Performance of Multihop Wireless Networks: Shortest Path is Not Enough

Douglas S. J. De Couto   Daniel Aguayo   Benjamin A. Chambers   Robert Morris

M.I.T. Laboratory for Computer Science
{decouto, aguayo, bac, rtm}@lcs.mit.edu

Abstract

Existing wireless ad hoc routing protocols typically find routes with the minimum hop-count. This paper presents experimental evidence from two wireless test-beds which shows that there are usually multiple minimum hop-count paths, many of which have poor throughput. As a result, minimum-hop-count routing often chooses routes that have significantly less capacity than the best paths that exist in the network. Much of the reason for this is that many of the radio links between nodes have loss rates low enough that the routing protocol is willing to use them, but high enough that much of the capacity is consumed by retransmissions. These observations suggest that more attention be paid to link quality when choosing ad hoc routes; the paper presents measured link characteristics likely to be useful in devising a better path quality metric.

1. Introduction

Ad hoc networking has grown into a large and diverse field of research, which spans topics from power control to privacy and security. There has been much work on routing in ad hoc networks, and protocols such as DSR [12], AODV [19], Grid [15], and DSDV [20] have been shown in simulation to work very well on small to medium networks [15, 3]. However, our experience with two wireless networks leads us to believe that there are significant challenges left in finding and choosing usable routes, even in small ad hoc networks which are static.

To explore how ad hoc protocols work when implemented as part of a complete system, we built two experimental wireless networks. The first, called the "indoor" net, has 18 small PCs as nodes on the fifth and sixth floors of our building, as shown in Figure 1. We chose node locations that would keep the network connected, while also providing spatial diversity. Each indoor PC has a Cisco Aironet 340 wireless adapter [1]; these adapters implement the IEEE 802.11b Direct Sequence Spread-Spectrum protocol [7], and have 30 mW of output power.

The second "rooftop" network [4] has seven nodes spread over one square kilometer near our lab, as shown in Figure 4. This is a dense residential area with primarily two- and three-story buildings. Nodes have external omni-directional antennas attached to chimney-tops, except for Node 30, which is on the ninth floor of our lab building and has a directional antenna aimed roughly at Node 7. This network uses Cisco 350 wireless adapters, which are like the 340s, but with 100 mW of output power.

This research was supported by grants from NTT Corporation under the NTT-MIT collaboration. More details about the Grid project are available at http://www.pdos.lcs.mit.edu/grid.

Both networks run our implementation of the DSDV protocol. We thought that DSDV would be a good choice because the networks are static, and DSDV seemed to be one of the simpler protocols to implement. Indeed, in simulation scenarios with little to no mobility, like our static networks, DSDV has been shown to deliver the same number of packets as protocols such as DSR and AODV [3]. Unfortunately, our daily experience with the indoor network has been disappointing: tasks like interactive logins and file transfers are often unusably slow.

Our first hypothesis was that there might be no high-quality paths connecting some parts of the network. Although there is a wide range of link qualities in the network, it turns out that the network is still well connected with links that deliver more than 90% of the packets. For example, Figure 2 shows the subset of inter-node radio links that had loss rates less than 10% at a particular time; with the exception of one node, these links form a connected graph.

To find an approximate lower bound on how well a routing protocol should be able to do, we tried to find the path with the highest throughput between each pair of nodes in the network. For each pair of nodes we generated all possible paths of length less than or equal to four hops. We pruned the large number of resulting paths by eliminating paths containing links with low delivery rates, as determined by earlier link measurements. We then randomly chose among the remaining paths so that each pair of nodes had a few paths of each length in each direction. For each path, the source node sent source-routed packets as fast as possible over that path to the destination node for 30 seconds; the destination measured the arrival rate for each path. The black points in Figure 3 indicate the best throughput for each pair of nodes. We ran a similar experiment using DSDV to find routes and forward packets between each pair of nodes. The grey points in Figure 3 show the throughput of traffic routed by DSDV for each node pair. For almost all of the node pairs, DSDV routed packets over paths that were considerably slower than the best path. DSDV sometimes performed better than the "best" route because the two experiments were necessarily run at different times, so network conditions were not identical; in addition, since we only tested some of the possible paths between each pair, we may have missed the actual best path in some cases.

