total states maintained network-wide. Average path length, as expected, remained fairly constant. This was due to relatively few paths to choose from to navigate around voids and therefore fairly consistent path choices were made in the connected network.

C. Affect of Number of Lines on Throughput

One of the key metrics in wireless mesh networks is network throughput. In wireless networks, throughput is dependent on a lot of factors like congestion, link quality, etc., which unfortunately become increasingly difficult to simulate. In this section, we try to understand the affect of transmitting along additional lines affect throughput. It is expected that with shorter paths and higher reachability, average throughput network-wide will increase. Table VI gives our simulation parameters and figure 7 illustrate our results.

TABLE VI

SIMULATION PARAMETERS: ADDITIONAL LINES ON THROUGHPUT

Parameter	Values
Number of Nodes	100
Average Number of Neighbors	10.52
Number of Random Connections	100
CBR Packet Size	512 KB
Transmission Duration	10.0 seconds

Our results in Figure 7 show that throughput increases with increase in lines. Looking at the reach and average path length graphs, this result is intuitive: with smaller reach probability, packets are not successfully delivered and with higher average path length, the delivery time increases dramatically. In short, the increase in lines of transmission lead to paths that are closer to shortest path, which lead to higher throughput. It is interesting to note that even with higher packet delivery success, higher throughput is not guaranteed.

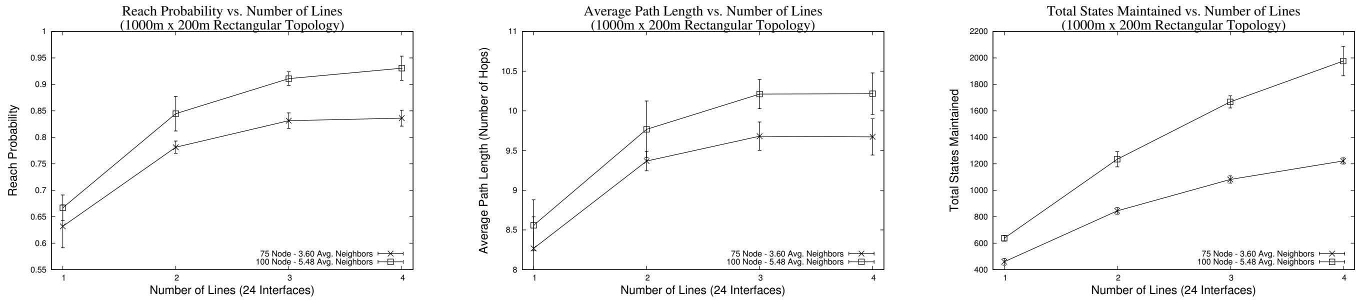

Fig. 5. Reach probability, total states maintained, and average path length vs. number of lines used for a rectangular topology. Reach is drastically affected by additional lines due to better paths in a slim topology.

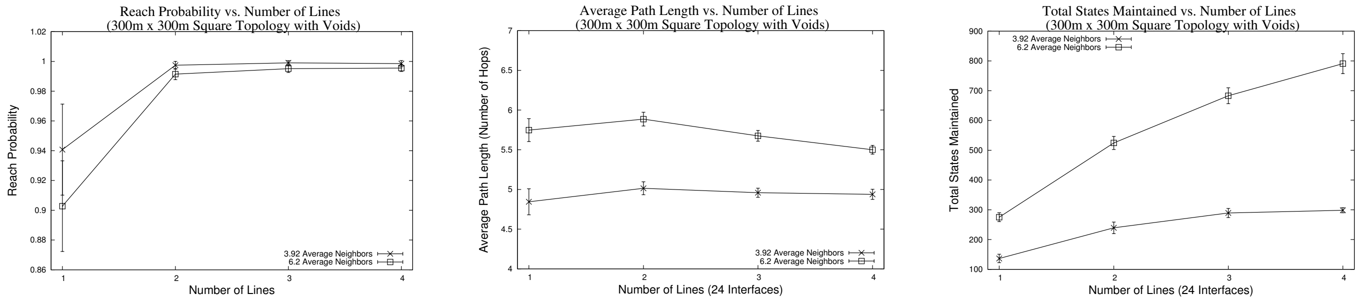

Fig. 6. Reach probability, total states maintained, and average path length vs. number of lines used for transmission for dense and sparse topologies with large voids present. As expected, with voids present, paths taken should be relatively equal due to less choices. At the same time, as more lines are used, the reach probability and total states maintained increased.

