Measurement-based Characterization of a Wireless Mesh Network

Saumitra M. Das, Dimitrios Koutsonikolas, Y. Charlie Hu
School of ECE, Purdue University, West Lafayette, USA
{smdas, dkoutson, ychu}@purdue.edu

1. Introduction

Wireless mesh networks (WMNs) are characterized by static mesh routers connected by wireless links to each other and to a few gateway nodes. Such networks have a significant application of providing cheap and easily deployable broadband Internet access. Recently, the deployment and use of WMNs has increased significantly ([1], [2], [3], [4], [5], [6]). Thus, improving WMN performance will have a direct impact on a growing population of users. To improve WMN performance, it is important to understand the fundamental characteristics of the network and study its behavior. Such an understanding allows network operators to understand what applications are suitable to the environment and what adaptation may be needed for different applications to operate over the network.

To this end, in this chapter, we present a measurement study of our 32-node wireless mesh network testbed at Purdue University. Several important characteristics of the network are measured ranging from packet latencies, loss rates, throughputs, and variability of link characteristics. Finally, we present a study of the interference characteristics of the testbed All these metrics are fundamental ones that affect a large variety of applications in WMNs. Latency and latency variance are important for VoIP [8], multicast applications [9], video streaming [10], and also small HTTP transfers, over mesh networks. On the other hand, loss rate affects web access and TCP performance, and can also be used by routing protocols to construct high-quality paths (e.g., via routing metrics such as ETX [11], ETT [12] or SPP [9]). Throughput measurements give an idea of the download speeds we should expect in wireless mesh networks. Finally, understanding and characterizing interference among neighboring nodes is important for a variety of purposes such as channel assignment in multi-radio multi-channel networks (e.g., [13], [14]), route selection (e.g., [15]), and fair scheduling (e.g., [16], [17]). Hence, this study is valuable to both network designers and deployers to get a feel for typical network characteristics they may have to deal with. It also provides guidelines for the design of new protocols for wireless mesh networks. Finally, it provides an indication of the performance achievable using currently available state-of-the-art routing protocols and hardware devices.

2. Setup and Methodology

For our measurement study we used our testbed, MAP [7], shown in Figure 1. MAP currently consists of 32 mesh routers spread out across four academic buildings on the Purdue University campus. Each router has two radios. For this study, we used one of them: the Atheros 5212 based 802.11a/b/g wireless card. Each radio is attached to a 2dBi rubber duck omnidirectional antenna with a low loss pigtail to provide flexibility in antenna placement. Each mesh router runs Mandrake Linux 10.1 and the open-source *madwifi*. IP addresses are statically assigned. We performed our measurements only during the night to minimize interference from other 802.11b networks and other sources such as microwaves.

Fig. 1: MAP Testbed

3. Latency Measurements

3.1 Methodology

For this set of measurements the OLSR routing protocol [18] was used in MAP, enhanced with ETX routing metric. However, we did not want our measurements to interfere with the control packets of OLSR. Hence we ran OLSR daemon for enough time to ensure that the best paths based on ETX have been set up and stabilized between all pairs of nodes and then we stopped it from sending control messages. Hence, the same routing paths have been used for all the experiments. Our methodology for measuring packet latencies was as follows: Each node in turn sent 100 1470-byte ping packets to each other node. In this way, we obtained the RTTs between any pair of nodes for the paths used by the OLSR protocol. The packet size used is typical for many Internet applications. The packet interarrival time was 0.01sec which is equivalent to a sending rate of about 1.1 Mbps. Note that although many measurement papers prefer broadcast packets (e.g., [19]), we preferred to use ping packets. Since ping packets are unicast, our measured latencies better represent latencies experienced by unicast traffic, which we believe will be the common case in WMNs. For example, the RTTs in our experiment take into account retransmissions at the MAC layer, which are not used for broadcast packets.

3.2. Results

The results of the latency measurements are shown in Figure 2. Figure 2(a) plots the average (over 100 packets), maximum and minimum RTT values for each of the 654 paths used by OLSR (i.e., for each of the 654 pairs of nodes that were well connected to each other)[1]. Figure 2(a) plots the CDF of RTT values for 1-hop, 2-hop, 3-hop, and 4-hop paths.