This result was surprising: all else being equal, multi-hop path capacity is determined by hop count [14], and DSDV finds shortest paths. But of course, all else is not equal. Figure 6 shows the packet transmission rates of three-hop paths from node 10 to 18, from an earlier experiment. These are often the shortest paths that a routing algorithm could find from 10 to 18. It is clear from the graph that if a routing protocol made a random choice among these paths, it

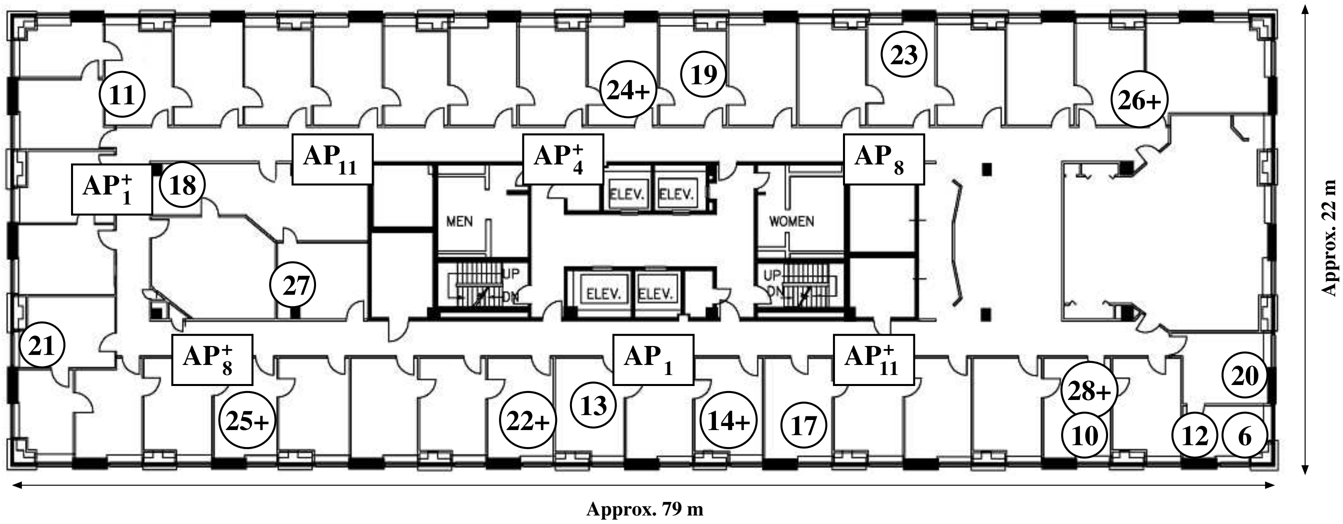

Figure 1: A map of the indoor network. The circles are nodes; the squares marked 'AP' are access points, labeled with their channel. Our experiments do not use the APs, but they may have affected the results. 6th floor nodes and APs are marked with '+'.