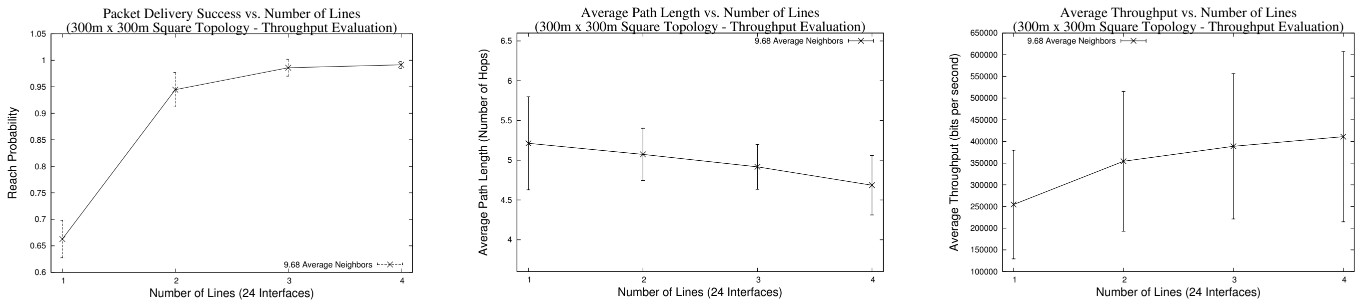

Fig. 7. Average throughput increases as number of lines increase. It can be seen that throughput is largely dependent on average hop count: as average hops increase, the throughput drops.

D. *Affect of Number of Lines on Varying Network Mobility*

Because ORRP was designed primarily for fixed wireless mesh networks, it is expected to fail under mobility because lines cannot be maintained in an efficient manner. Adding additional lines, however, could lead to better paths and increased delivery success even in mobile and/or disruption tolerant environments. In this section, we seek to understand whether addition of lines helps in a mobile environment. We suspect that the addition of lines should not affect reach probability much because all paths are moving. Table VII gives our simulation parameters and figure 8 show our results.

Our results in figure 8 show that for a mobile network, directional routing protocols like ORRP have severe issues without decreasing the announcement interval and route time-out. However, there seems to be a fairly large increase in reach probability as number of lines increased from 1 to 2 but the gains trail off afterwards. We attribute this increase to having additional and better paths to choose from which in-turn lead to less number of hops and less number of nodes that have moved away providing for a higher reach probability. In the same way, average path length, as expected, decreased with additional lines as better path options were available.

TABLE VII
SIMULATION PARAMETERS: ADDITIONAL LINES ON MOBILE NETWORKS

Parameter	Values
Number of Interfaces	12
Topology Boundaries	300m x 300m
Number of Nodes	100
Mobility	RWP Model: 2.5m/s, 5.0m/s, 7.5m/s
Simulation Time	100s
Connectivity Sampling Frequency	Every 20s

V. Future Work and Conclusion

In this paper, we extended Orthogonal Rendezvous Routing Protocol (ORRP) to send packets out additional directions to measure the tradeoff between delivery success, average path length, total states maintained, and aggregate throughput. Our

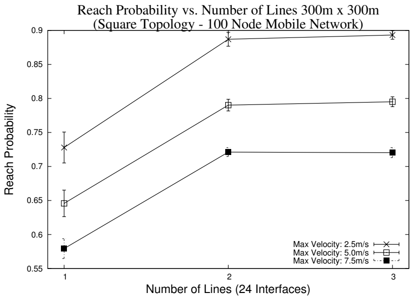

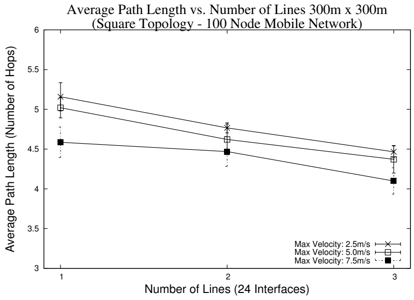

Fig. 8. Reach probability and average path length vs. number of lines used for transmission for mobile networks. ORRP was never designed for mobility but as it can be expected, addition of lines increases reach probability. This is most likely due to shorter paths chosen.

analysis in section III showed that the jump between one line and two lines yields significant increases in reach probability and path stretch while the addition of more lines gives only *marginal gains* in reach probability but should choose much better paths resulting in smaller path stretch. Because the analysis was performed with straight line paths without angle correction deviations, packetized simulations were necessary.

We simulated the affect of number of lines of transmission had on reach probability, average path length, total states maintained network-wide, and aggregate throughput on various topologies, network densities, void conditions, and mobility. Our results indicated that in non-void, non-mobile scenarios, there is a significant increase in delivery success and throughput from one to two lines but as suggested by our analysis, the gains after adding additional lines are slim. Average path length was also shown to decrease until shortest path was almost reached in increasing number of lines. Additionally, as the number of lines increased, total states maintained in the network increased fairly linearly (but still order $N^{3/2}$. As voids were added, however, average path length remained fairly constant due to similar paths taken despite seemingly more paths to choose from. With mobility, it was shown that the addition of lines had very little affect on delivery success but dropped average path length marginally as expected.

Overall, the addition of lines yields only marginal gains over the two orthogonal lines scenario suggested in [1] and it would be interesting to explore additional methods for deviation correction, perimeter routing, and void traversals to account for the few percentage of unsuccessful packets delivered. Furthermore, since ORRP fails drastically in mobile environments even with decreased announcement intervals and route lifetime, it would be interesting to look at the possibility of extending ORRP to mobile adhoc networks.

