Rendezvous-Based Directional Routing: A Performance Analysis (Invited Paper)

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*— The increased usage of directional methods of communications to improve medium reuse, network capacity, and bandwidth has prompted research into leveraging directionality in every layer of the network stack. Recently, there has been work on bringing the apparent capacity gains on layer 2 using directional communications methods to layer 3 by using directionality to route packets scalably in unstructured, flat networks. In their protocol, Orthogonal Rendezvous Routing Protocol, Cheng et al. [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 O($N^{3/2}$) evenly distributed states at a cost of only 1.2 path stretch. In this paper, we seek to provide more in-depth performance analysis by tuning additional factors such as the number of directions to transmit, the number of interfaces per node, among others, to understand its affect on varying network densities, topologies, connections, and traffic patterns. We show that by sending packets out in more directions, increased connectivity, smaller average path length, better goodput results only up to a point as compared to other routing protocols. The trade-off, however, is added state information maintained at each node and additional control packets received. We also show that the addition of more interfaces generally yields better packet delivery success, average path length, and goodput.** [1]

I. Introduction

A recent trend in wireless communications has been the desire to leverage directional forms of communications (e.g. directional smart antennas [13], Free-Space-Optical transceivers [26], and sector antennas) for more efficient medium reuse, increased scalability, enhanced security and potential for higher achievable bandwidth. Previous work in directional antennas focused heavily on measuring network capacity and medium reuse [13] [14]. In these works, it was shown that with proper tuning, capacity *improvements* using directional over omnidirectional antennas are dramatic - ranging from a factor of $\frac{2\pi}{\sqrt{\alpha\beta}}$ for planned networks to a factor of $\frac{4\pi^2}{\sqrt{\alpha\beta}}$ for random networks where α and β are the beamwidths of the transmitting and receiving antennas.

[1]An extended abstract of this paper is to appear in LANMAN 2007 [2]. 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.

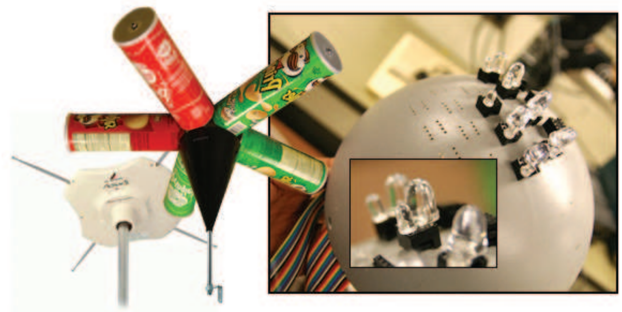

Fig. 1. Wireless directional communications methods such as directional antennas and free-space-optical transceivers have become increasingly available.

Additionally, there has been a large push in the free space optical (FSO) community to use FSO to compliment traditional RF methods in the wireless mesh context [27]. Currently available in point-to-point links in terrestrial last mile applications and in infrared indoor LANs [23] [22], FSO has several attractive characteristics like *(i) dense spatial reuse*, *(ii) low power usage*, *(iii) license-free band of operation*, and *(iv) relatively high bandwidth compared to RF*. Conversely, FSO suffers from *(i) the need for line of sight (LOS) alignment between nodes* and *(ii) reduced transmission quality in adverse weather conditions*. Yuksel et al. [26] proposed several ways to mitigate these issues by tessellating low cost FSO transceivers in a spherical fashion (see Figure 1) and replacing long-haul point-to-point links with short, multi-hop transmissions.

Given the seemingly large increases in medium reuse and potential for higher bandwidth in directional forms of communications, it becomes interesting to investigate how directionality can be used to facilitate and even improve wireless networks in all layers of the stack. The leveraging of directionality in wireless communications (e.g. directional smart antennas [13] [12]) for more efficient medium usage [12] [13] [14], routing [1] [3] and scalability in prior work has laid much of the foundations for extending directional communications to FSO.

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, state complexity, control packet overhead, and aggregate goodput. 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.