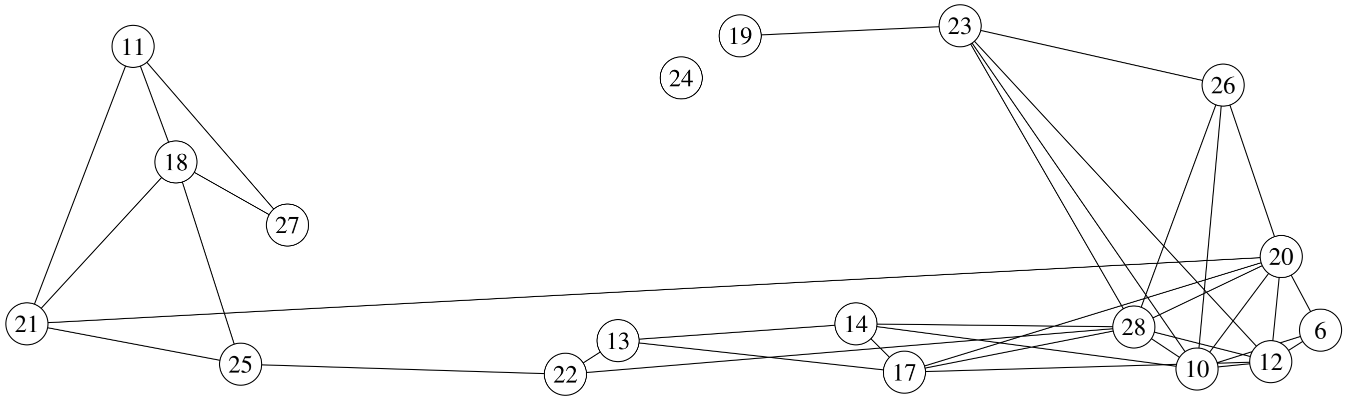

Figure 2: Indoor network links with greater than 90% delivery rate in both directions show that the network is well connected (8-Feb-13:30-50-byte). Node 24 was not functioning during this experiment. For clarity, node positions have been slightly changed from Figure 1.

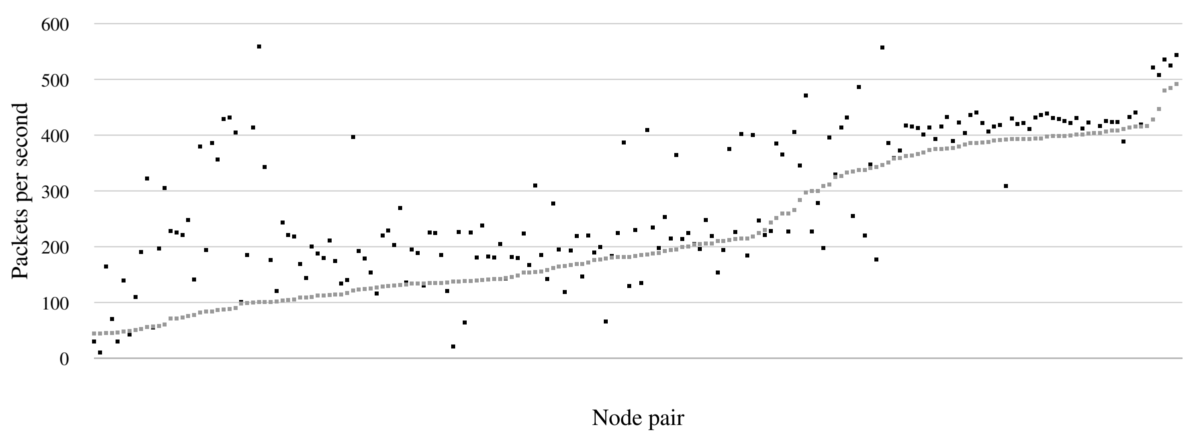

Figure 3: DSDV and best static route throughput for pairs of nodes in the indoor network, sorted by the DSDV throughput of each pair. DSDV does worse on 152 out of the 185 pairs. DSDV throughput is shown in gray, static route throughput in black (DSDV data from 30-Sep-22:10; static route data from 19-Aug-28-Aug). Packets had 124 bytes of 802.11 payload.

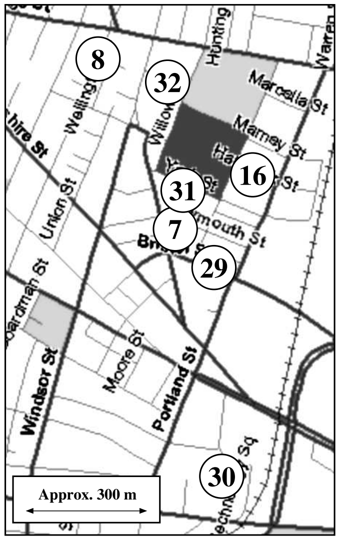

Figure 4: A map of the outdoor rooftop network.