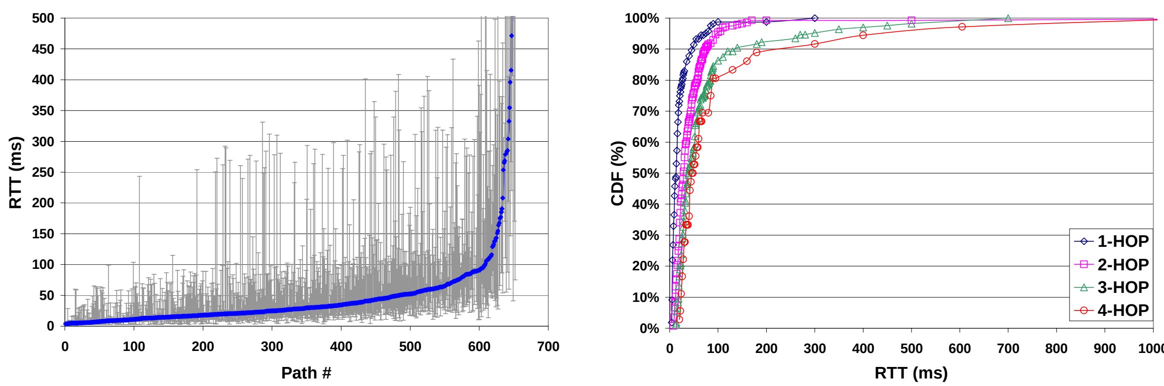

Fig. 2: (a) RTT, (b) CDFs of RTTs

From Figure 2(a) we observe that RTTs span a large range of values between 3.45 and 1374 msec (we kept the maximum value of the y-axis in the graph at 500msec, for better visibility). Also, the variations for many paths are very large, and this is mainly true for the maximum values. For example, paths with average RTT below 50msec have maximum values larger than 250msec. This shows that network conditions experience steep changes over time and there exist periods of very poor connectivity. This can be a significant problem for applications such as VoIP or video streaming which require constant jitter. On the other hand, minimum

[1] The remaining node pairs had some bad link in the paths found and were not well connected.

values (the true RTT) are not in general much smaller than the average ones, which shows that the quality of the paths is not very unstable. One more encouraging observation is that for the majority of the paths (600 out of 654 paths, about 91% of all paths), the average RTTs are lower than 100msec which is tolerable. Figure 2(b) confirms these observations, showing that a large fraction of paths have low RTTs. About 99% of the 1-hop paths and 95% of the 2-hop paths have RTTs lower than 100msec. Also, 50% of them have RTTs lower than 50msec. The majority of longer paths also experiences low RTT values – 86% of the 3-hop paths and 82% of the 4-hop paths have RTTs lower than 100msec. The problem with paths longer than 1 hop is that their CDFs have very long tails. For example, there is one 2-hop path, from node 12 to node 17, which had an average RTT of 1374msec.

4. Loss Measurements

4.1 Methodology

For the loss measurements, we used the same ping messages used in the previous section and we measured the average loss rate over 100 packets for each path. The results are shown in Figure 3. Figure 3(a) shows the average loss rate for each of the 654 paths and Figure 3(b) shows the CDFs of the loss rates for the 1-hop, 2-hop, 3-hop, and 4-hop paths.

4.2. Results

From Figure 3(a) we observe that loss rates cover the whole range between 0 and 100%. The results are not as encouraging as in the case of RTT measurements. More specifically, only 300 out of 654 paths (46%) have loss rates lower than 20%. On the other hand, for 33% of the paths loss rates are higher than 50%.