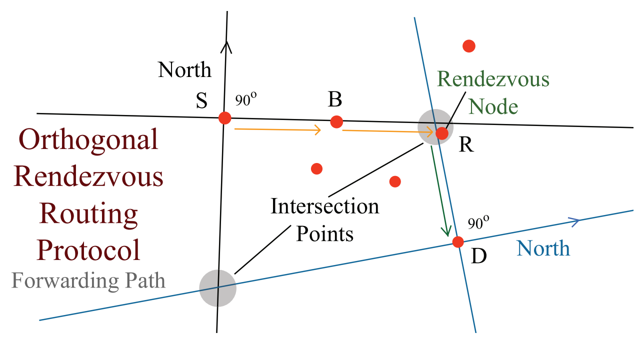

Fig. 2. ORRP Basic Example: Source sends packets to Rendezvous node which in turn forwards to Destination

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 as compared to our analytical bounds without MAM.
- In asymmetric topologies (e.g. rectangular) and using the Multiplier Angle Method (MAM) heuristic suggested in [1], increasing the number of lines yields better reach probability and average path lengths.
- Addition of lines yields significantly *diminishing returns* from a connectivity-state maintenance perspective.
- Addition of lines yields better paths from source to destination.
- Addition of lines yields better aggregate goodput overall and about 20x more goodput than DSR and AODV.
- Increasing the *number of interfaces* per node yields better results for reachability, average path length, and average goodput up to a certain point that is determined by network density.
- As *number of continuous flows* increase, ORRP with increased lines delivers more packets successfully utilizes the medium much more efficiently resulting in higher goodput network-wide.

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 expected 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 2. 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.

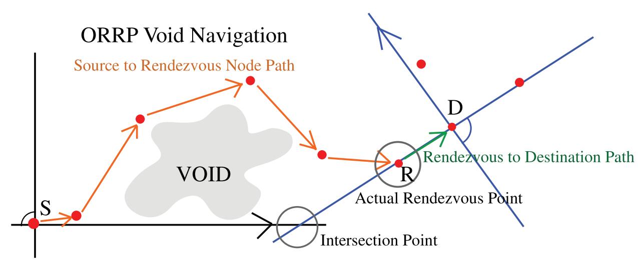

Fig. 3. Traversing voids in sparse networks with differing intersection points

To handle perimeter, void, and path deviation issues, ORRP implements 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 3. Cheng et al. [1] showed that ORRP (2 lines) achieves connectivity with high probability even in sparse networks with voids, scales well without imposing DHT-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 sub-optimality 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 redirected 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 (2 line) case.

III. Numerical 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 presents 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 within the boundaries of the area, 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.

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%	67.61%

It can be seen that the addition of more lines yields significant gains from the one to two line case but only slight gain 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. Cheng et al. [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. It is important to understand why increasing from 2 to 3 lines only marginally increases the reach for the rectangular case. In the rectangular case, the width is around 6x more than the height. Thus, even if there are more lines, the probabilities of reach for source and destination nodes at the outer most edges remains fairly low while only those in the center benefit from additional lines.

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 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

Table I and II show the reachability and path stretch numerical analysis 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.

Figure 4 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*, *packet delivery success*, *end to end latency* and *aggregate goodput* under conditions of varying *network densities*, *number of interfaces*, *network topologies*, *network saturation*, transmission rates, and *void conditions*. Unless otherwise noted, all simulations were performed using Network Simulator [16] with n interfaces (divisible by 4) and each interface having a beamwidth of $360/n$ degrees. Unless otherwise noted, all nodes are

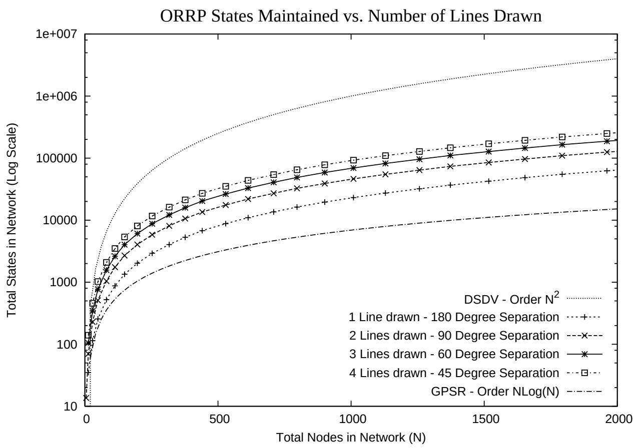

Fig. 4. Total states maintained in network with respect to the number of transmission lines used. As number of lines increase, the number of states maintained throughout network increases.

outfitted by default with 24 interfaces and simulation results averaged over 10 runs each under random node orientation.