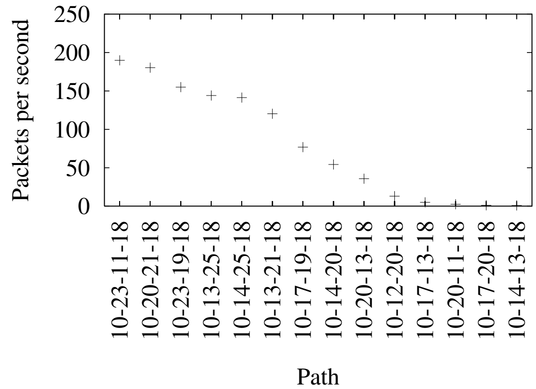

Figure 6: Measured capacity on various 3-hop paths from node 10 to node 18 (27-Jun-17:00). The units are packets per second; each packet contains 124 bytes of 802.11 payload.

would be unlikely to choose the path with the best packet throughput. In fact, in this case a randomly chosen path would achieve less than half of the maximum path throughput more than half of the time. Although we do not have as detailed data for all other node pairs, initial results imply that we can expect to see similar behavior for other node pairs: a wide range in packet throughput among the shortest paths between each pair.

The rest of the paper uses link quality measurements to explain the observed differences in quality of the paths between given node pairs. It also presents detailed wireless link measurements which outline some of the difficulties that would be involved in finding good paths. The measurements include the distribution of link-by-link loss statistics, to see how accurately we need to distinguish links; the extent to which link quality differs depending on link direction; the rate at which link quality changes with time, which determines the usefulness of averaging techniques; and the relationship between signal strength reported by 802.11 cards and loss rate. We conclude the paper by outlining some potentially fruitful avenues of research.

2. Related Work

One solution to low link quality is to improve the apparent quality with some form of redundancy. Forward error correction, MAC-level acknowledgment and retransmission, and solutions such as Snoop-TCP [2] and Tulip [18] all take this approach. Even with these techniques it is preferable to use high-quality rather than low-quality links: retransmissions (or other redundancy) reduce useful link capacity and generate interference. For these reasons, error correction should be combined with route selection that avoids low-quality links when possible.

A number of existing ad hoc wireless routing algorithms collect per-link signal strength information and apply a threshold to eliminate low-quality links [10, 9, 21, 8, 13]. This approach has two drawbacks. First, thresholding may eliminate links that are necessary for connectivity, or include links of dubious quality; both of these are likely to be issues if many links are of intermediate quality. Second, Section 3.3 shows that, in the case of some 802.11 cards, reported signal strength is a poor predictor of loss rate.

Wireless Quality-of-Service (QoS) algorithms approach route selection from the top down. Some techniques explicitly schedule transmission slots in time or frequency division MAC layers to provide bandwidth guarantees [11, 6, 16, 17, 24], while others treat the MAC as opaque, and rely upon it for bandwidth and delay information and constraints [5, 23, 22]. These approaches are only successful if the lower layers can provide accurate information about the actual links, such as the average number of usable transmission slots, or the achievable throughput. However, none of these approaches consider the case of lossy links.

3. Link Behavior

We carried out several experiments to evaluate the link characteristics, particularly loss rate, between every pair of nodes. In each complete experiment, each node in turn broadcasts a series of equally-sized packets at a constant rate; the other nodes record which packets they receive. Because we wanted to discover the underlying link behavior, we used broadcasts to avoid the 802.11 ACK and RTS/CTS mechanisms, which hide the real loss rates.

No routing protocol is running during these experiments: only experiment packets are sent or received on each node's wireless interface. The interfaces are running in 802.11 ad hoc mode. Packets were sent at about 0.4 Mbps, which is well below the minimum 802.11 capacity of 1 Mbps. However, on some occasions nodes were not able to broadcast at the desired rate, perhaps because of 802.11 traffic outside our control, or interference appearing to the card as carrier.

3.1 Distribution of Link Quality

The link quality distribution will affect the way we distinguish good links from bad. Most current ad hoc routing protocols assume that link quality follows a bi-modal distribution, where links are either very good or very bad. Protocols assume that if a link can successfully deliver a routing control packet, then the link is useful for delivering data. In this case, protocols don't need a very accurate link classifier because if a link seems at all usable, it is likely to be very good. On the other hand, the link quality distribution may be very spread out. In this case, protocols will need to accurately differentiate good links from links that aren't suitable for data, but still deliver some control packets. Our experiments indicate that the distribution is spread out.