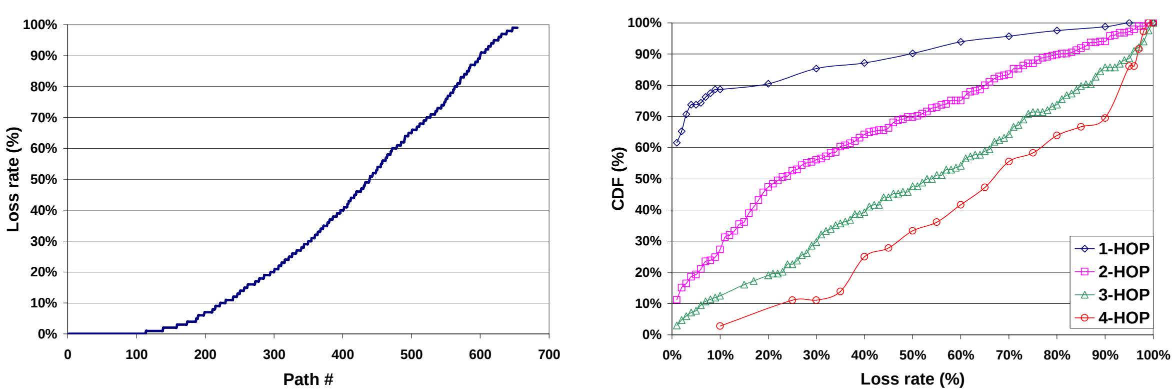

Fig. 3: (a) loss rate, (b) CDFs of loss rates

From Figure 3(b) we can see the correlation between loss rate and number of hops. The majority of 1-hop paths are paths with very low loss rate. 78.6% of them have loss rates lower than 10%. For paths longer than 1 hop, the loss rates increase rapidly. Only 47% of the 2-hop paths, 19% of the 3-hop paths and 8% of the 4-hop paths have low loss rates of 20% or less. In general, with this set of measurements we observed that the majority of the paths in our testbed have satisfactory average RTTs but high loss rates. In the next session we are interested in the transport layer performance achieved under such link conditions.

5. Transport Measurements

5.1 Methodology

In this section we are interested in characterizing the performance of transport layer protocols, UDP and TCP. Our methodology is as follows: Each node in turn initiated a UDP session to each other node, one at a time. We used a 5-sec idle interval between any two sessions to make sure they do not interfere. The sending rate was 5Mbps. The *iperf* tool was used for this experiment. In the second part of this experiment we used *netperf* to initiate TCP sessions between any two pairs of nodes. Again only one TCP session was active at any time. We found that for some of the paths iperf or netperf was unable to initiate a connection. We did not include these paths in the results shown below. The results are shown in Figure 4.

5.2. Results

Figure 4 shows that UDP outperforms TCP significantly. Figure 4(a) shows that 50% of the paths have UDP throughput higher than 1700Kbps, while from Figure 4(b) the percentage of paths that have TCP throughput higher than 1700Kbps is only 30%. Moreover, for half of the paths, TCP throughput is lower than 600Kbps. Even worse, for 31% of the paths, TCP throughput is less than 200Kbps; on the other hand, the minimum UDP throughput observed is 298Kbps. This result shows that TCP performs poorly in multihop wireless networks. From *tcpdump* and *tcptrace* we found the reason for this to be the lossy environment (Section 4), which prevents TCP from increasing its window.

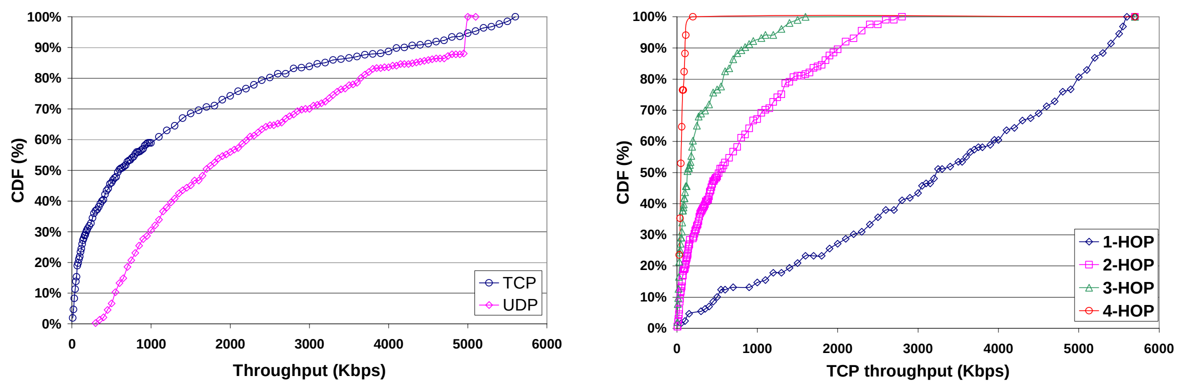

Fig. 4: (a) CDF of TCP/UDP throughputs, (b) CDF of TCP throughputs for different hop count paths