A. *Simulation Environment Specifics*

Default NS2 simulation parameters are listed in Table III. For evaluating the effect of additional lines on various topologies and network voids (Sections IV-B-IV-C), 24 interfaces were used as 24 interfaces allowed for evaluating 1-4 lines (needing 2-8 transceivers respectively to send). In these sections, because we were only interested in determining reach probability, average path length, and total states maintained network-wide, it was more important to check the connection from every node in the network to every other node. To do this, each node simply sent a short burst (1-2 CBR packets) to every other node in the network. Reach probability was measured by the number of received vs. sent CBR packets and average path length was calculated by averaging the number of hops from source to destination. In these subsection, total states maintained network-wide were calculated by measuring the size of each routing table (with each entry counted as a single state) *before* any CBR packets are sent and totaling the associated values.

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

Network voids in Section IV-C were generated by taking a fully connected 100 node network and “removing” nodes using scripts that took inputs to an elliptical area and removed all nodes in that area. Two voids are present in both void networks evaluated. In evaluating total control packets, average path length, aggregate network goodput and end to end packet delay in Sections IV-D-IV-F connection patterns were generated by randomly choosing a source and destination. Simulations were run over 10 trials and results averaged with standard deviations given in graphs. For evaluating effect of transmissions rate on network capacity, CBR connections were made from every node to every node and the rate was increased steadily to find the aggregate network capacity. It’s important to note that even though we show that there are high gains in reachability and goodput using ORRP with only 8 interfaces, the reason we choose to evaluate ORRP against AODV and DSR with 12 interfaces is that 12 is easily divisible by the number of directions to transmit (2, 4, and 6) corresponding to the number of lines.

B. *Effect 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, 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 5 and Figure 6 show our results

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

As illustrated in Figure 5, for square topologies, there is a large gain in reach probability going from one line to two 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

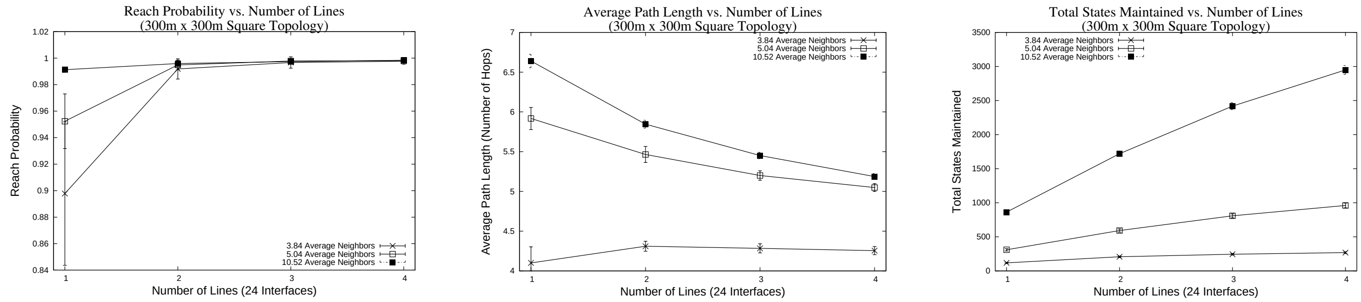

Fig. 5. Reach probability, total states maintained, and average path length vs. number of lines used for transmissions for dense and sparse with no voids present. As expected, as number of lines increased, the reach probability and total states maintained increased while average path length decreased.

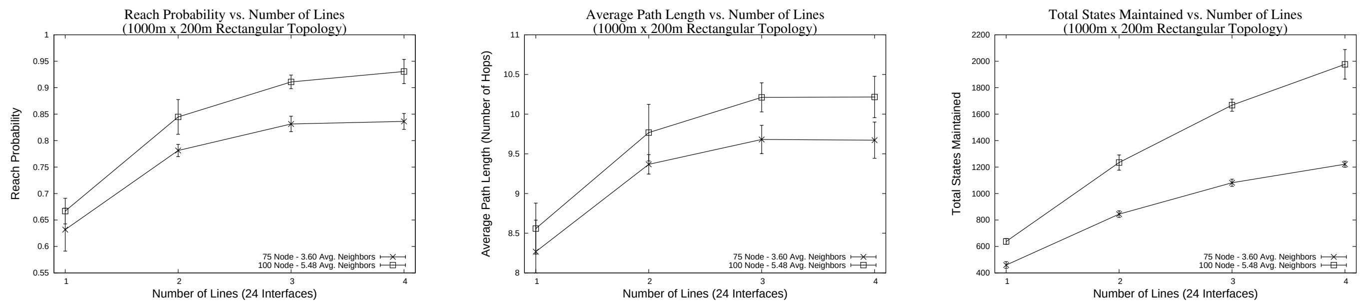

Fig. 6. Reach probability, total states maintained, and average path length vs. number of lines used for a rectangular topology. Reach is drastically affected by additional lines due to better paths in a slim topology.