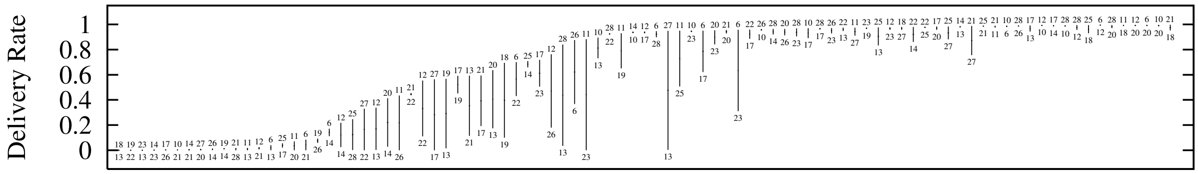

Figure 5: Delivery rates for each link pair in the indoor network (8-Feb-13:30-50-byte). The y values of the two ends of each line indicate the delivery rate in each direction; the numeric labels indicate the sending node for that delivery rate. Links with zero delivery rate in both directions are omitted. 91 links are shown.

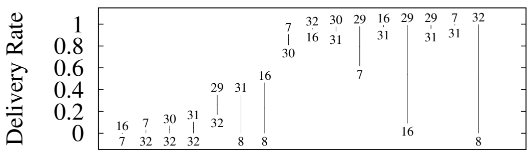

Figure 7: Delivery rates on each link pair for the outdoor rooftop network (6-Mar-18:30-50-byte), as in Figure 5. 16 links are shown.

We conducted a set of experiments with the indoor testbed using 50-byte UDP packets (8-Feb-13:30-50-byte). Each node transmitted 1024 packets per second for 300 seconds. Figure 5 shows the results for each link pair, excluding link pairs which were not able to communicate at all. About 30% of the link pairs shown are completely unusable, although they might deliver a few routing packets. The best 40% of link pairs deliver about 90% or more of their packets; these are the links we would like to use in routes. The delivery rates of the remaining links are spread out. Other experiments on different days, at different times, and with different parameters confirm that in general the links in the network exhibit a wide range of delivery rates.

Link pairs that are very good in one direction tend to be good in both directions, and pairs that are very bad in one direction tend to be bad in both directions. However, at least 30% of the link pairs shown have asymmetric delivery rates, defined as a difference of more than 20% between the rates in each direction.

Figure 7 summarizes an identical set of experiments carried out on our rooftop network (6-Mar-18:30-50-byte). Like the indoor network, the rooftop network has widely varying delivery rates, with noticeable asymmetry. Experiments over several days exhibited similar distributions of delivery rates. The wide variation in delivery rates for both networks suggests that routing protocols may often choose links that are high enough quality to pass routing protocol packets, but which still have substantial loss rates.

3.2 Link Variation over Time

One way to determine link quality is to measure it by counting the number of packets received over a period of time. However, the accuracy of this approach is sensitive to length of time over which the delivery rate is measured. For example, Figure 8 shows the second-by-second delivery rates for two links (8-Feb-13:30-50-byte). The

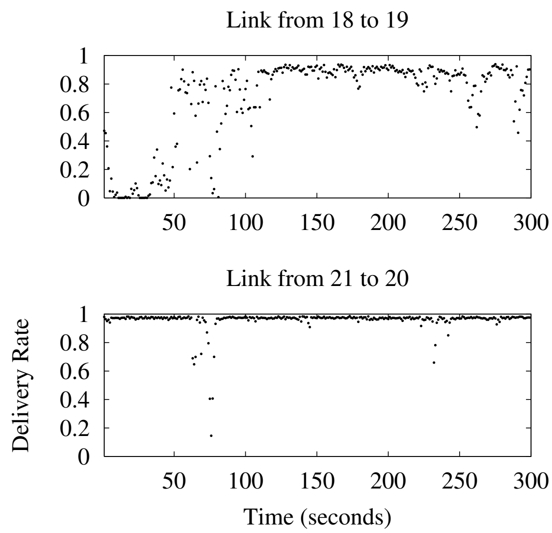

Figure 8: Example per-second variation in link delivery rates (8-Feb-13:30-50-byte). Each point is the delivery rate over one second. The delivery rate of the link from 18 to 19 fluctuates quickly, while the link from 21 to 20 is comparatively stable.

graphs show that while delivery rates are generally stable, they can sometimes change very quickly. Averaging may work well on the link from node 21 to 20, but it is likely to hide much of the detailed behavior of the link from node 18 to 19.