6. Interference Measurements

6.1 Methodology

The 32 nodes of the testbed can form up to 992 (directional) links. However many of these links may not exist, for example, if two nodes are far away from each other or if they are separated by many obstacles such as walls and metallic objects. To find out the exact number of links in our testbed and to get a qualitative estimate of their quality, we first we first transferred a series of 100 ping messages between each pair of nodes. Note that in this case no routing protocol was used, since we are interested in links (two nodes being able to communicate directly with each other) and not in multihop paths. Hence we could not use the ping messages from the previous sections. If all 100 ping messages between a pair of nodes were lost, we classified this link as non-existing. This experiment made sure that our network is not partitioned. As expected, the link quality and the node distances do not directly correspond. For example, there is a perfect link (0% loss) between nodes 1 and 13, which are placed in different buildings, because these nodes are close to windows and signal is propagated in the low-loss outdoor environment. On the other hand, there is no link between nodes 19 and 20, although their distance is half compared to the distance of link 1 – 13. In total, we found 257 links with loss rate less than 100%.

To measure interference, we executed the following series of experiments. In the first set of experiments, each node in turn broadcasted 1470-byte packets as fast as it could for 30 seconds. All other nodes that could receive packets measured the throughput in these 30 seconds. This experiment gave us the throughput of all existing links in our testbed. Figure 5(a) shows the cumulative distribution function (CDF) of the throughput for the 257 existing links.

In the second series of experiments, we measured the pairwise interference among all nodes in our testbed, following a methodology similar to [20]. For this experiment, we had each *pair* of nodes broadcast 1470-byte packets *togethe*r for 30 seconds. In each 30 second interval, all other nodes except for the two senders measured the throughput from the two senders. In the end of this experiment, for each pair of nodes, one viewed as the sender S and the other the interferer I (and vice versa), we calculated the receiver throughput ratio (RTR) at each one of the rest 30 nodes (R) as follows:

$$RTR_R^{S,I} = Throughput_R^{S,I} / Throughput_R^S$$

where $RTR_R^{S,I}$ is the receiver throughput ratio at receiver R when S is the sender and I the interferer node, $Throughput_R^S$ is the throughput at R from node S when only S transmits and $Throughput_R^{S,I}$ is the throughput at R from node S when both S and I transmit simultaneously. If $RTR_R^{S,I} < 0.9$, we consider that node I is an interferer for link S→R, else it is not. In this way we found out all nodes that are not interferers for each particular link S→R according to the pairwise interferer model.

In the third set of experiments, we measured multi-way interference[2] as follows. We considered each of the 257 links in turn, along with their non-interferers (say *n*). For each of these links we had again the sender broadcast packets for 30sec along with 1, 2, 3, ... up to *n* interferers simultaneously and every time we measured the sender's throughput at the receiver and calculated the receiver's throughput ratio. Comparing the ratio in case of 2, 3, ...*n* interferers at a time with the ratio in case of 1 interferer will tell us for which links the pairwise model cannot accurately predict interference.

6.2 Results

6.2.1 Pairwise interference

Figure 5(b) shows the CDF of the number of interferer nodes for the 257 existing links. We observe that the number of interferers varies but, in general, pairwise interference is a widespread phenomenon in our testbed. About 15% of the links have no interferers. On the other hand, 60% of the links have 10 or more interferers and 50% have 13 or more interferers. Finally,

---

[2] Interference from the cumulative effect of multiple transmitters on a given link. Such interference cannot be inferred by pair wise measurements since they only consider two transmitters at a time.

the last 17% of the links have more than 20 interferers, and, as we mentioned before, we removed these heavily affected links from the rest of our experiments.