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 6 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.

C. Effect of Number of Lines on Network Voids

It is interesting to see how the number of lines of transmission effect 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 7 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.

D. Effect of Varying the Number of Interfaces

Adding more interfaces to a node increases the diversity of directions to send with the finer granularity of spread resulting in less neighbors associated with a single interface. It is expected that the gains in delivery success, average path length, and goodput will increase with the number of interfaces up until a point and that this point is determined by the network density. Table VI lists our simulation parameters that differ from the default and Tables VII-IX give our results. Because it is important to transmit symmetrically (i.e. the angles between each transmission interface must be equal), certain number of interfaces can only transmit along 1, 2, 3 lines while others can only transmit along 1, 2, 4 lines. The N/A values in the tables represent the cases when transmission is not possible.

TABLE VI

SIMULATION PARAMETERS: AFFECT OF NUMBER OF INTERFACES

Parameter	Values
Number of Nodes	100 (Avg Neighbors: 10.52)
Number of Random Connections	100
CBR Packet Size	512 KB
Transmission Duration	10.0 seconds
Number Interfaces	8, 12, 16, 24
Simulation Time	100s

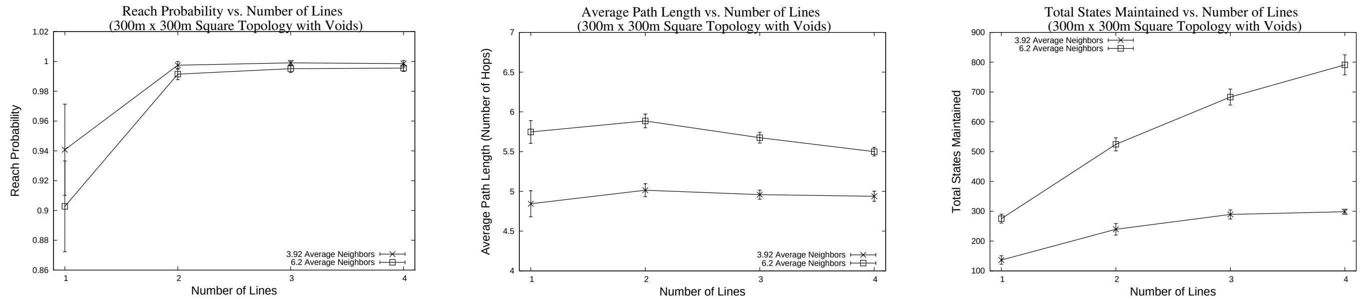

Fig. 7. Reach probability, total states maintained, and average path length vs. number of lines used for transmission for dense and sparse topologies with large voids present. As expected, with voids present, paths taken should be relatively equal due to less choices. At the same time, as more lines are used, the reach probability and total states maintained increased.

TABLE VII
DELIVERY SUCCESS VS. NUMBER OF INTERFACES

	1 Line	2 Lines	3 Lines	4 Lines
8 Interfaces (Avg NBs: 9.7)	64.7%	88.3%	N/A	98.0%
8 Interfaces (Avg NBs: 20.9)	65.4%	93.8%	N/A	98.4%
12 Interfaces (Avg NBs: 9.7)	65.3%	93.1%	98.3%	N/A
12 Interfaces (Avg NBs: 20.9)	71.6%	97.3%	99.1%	N/A
16 Interfaces (Avg NBs: 9.7)	65.7%	94.6%	N/A	98.7%
16 Interfaces (Avg NBs: 20.9)	76.4%	98.4%	N/A	99.7%
24 Interfaces (Avg NBs: 9.7)	67.2%	95.6%	99.0%	99.4%
24 Interfaces (Avg NBs: 20.9)	77.2%	99.6%	99.9%	99.9%

Table VII shows the packet delivery success for varying number of interfaces and network densities. It can be seen that in general, when number of interfaces increases, there is a large effect on delivery success going from 8 to 12 interfaces for a network density of 9.7 average neighbors and 12 to 16 for a network density of 20.9 average neighbors. Afterwards, the gains taper off. It is interesting to note that a network density of 9.7 average neighbors per node equates to approximately 1 neighbor per interface. It makes sense that the affect on delivery success would be most affected by the network density as there is approximately one node per network interface. The lower the number of interfaces, the more neighbors are associated with a specific interface and therefore, there is higher risk of announcement and RREQ packets "missing" each other. Additionally, "matching" one neighbor to a specific interface allows MAM to operate to the best efficiency because it can be consistent when choosing random nodes to send to in a specific direction.