Figure 9 summarizes variation in loss rate over time for all links. For each link, we calculated the mean and standard deviation of the 1- and 10-second loss rates over the whole experiment. The graph shows the cumulative distribution of these standard deviations, normalized by the respective means. We use loss rates rather than delivery rates for this analysis to emphasize the changes in the delivery rate on links with low loss, since very lossy links are useless for data traffic regardless of their variation.

Results for 1- and 10-second windows show that quite a few links vary greatly on these times scales. For example, half of the links had standard deviations in their 1-second loss rates that exceeded half of the mean 1-second loss rate. This suggests that wireless routing protocols should use agile predictors of link loss rates.

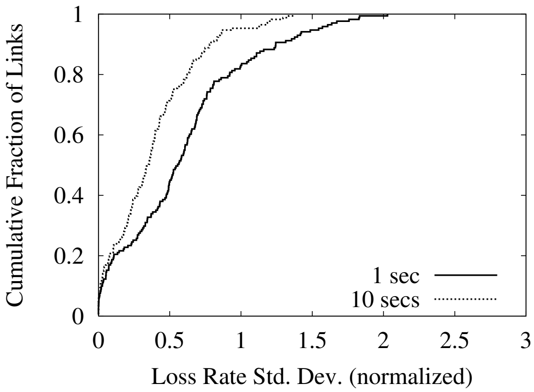

Figure 9: The indoor network's cumulative distribution of the standard deviation of short-term link *loss* rates. For each link, the loss rate is calculated over 1- and 10-second intervals, and the standard deviation is normalized by the link's mean loss rate (8-Feb-13:30-50-byte). Many links show significant variation in short-term loss rates over time.

Figure 10 shows the variation in short-term loss rates from the same experiment as in Figure 9, but carried out on the rooftop network (6-Mar-18:30-50-byte). This figure shows that short-term loss rates in the rooftop network vary much like those in the indoor network.

In addition to looking at short-term loss rates, we measured how link delivery rates change throughout the day. Figure 11 shows delivery rates for two links over a 24-hour weekday period in January. Every half-hour, each node tried to broadcast 100 1024-byte packets per second for 30 seconds. The results for the link from node 6 to node 23 are particularly interesting; the fact that the quality increases dramatically at 8 am suggests that opening office doors in the morning increases link quality.

3.3 Link Signal Strength

Signal strength could potentially be helpful in predicting link quality. To explore this possibility, we recorded signal strength (dBm) from the radio interface for each received packet during the link experiments. Figure 12 shows how the short-term delivery rate varies with these values for a few example links. Unfortunately there is no good correlation between delivery rate and the radio's measurements. Instead, the data reflect the physical fact that received signal strength is mostly a function of the distance between nodes. The link from 18 to 19 is a long link with low signal strength, and as a result is very susceptible to noise. Since the successful reception of a packet depends on the signal to noise ratio (SNR) at the receiver, this link's delivery rate varies significantly. In contrast, the link from 18 to 11 is a short link. Because it has a very high received signal strength, it is robust to noise and has high delivery rates. The links from 23 to 19 and 27 to 11 are medium range links which still deliver most packets. Since our radios don't provide a noise estimate with which to compute the SNR, we cannot determine much about the links from the signal strength estimate.

4. Research Agenda

Based on the measurements presented here, we intend to develop techniques to help ad hoc routing protocols choose high-quality routes. The main challenges involve practical estimates of link quality and techniques to combine link metrics into useful path metrics.

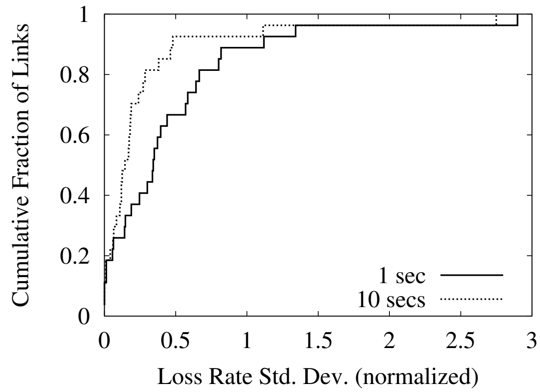

Figure 10: The rooftop network's cumulative distribution of the normalized standard deviation of short-term link loss rates over 1- and 10-second intervals, calculated as in Figure 9 (6-Mar-18:30-50-byte).