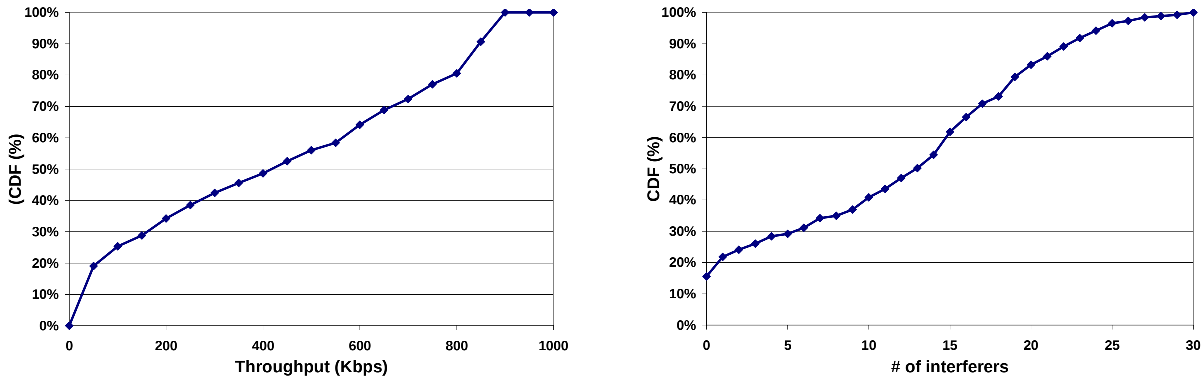

Fig. 5: (a) CDF of individual throughputs, (b) number of interferers for the 257 links of the testbed.

According to [20], the pairwise interference model is accurate in predicting interference in the majority of practical cases. Our results seem to verify it. However, we still wanted to study further what happens to those links that are not affected by pairwise interference. This is the subject of the following paragraphs.

6.2.1 Does multi-way interference occur?

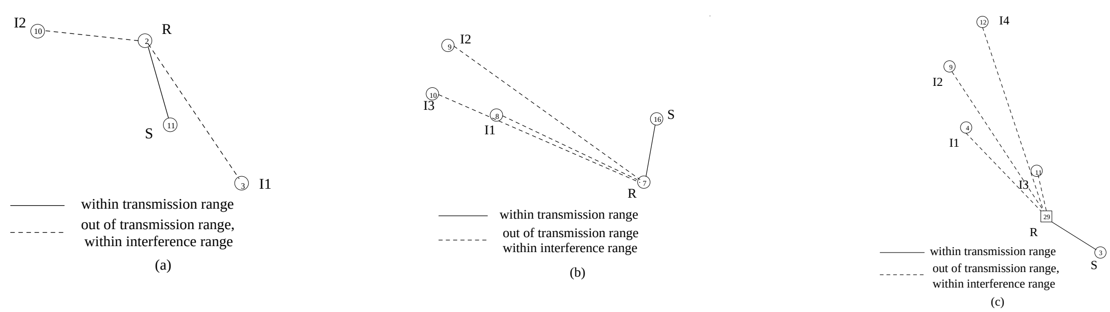

(a) (b) (c)

Fig. 6: Three examples in our testbed showing that pairwise interference is not always enough.

Figure 6 shows three examples in our testbed where the pairwise interference model is not enough. In these figures S is the sender, R is the receiver, and I_1, I_2, I_3, I_4 are nodes that were found not to interfere according to the pairwise model, but they might interfere if two, three, or four of them transmit together. Table 1 shows the receiver throughput ratio for different number of interferers in the three examples of Figure 6. The second column shows the minimum ratio at the receiver when interferers are considered one at a time, the third one the minimum ratio when interferers are considered in pairs, and the fourth one the minimum ratio when interferers are considered three at a time. Finally, the last column shows the ratio when all four interferers transmit simultaneously. As we observe, the ratio for the second column is above 0.9 in all the three examples (no interference according to the pairwise model).

In the first example, node 11 is the sender; node 2 is the receiver, and nodes 3, 10 the interferers. As Table 1 shows, if we allow both nodes 3 and 10 to transmit simultaneously with node 11, the receiver throughout ratio is reduced to 0.74, quite below the threshold of 0.9 and the quality of link 11→2 worsens significantly. Note again that the distance of the nodes is not directly mapped to the quality of the links. For example, nodes 11 and 10 are almost equal distance away from node 2, however only link 11→2 exists.

In the second example, node 16 is the sender; node 7 is the receiver, and nodes 8, 9, 10 the interferers. In this example, all the three interferers are very far from link 16→7 and they cannot affect it (throughput ratio is 1 when each of these three nodes transmits simultaneously with the sender). Even when two of these nodes transmit together, the ratio is not reduced a lot -- it only drops to 0.88, very close to the threshold. But when all the three nodes 8, 9, 10 transmit simultaneously, the ratio becomes 0.74 which makes the link quality unacceptable.