TABLE VIII
AVERAGE PATH LENGTH (# OF HOPS) VS. NUMBER OF INTERFACES

	1 Line	2 Lines	3 Lines	4 Lines
8 Interfaces (Avg NBs: 9.7)	5.29	5.43	N/A	4.71
8 Interfaces (Avg NBs: 20.9)	6.89	6.12	N/A	5.22
12 Interfaces (Avg NBs: 9.7)	5.35	5.20	4.62	N/A
12 Interfaces (Avg NBs: 20.9)	6.71	6.18	5.34	N/A
16 Interfaces (Avg NBs: 9.7)	5.69	5.13	N/A	4.71
16 Interfaces (Avg NBs: 20.9)	6.56	5.81	N/A	4.48
24 Interfaces (Avg NBs: 9.7)	5.18	5.21	4.80	4.44
24 Interfaces (Avg NBs: 20.9)	6.28	5.50	4.64	4.44

As can be seen from Table VIII, as number of interfaces increase, the average path length generally decreases. The affect is more noticeable with denser networks and more lines as having more interfaces increases the granularity of neighbors associated with a specific interface. This refines the neighbor selection and allows for better paths. Because increase in node density leads to *shorter* distances to neighbors and *more* hops to go from source to destination, it makes sense that with *less* interfaces (more neighbors associated with a specific interface), paths chosen would be worse.

TABLE IX
THROUGHPUT (KBPS) VS. NUMBER OF INTERFACES

	1 Line	2 Lines	3 Lines	4 Lines
8 Interfaces (Avg NBs: 9.7)	71.2	92.7	N/A	200.5
8 Interfaces (Avg NBs: 20.9)	44.0	77.1	N/A	155.9
12 Interfaces (Avg NBs: 9.7)	130.4	181.4	321.0	N/A
12 Interfaces (Avg NBs: 20.9)	60.5	135.1	246.7	N/A
16 Interfaces (Avg NBs: 9.7)	144.0	187.5	N/A	387.0
16 Interfaces (Avg NBs: 20.9)	79.0	192.8	N/A	508.4
24 Interfaces (Avg NBs: 9.7)	213.9	300.8	407.7	503.2
24 Interfaces (Avg NBs: 20.9)	123.6	453.6	723.4	666.9

Table IX shows the throughput vs. number of interfaces. It can be seen that throughput for denser networks is generally smaller for the same number of interfaces because denser networks incur additional hops from source to destination. However, as number of interfaces and number of lines increase, it seems that throughput becomes much better with denser networks. We suspect this is not only due to better paths, but also less interference between transmissions as more interfaces localize affected nodes better.

E. Network Density Evaluation vs. AODV and DSR

It is interesting to understand how network density affects packet delivery success, average path length, total control packets, and average throughput network-wide for ORRP with multiple lines compared to other routing protocols like AODV and DSR. It is expected that with broadcast protocols that use omni-directional antennas such as AODV and DSR, as density increases, less packets will be delivered resulting in lower goodput. Table X gives the simulation parameters which differ from the default and Figures 8 and 9 show our results.

As can be seen from Figure 8, as network density (number of nodes) is increased, ORRP generates a higher number of control packets with increasing number of lines. Because ORRP is a proactive/reactive hybrid, it makes sense that under conditions of no node mobility, ORRP sends more control