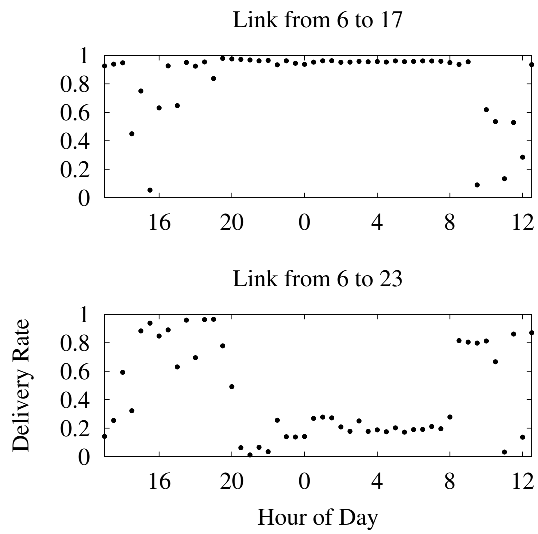

Figure 11: Example variations in link delivery rates throughout the day (10-Jan-24h-1024-byte).

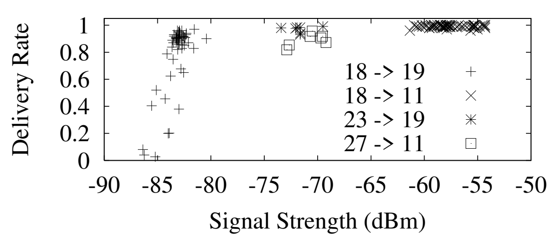

Figure 12: Delivery rates over 100 msecs versus average signal strength of received packets for four example links on the indoor network (6-Mar-18:30-50-byte). For clarity, only 1 in 50 of the data points are shown for each link. Signal strength more accurately reflects link distance than delivery rate.

One obstacle to using delivery rate as a link metric is that it requires many packet transmissions to measure; this is a problem if nodes move or if the environment changes rapidly. Signal-to-noise ratio might be useful as a fast predictor of delivery rates, if it is available from the radio hardware.

Combining route metrics to form a path metric is not straightforward. For example, the product of the delivery rates of the links in a path does not predict anything useful when 802.11 link-layer retransmissions are used. Self-interference by successive nodes in a route [14] may sometimes make long routes with good links less attractive than short routes with less good links. We are currently evaluating use of the expected total number of transmissions of a packet along a path (including forwarding and retransmission) as a path metric. This metric has a number of advantages: it captures the route's impact on the spectrum, it can be computed incrementally from link delivery rates, and it penalizes longer routes.

Finally, we plan to explore how protocols such as DSR and AODV handle the link quality distribution seen on our testbeds. These protocols are not simply shortest-path protocols: they notice when a route is using a very low-quality link and try to find a different route. We intend to compare the quality of the routes they find with the quality of "best" routes found by exhaustive search.

5. Acknowledgments

The authors would to thank Jinyang Li for her insightful comments, and for her work on the ad hoc capacity problem.

6. References

[1] *Using the Cisco Aironet 340 Series PC Card Client Adapters*. Cisco Systems Inc., March 2000. Part number: OL00379-01.

[2] Hari Balakrishnan, Venkata N. Padmanabhan, Srinivasan Seshan, and Randy H. Katz. A comparison of mechanisms for improving TCP performance of wireless links. *IEEE/ACM Transactions on Networking*, 6(5), December 1997.

[3] Josh Broch, David A. Maltz, David B. Johnson, Yih-Chun Hu, and Jorjeta Jetcheva. A performance comparison of multi-hop wireless ad hoc network routing protocols. In *Proc. ACM/IEEE MobiCom*, pages 85–97, October 1998.

[4] Benjamin A. Chambers. The Grid roofnet: a rooftop ad hoc wireless network. Master's thesis, Massachusetts Institue of Technology, May 2002.

[5] Shigang Chen and Klara Nahrstedt. Distributed quality-of-service routing in ad hoc networks. *IEEE Journal on Selected Areas in Communcations*, 17(8), August 1999.