Finally, the third example shows that in some cases even 5-way interference has to be considered. In this example node 3 is the sender, node 29 is the receiver and nodes 4, 9, 11, and 12 the interferers. As Table 1 shows, the throughput ratio remains above the threshold when two or three of the interferers transmit simultaneously (with minimum values equal to 0.96 and 0.91 respectively) but it drops to 0.86 when all four interferers transmit simultaneously.

Example	**Throughput ratio**			
	Min 1 intf.	Min 2 intf.	Min 3 intf.	4 intf.
(a)	0.92	0.74	-	-
(b)	1	0.88	0.74	-
(c)	0.93	0.96	0.91	0.86

Table 1: Three examples in our testbed showing that pairwise interference is not always enough.

6.2.2 How widespread is multi-way interference?

Out of 257 links in our testbed, we found 16 for which the pairwise interference model could not accurately predict interference. Out of these 16 links, 4 had throughput between 200 and 400Kbps, 9 had throughput between 600 and 800 Kbps, and the rest three had throughput higher than 800Kbps without any interferers present. Some general observations are as follows.

As we observe in Figure 6(a), about 20% of the links have throughput lower than 50Kbps. For those links, no observation can be made about interference, since the quality is so bad that adding more interferers cannot make it worse. Actually, when we repeated our experiments, many of them got zero throughput in some cases. Such links will probably not be selected by a routing protocol that uses a link-quality based routing metric (e.g., ETX or SPP). Hence, in the rest of the chapter, we ignore these links. Similar methodology is followed in [20] where low quality links are rejected using an ETX-based threshold.

We did not find any case of 3-way interference for the 36 links with throughput between 50 and 200Kbps (about 14% of the total 257 links). For these links, throughput is still very low, although non-zero. For many of them we observed large variations in throughput when the number of interferers changed. For example, in many cases throughput was increased when we added interferers, compared to the case where only the sender transmitted. For the rest of them one interferer was enough to reduce the throughput ratio to very low levels, hence the pairwise model was enough. Note that a reasonable routing algorithm should also avoid most of these links. For the 77 links of medium quality with throughput between 200 and 600 Kbps (30% of the total 257), the pairwise interference model was successful in predicting interference in almost all cases. For those links, we did not observe the strange variations described in the previous paragraph, but in most cases one interferer was enough to change the link quality from medium to low and reduce the throughput ratio below acceptable levels. Only 4 links remained unaffected by single interferers, but were affected when two or three interferers transmitted simultaneously.

The majority of cases (9 out of the 16 links) where pairwise interference model was not enough were observed for the 41 links (16% of the total 257) of medium to high quality with throughput varying between 600 and 800 Kbps. Since throughput is high enough for these links, there is margin for gradual decrease by adding more interferers. Hence we had cases where one interferer reduced the ratio only slightly but without crossing the 0.9 threshold, the second interferer sent the ratio close to the threshold, and three or more interferers resulted in large throughput reduction.

Finally, for the 51 high quality links (20% of the total 257) with throughput higher than 800Kbps, the common case is that if such a link is not affected by other nodes, when they are considered one at a time, it is also not affected when the other nodes are considered more than one at a time. Hence, again the pairwise interference model gives the correct answer in most cases. But we still found three cases where 3-way interference should be considered. Thus, while multi-way interference does occur and when it does it significantly affects throughput, we found that the phenomenon is not widespread and depends on the original link quality.

7. Conclusions

In summary, we can make the following observations from our measurement study regarding the various network level metrics expected in an operational wireless mesh network.