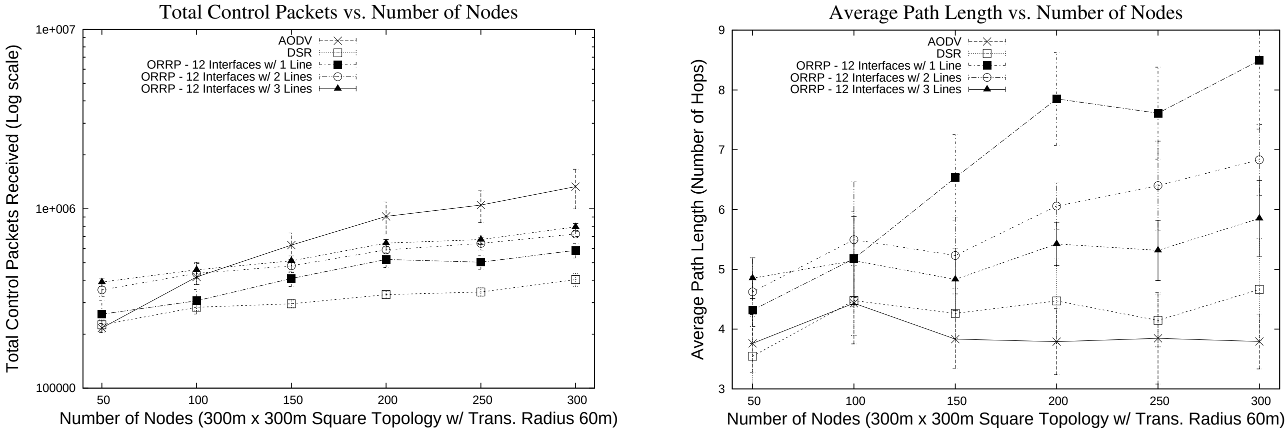

Fig. 8. Control packets received and average path length vs. number of nodes in the network for various routing protocols. ORRP with more lines sends out more control packets as expected. Because ORRP is a hybrid proactive/reactive protocol, it is expected to disseminate more packets than DSR and AODV in non-mobile situations. Also, ORRP delivers more packets with shorter average path length for additional lines. Because AODV and DSR utilize shortest path algorithms, it is expected that average path length is smaller compared to ORRP.

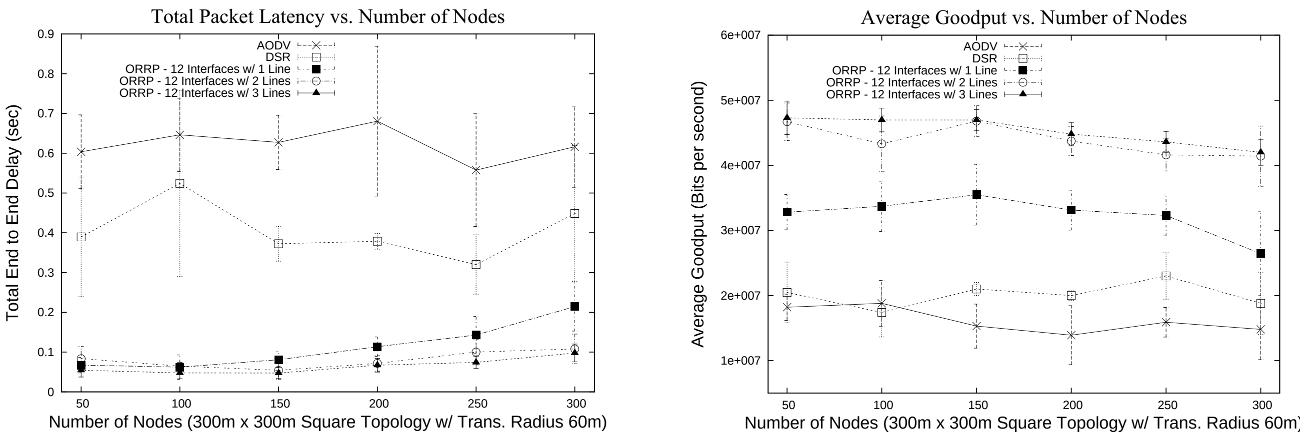

Fig. 9. Average end to end delay and average throughput vs. number of nodes in the network for various routing protocols. ORRP with more lines generates better paths and therefore incurs smaller end to end delay. The delay is much less than DSR or AODV because those protocols utilize the medium less efficiently. Goodput, as expected, is higher for addition of lines because of better paths and more efficient use of the medium.

TABLE X

SIM PARAMETERS: NETWORK DENSITY EVAL. VS. AODV AND DSR

Parameter	Values
Number of Interfaces	12
Topology Boundaries	300m x 300m
Number of Nodes	50, 100, 150, 200, 250, 300
Avg. # of Neighbors	5.1, 9.7, 15.6, 20.9, 26.1, 29.7
Simulation Time	100s
Connection Pattern	50 Random Connections - 10s each
TTL for Control Pkts	15
Simulation Time	100s

packets than reactive protocols like DSR and AODV. As expected, ORRP with one line yields the highest average path length as having only 1 intersection point results in longer paths. As more lines are introduced, the average path length declines steadily. Because AODV and DSR use shortest path algorithms for route discovery, it is expected that they yield the smallest average path length even despite packet losses.

Figure 9 shows that with increased network density (number of nodes), the end to end latency increases for ORRP. This is expected as additional hops yield additional delays. Packet latency under ORRP is much less than DSR or AODV because it uses the medium more efficiently. Aggregate network goodput is affected mostly by the average path length graph and the amount of packets that reach given usage of the medium. Goodput decreases with higher density in ORRP because intersections between announcement and RREQ packets are "missed" with increased density. In the case of DSR and AODV, network saturation with omnidirectional antennas prevents packets from being successfully distributed resulting in lower goodput.