[6] T.-W. Chen, J.T. Tsai, and M. Gerla. QoS routing performancs in multihop, multimedia, wireless networks. In *Proceedings of IEEE ICUPC '97*, 1997.

[7] IEEE Computer Society LAN MAN Standards Committee. *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*. New York, New York, 1997. IEEE Std. 802.11–1997.

[8] Brian H. Davies and T. R. Davies. The application of packet switching techniques to combat net radio. *Proceedings of the IEEE*, 75(1), January 1987.

[9] Rohit Dube, Cynthia D. Rais, Kuang-Yeh Wang, and Satish K. Tripathi. Signal stability-based adaptive routing (SSA) for ad hoc mobile networks. *IEEE Personal Communications*, February 1997.

[10] Tom Goff, Nael B. Abu-Ghazaleh, Dhananjay S. Phatak, and Ridvan Kahvecioglu. Preemptive routing in ad hoc networks. In *Proc. ACM/IEEE MobiCom*, July 2001.

[11] Yu-Ching Hsu, Tzu-Chieh Tsai, Ying-Dar Lin, and Mario Gerla. Bandwidth routing in multi-hop packet radio environment. In *Proceedings of the 3rd International Mobile Computing Workshop*, 1997.

[12] David B. Johnson. Routing in ad hoc networks of mobile hosts. In *Proc. of the IEEE Workshop on Mobile Computing Systems and Applications*, pages 158–163, December 1994.

[13] John Jubin and Janet D. Tornow. The DARPA packet radio network protocols. *Proceedings of the IEEE*, 75(1), January 1987.

[14] Jinyang Li, Charles Blake, Douglas S. J. De Couto, Hu Imm Lee, and Robert Morris. Capacity of ad hoc wireless networks. In *Proceedings of the 7th ACM International Conference on Mobile Computing and Networking*, pages 61–69, Rome, Italy, July 2001.

[15] Jinyang Li, John Jannotti, Douglas S. J. De Couto, David R. Karger, and Robert Morris. A scalable location service for geographic ad hoc routing. In *Proc. ACM/IEEE MobiCom*, August 2000.

[16] Chunhung Richard Lin. On-demand QoS routing in multihop mobile networks. In *Proc. IEEE Infocom*, April 2001.

[17] Anastassios Michail and Anthony Ephremides. Algorithms for routing session traffic in wireless ad-hoc networks with energy and bandwidth limitations. In *Proceedings of 12th IEEE International Symposium on Personal, Indoor and Mobile Radio Communications*, 2001.

[18] Christina Parsa and J. J. Garcia-Luna-Aceves. TULIP: A link-level protocol for improving TCP over wireless links. In *Proc. IEEE Wireless Communications and Networking Conference 1999 (WCNC 99)*, September 1999.

[19] Charles Perkins, Elizabeth Royer, and Samir R. Das. Ad hoc On-demand Distance Vector (AODV) routing. Internet draft (work in progress), Internet Engineering Task Force, January 2002. `http://www.cs.ucsb.edu/~ebelding/txt/aodvid.txt`.

[20] Charles E. Perkins and Pravin Bhagwat. Highly dynamic Destination-Sequenced Distance-Vector routing (DSDV) for mobile computers. In *Proc. ACM SIGCOMM Conference (SIGCOMM '94)*, pages 234–244, August 1993.

[21] Ratish J. Punnoose, Pavel V. Nitkin, Josh Broch, and Daniel D. Stancil. Optimizing wireless network protocols using real-time predictive propagation modeling. In *Radio and Wireless Conference (RAWCON)*, August 1999.

[22] Samarth H. Shah and Klara Nahrstedt. Predictive location-based QoS routing in mobile ad hoc networks. In *Proceedings of IEEE International Conference on Communications*, 2002.

[23] Prasum Sinha, Raghupathy Sivakumar, and Vaduvur Bharghavan. CEDAR: A core-extraction distributed ad hoc routing algorithm. In *Proc. IEEE Infocom*, March 1999.

[24] Chenxi Zhu and M. Scott Corson. QoS routing for mobile ad hoc networks. In *Proc. IEEE Infocom*, June 2001.