REFERENCES

[1] Bow-Nan Cheng, Murat Yuksel, Shivkumar Kalyanaraman, "Orthgonal Rendezvous Routing Protocol for Wireless Mesh Networks," Proceedings of the 17th IEEE International Conference on Network Protocols (ICNP), November 2006.

[2] Hrishikesh Gossain, Tarun Joshi, Carlos De Morais Cordeiro, and Dharma P. Agrawal, "DRP: An Efficient Directional Routing Protocol for Mobile Ad Hoc Networks," IEEE Trans. Parallel and Distributed Systems, vol. 17, no. 12, 2006, pp. 1439-1451.

[3] B. Karp and H.T. Kung, "GPSR: Greedy Perimeter Stateless Routing for Wireless Networks," Proceedings of ACM/IEEE MOBICOM, Boston, MA, August 2000.

[4] J. Li, J. Jannotti, D.S.J. De Couto, D.R. Karger, and R. Morris, "A Scalable Location Service for Geographic Ad Hoc Routing," Proceedings of ACM/IEEE MOBICOM, pp. 120-130, Aug. 2000.

[5] John Bicket, Daniel Aguayo, Sanjit Biswas, and Robert Morris, Architecture and Evaluation of an Unplanned 802.11b Mesh Network, *ACM/IEEE MOBICOM* 2005.

[6] Charles Perkins and Pravin Bhagwat. Highly Dynamic Destination-Sequenced Distance-Vector Routing (DSDV) for Mobile Computers. Proceedings of ACM SIGCOMM, pages 234-244, 1994.

[7] Charles E. Perkins and Elizabeth M. Royer. "Ad hoc On-Demand Distance Vector Routing." Proceedings of the 2nd IEEE Workshop on Mobile Computing Systems and Applications, New Orleans, LA, February 1999, pp. 90-100.

[8] David B. Johnson, David A. Maltz, Josh Broch. DSR: The Dynamic Source Routing Protocol for Multi-Hop Wireless Ad Hoc Networks. *Ad Hoc Networking*, edited by Charles E. Perkins, Chapter 5, pp. 139-172, Addison-Wesley, 2001.

[9] T. Clausen and P. Jacquet. OLSR RFC3626, October 2003. http://ietf.org/rfc/rfc3626.txt.

[10] W. T. Tsai , C. V. Ramamoorthy , W. K. Tsai, O. Nishiguchi, An Adaptive Hierarchical Routing Protocol, *IEEE Transactions on Computers*, v.38 n.8, p.1059-1075, August 1989.

[11] Romit Roy Choudhury and Nitin Vaidya. "Impact of Directional Antennas on Ad Hoc Routing", Proceedings of the Eighth International Conference on Personal Wireless Communication (PWC), Venice, September 2003.

[12] R. Ramanathan. "On the performance of ad hoc networks using beamforming antennas", Proceedings of ACM MOBIHOC, October 2001.

[13] Su Yi, Yong Pei, and Shivkumar Kalyanaraman, "On the Capacity Improvement of Ad Hoc Wireless Networks Using Directional Antennas," Proceedings of ACM MOBIHOC, Pages 108-116, Annapolis, MD, June 2003.

[14] Jayasri Akella, Chang Liu, David Partyka, Murat Yuksel, Shiv Kalyanaraman, and Partha Dutta, "Building Blocks for Mobile Free-Space-Optical Networks," Proceedings of IFIP/IEEE WOCN, pages 164-168, Dubai, United Arab Emirates, March 2005.

[15] D. Braginsky and D. Estrin, "Rumor routing algorithm for sensor networks," Proceedings of WSNA, Atlanta, GA, October 2002.

[16] The Network Simulator. ns-2. http://www.isi.edu/nsnam/ns.

[17] Romit Roy Choudhury, Nitin H. Vaidya "Performance of Ad Hoc Routing using Directional Antennas" Journal of Ad Hoc Networks - Elsevier Publishers, November, 2004.

[18] Amit Kumar Saha and David B. Johnson. Routing Improvements Using Directional Antennas in Mobile Ad Hoc Networks. IEEE Globecom, Dallas, Texas, November 2004.

[19] S. Ratnasamy, B. Karp, L. Yin, F. Yu, D. Estrin, R. Govindan, and S. Shenker, "GHT: A geographic hash table for data-centric storage in sensornets," In Proceedings of the First ACM International Workshop on Wireless Sensor Networks and Applications (WSNA) 2002.

[20] S. Kalyanaraman S. Yi, Y. Pei and B. Azimi-Sadjadi, "How is the capacity of ad hoc networks improved with directional antennas," Accepted for Publication in Wireless Networks, 2005.

[21] R. R. Choudhury, X. Yang, R. Ramanathan and N. H. Vaidya, "Using Directional Antennas for Medium Access Control in Ad Hoc Networks," In Proceedings of MobiCom, pp. 59-70, Sep. 2002.