F. Number of Connections Evaluation vs. AODV and DSR

It has been shown that network congestion can be controlled and limited by routing packets using two-phase routing algorithms [25] [24]. Current wireless networks measure route cost through hop count. In high-traffic networks, by choosing the shortest path, nodes with many connections will become saturated with packets. Busch et al. [25] has shown that by drawing a perpendicular bisector between source and destination and forwarding packets from source to a random point on the perpendicular bisector which in-turn forwards to destination when that point is reached, load can be balanced across the network. In much the same way, ORRP inherently implements a seemingly two-phase routing algorithm because it provides rendezvous abstractions whereby the source sends to the rendezvous node and the rendezvous node sends to the destination. In this section, we seek to understand how the number of connections effect the packet delivery success and average path length throughput network-wide with ORRP, AODV, and DSR. Table XI gives the simulation parameters that differ from the default.

TABLE XI
SIMULATION PARAMETERS: NETWORK OF CONNECTIONS EVALUATION VS. AODV AND DSR

Parameter	Values
Number of Interfaces	12
Topology Boundaries	300m x 300m
Number of Nodes	100 (Avg # Neighbors: 9.68)
Simulation Time	100s
Connection Pattern	10-100 Random Connections - 10s each
TTL for Control Pkts	15
Simulation Time	100s

As can be seen from Figure 10, ORRP with more lines delivers far more packets than AODV or DSR and is fairly consistent in number of packets delivered despite number of connections. This is due to more efficient medium usage by directional communications methods. AODV and DSR suffer when the network becomes more saturated. Average path length for all cases seems fairly constant as number of connections shouldn't affect the path length chosen.

G. Transmission Rate Evaluation vs. AODV and DSR

One of the key metrics in wireless networks is network goodput. In wireless networks, goodput 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 on aggregate network goodput in comparison to DSR and AODV. To do so, we make all-to-all connections and slowly increase the CBR rate until network capacity is reached. It is expected that with shorter paths, a higher aggregate goodput will result network-wide. We expect that DSR and AODV will reach network capacity much quicker due to omni-directional transmission. Table XII gives our simulation parameters and Figure 11 illustrate our results.

TABLE XII
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 / Rate	512 KB / Varied from 2KB-20KB
Transmission Duration	5.0 seconds

Our results in Figure 11 show that network goodput increases with increase in lines despite an increase in control packets sent. As the rate is increased, DSR and AODV reach the network capacity much faster. We see that latency starts out much higher for DSR and AODV but drops sharply toward the end. The reason for this is that most packets are now not being successfully delivered. Therefore, only the successfully delivered packets have latency measured resulting in a latency drop as rate is increased. As expected for ORRP, the increase in lines increases latency simply because of additional saturation of the network. This is not counterbalanced by shorter paths as in the case with more lines.

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, end to end latency, and aggregate goodput. Our 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 numerical 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, control packet overhead, end to end latency and aggregate goodput on various topologies, network densities, void conditions, number of connections, number of interfaces, and transmission rates. 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.

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] B. Cheng, M. Yuksel, S. Kalyanaraman, "Orthgonal Rendezvous Routing Protocol for Wireless Mesh Networks," Proceedings of the 17th IEEE International Conference on Network Protocols (ICNP), November 2006.

[2] B. Cheng, M. Yuksel, S. Kalyanaraman, "Directional Routing for Wireless Mesh Networks: A Performance Evaluation," To appear in the Proceedings of IEEE Workshop on Local and Metropolitan Area Networks (LANMAN), Princeton, NJ, June 2007.

[3] H. Gossain, T. Joshi, C. 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.