- Despite the small scale of such networks, latency observed by the applications can range from miniscule to 100ms with longer paths having longer latency. Interestingly, large latency variations can occur across a large number of paths. Multicast and streaming applications will need to be robust to the scale of these latencies and jitter.
- While around 100 (mostly one-hop) links have almost no loss, all other links have loss rates almost uniformly varying between 0 and 100%. Longer paths have significantly larger loss rates. Note that these are loss rates observed on actual paths selected by the routing protocol that applications will use (not random links) and occurred despite MAC layer retransmissions. This indicates that despite the selection of paths that minimize the expected transmission count, higher-layer protocols are still likely to experience losses and need to have mechanisms to deal with them efficiently.
- While the pure transport layer throughput (using a backlogged UDP flow) is high, >50% of paths have >1500Kbps throughput, TCP cannot attain this available bandwidth and over the same set of paths has lower throughput. This was attributed to frequent loss events and subsequent reduction of window sizes. In addition, the RTT experienced by the TCP flow increased with congestion. Finally, TCP throughput reduces drastically with hop count.
- Finally, significant interference exists from nodes that are out of transmission range. More the 70% of nodes had more than 5 interferers. We also found that pair-wise interference measurements are not sufficient and multi-way interference can occur when two or more nodes transmit simultaneously. Thus, schemes that assign channels or perform scheduling need to be careful in deciding which transmissions are non-interfering.

While these measurements characterize only our specific testbed (topology, environment and hardware), we believe that the overall take-home message is more general and we expect that the trends we found will hold in most other wireless mesh network deployments.

8. References

[1] MIT Roofnet. http://www.pdos.lcs.mit.edu/roofnet.

[2] Bay area wireless users group. http://www.bawug.org.

[3] Seattle wireless. http://www.seattlewireless.net.

[4] Champaign-Urbana community wireless network. http://www.cuwireless.net.

[5] Southampton wireless network. http://www.sown.org.u.

[6] Wireless leiden. http://www.wirelessleiden.nl.

[7] Mesh@Purdue. http://www.engineering.purdue.edu/MESH.

[8] Dragoş Niculescu, Samrat Ganguly, Kyungtae Kim and Rauf Izmailov, Performance of VoIP in an 802.11-based Wireless Mesh Network. In *Proc. of IEEE INFOCOM*, 2006

[9] S. Roy, D. Koutsonikolas, S. M. Das, and Y. C. Hu. High-Throughput Multicast Routing Metrics in Wireless Mesh Networks. In *Proc. of ICDCS*, 2006.

[10] Yiannis Andreopoulos, Nicholas Mastronarde, and Mihaela van der Schaar. Cross-Layer Optimized Video Streaming Over Wireless Multihop Mesh Networks. *IEEE Journal on Selected Areas in Communications, vol. 24, No. 11*, November 2006.

[11] D. S. J. De Couto, D. Aguayo, J. C. Bicket, and R. Morris. A high-throughput path metric for multi-hop wireless routing. In *Proc. of ACM MobiCom*, 2003.

[12] R. Draves, J. Padhye, and B. Zill. Routing in multi-radio, multi-hop wireless mesh networks. In *Proc of ACM Mobicom*, 2004.

[13] A. Raniwala and T. Chiueh. Architectures and algorithms for an IEEE 802.11-based multi-channel wireless mesh network. In *Proc. of IEEE Infocom*, 2005.

[14] K. Ramachandran, E. Belding, K. Almeroth, and M. Buddhikot. Interference-aware channel assignment in multi-radio wireless mesh networks. In *Proc. of IEEE Infocom*, 2006.

[15] A. P. Subramanian, M. M. Buddhicot, and S. Miller. Interference aware routing in multi-radio wireless mesh networks. In *Proc. of WiMesh*, 2006.

[16] H. Luo, S. Lu, and V. Bharghavan. A new model for packet scheduling in multihop wireless networks. In *Proc. of ACM MobiCom*, 2000.

[17] N. B. Salem and J.-P. Hubaux. A fair scheduling for wireless mesh networks. In *Proc. of WiMesh*, 2005.

[18] T. Clausen, P. Jacquet, C. Adjih, A. Laouiti, P. Minet, P. Muhlethaler, A. Qayyum, and L.Viennot. Optimized link state routing protocol (OLSR). RFC 3626, Oct 2003.

[19] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris. Link-level Measurements from an 802.11b Mesh Network. In *Proc. of ACM SIGCOMM*, 2004.

[20] J. Padhye, S. Agarwal, V. Padmanabhan, L. Qiu, A. Rao, and B. Zill. Estimation of Link Interference in Static Multi-hop Wireless Networks. In *Proc. of IMC*, 2005.