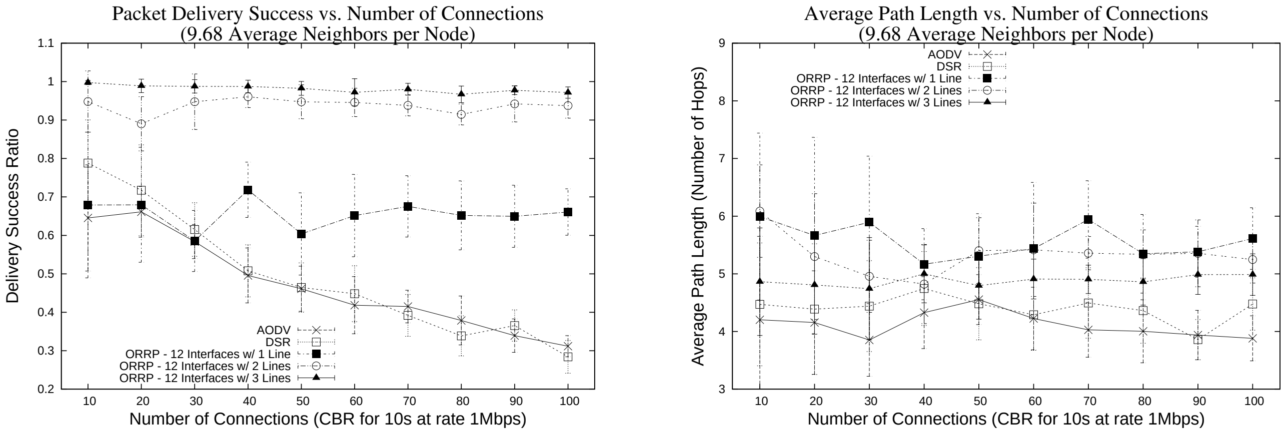

Fig. 10. Data delivery success and average path length vs. number of connections. As connections increase, it can be seen that the network becomes saturated faster with AODV and DSR. Average path length is fairly constant throughout. There is a high variability in path length after each run especially with less connections because source and destination nodes are randomly chosen.

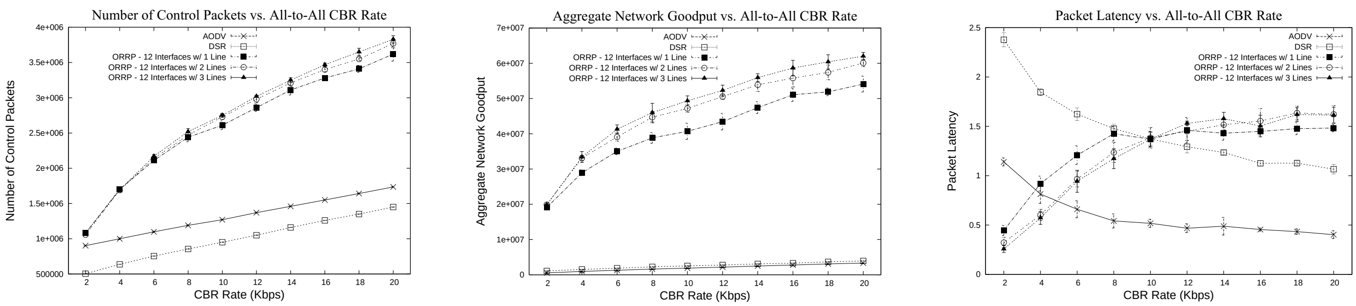

Fig. 11. Average goodput increases as number of lines increase. DSR and AODV reach the network capacity much faster because they use omnidirectional communications.

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

[5] 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.

[6] J. Bicket, D. Aguayo, S. Biswas, and R. Morris, Architecture and Evaluation of an Unplanned 802.11b Mesh Network, *ACM/IEEE MOBICOM* 2005.

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

[8] 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.

[9] D. B. Johnson, D. A. Maltz, J. 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.

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

[11] 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.

[12] R. R. 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.

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

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

[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] R. R. 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.

[22] H. Willebrand and B. S. Ghuman, Free Space Optics (Sams Pubs, 1st edition, 2001).

[23] D. J. T. Heatley, D. R. Wisely, I. Neild, and P. Cochrane, Optical Wireless: The story so far, IEEE Communications (December 1998), Volume 36, pp. 72 - 74.

[24] M. Kodialam, T. V. Lakshman, and Sudipta Sengupta, "Efficient and Robust Routing of Highly Variable Traffic", Third Workshop on Hot Topics in Networks (HotNets-III), San Diego (USA), November 2004.

[25] C. Busch, M. Ismail, and J. Xi, "Oblivious Routing on Geometric Networks", Proceedings of the 17th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA), pp 316-324, Las Vegas, Nevada, July 2005.

[26] M. Yuksel, J. Akella, S. Kalyanaraman, and P. Dutta, "Free-Space-Optical Mobile Ad-Hoc Networks: Auto-Configurable Building Blocks", To appear in ACM/Springer Wireless Networks, 2007.

[27] D. Britz, R. Miller, "Mesh Free Space Optical System: A method to Improve Broadband Neighborhood Area Network backhaul", Proceedings of IEEE LANMAN, Princeton, NJ, June 2007.