

Experimental Characterization of an 802.11b Wireless Mesh Network

Stephanie Liese, Daniel Wu and Prasant Mohapatra
Department of Computer Science
University of California, Davis
Davis, CA 95616

{slliese, danwu}@ucdavis.edu, prasant@cs.ucdavis.edu

ABSTRACT

Wireless Mesh Networks are being deployed everywhere as an alternative to broadband connections. Their ease of setup and large coverage are attractive attributes. However, there are very few performance studies of mesh networks especially in the multiple channel arena. Our objective is to study the performance and characterize the 802.11b wireless mesh backbone as a linear topology with respect to multiple channel usage. We look at the relative performances of single and multiple channels, as well as the number of hops utilized. We introduce a number of communication flows into the network to study the interactions. Finally, we look at the physical placement of the antennas on an access point to determine its impact on performance. Several design decisions for the configuration of wireless mesh network deployments have been inferred from our experimental testbed.

Categories and Subject Descriptors: C.2.0 [General]: Wireless Communications

General Terms: Measurement, Performance

Keywords: Multi-Hop Wireless Mesh Network, Performance, Characterization, 802.11b, Multiple Channel

1. INTRODUCTION

Multi-hop Wireless Mesh Networks (WMN) are becoming a popular alternative in extending the typical Wireless Local Area Networks (WLANs) we use today. Mesh networks consist of heterogeneous wireless clients and stationary wireless mesh routers. Mesh routers are wirelessly linked to one another to form a large scale wireless network. Mesh networks can be used in a variety of applications, including enterprise networking, building automation, extending provider service coverage, and wireless community networks [1]. These networks have the advantage of low-cost incremental deployment, where wireless routers can be added at any time in proportion to the number of users. In addition, the lack of a wired infrastructure lends mesh networking technology to

“last-mile” solutions for Internet Service Providers (ISP) or simply to aid in increased coverage area.

The gaining popularity and low cost tetherless deployment could facilitate widespread adoption of WMNs. Thus arises the need for the design and efficient configuration of these networks. Several aspects, such as the number of hops, number of channels, node/antenna placement, throughput and delay requirements, and other application demands will have an impact on the performance and cost of these networks. Thus, it is important to have an accurate physical characterization of wireless mesh networks. As we will detail below, there has been much work in the simulation of these types of networks but less work on the analysis of a physical implementation. Simulations cannot always accurately account for many physical layer issues such as ground effect, antenna proximity, and variations in interference and delay. As we noticed in the process of performing our experiments, unexpected physical layer issues often arise that are not addressed in the simulated environment. Our work accounts for these discrepancies, reporting what can be achieved as opposed to what can be expected. In addition, measurements taken on existing wireless mesh networks often yield results affected by background user traffic. Our multi-hop wireless network was isolated from background traffic and background interference in order to provide more accurate numbers. By taking out the background traffic we eliminate their ability to hide the causes/effects of various parameters pertinent to our experiment.

In our experiments, we measure the performance of various topological decisions and channel assignments in order to better understand which would yield the best performing network. There are many aspects to consider in the design of a wireless mesh network. How large of a performance advantage can be achieved from multi-radio nodes? Which channel assignment yields optimal performance? How should multiple antennas be placed for a single node? How does an increase in users affect aggregate goodput? Our goal is to understand these questions and to quantify the results. The conclusions that can be drawn from our work will aid future design decisions in wireless mesh networks.

The rest of the paper is organized as follows: We will go over the related work and background concepts in Section 2. In Section 3, we will describe in detail the experimental setup and configurations. The results of the experiment are shown in Section 4. Finally we’ll discuss the impact of the results in Section 5.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

IWCMC’06, July 3–6, 2006, Vancouver, British Columbia, Canada.
Copyright 2006 ACM 1-59593-306-9/06/0007 ...\$5.00.

2. RELATED WORK

Wireless networks research can be classified into four main categories. These are single/multiple channel, single/multiple hop networks.

2.1 Single Channel, Single Hop networks

Vasan and Shankar [2] characterized the hotspot model using a wireless sniffer and trace-based analysis. They focused on instantaneous throughput by sending UDP flows from a few clients to a single server on the wired side.

2.2 Single Channel, Multiple Hop networks

Gambiroza et al. [3] used a simulation-based linear topology model. They placed APs in a straight line configuration with multiple clients off of each AP. Thus the farthest AP's client to the sink gets the least amount of throughput.

Aguayo et al. [4] deployed and studied an actual 802.11b mesh network. Their study focused on the whole mesh network, whereas ours is a detail study on a single specific path. By looking at a specific topology, we can infer specific things with less variation in the environment.

2.3 Multiple Channel, Single Hop networks

With a slight twist to the regular single hop model, there are now commercial products with two or more wireless radios. These radios can be tuned to different frequencies and in-effect give us a multiple channel network. Proxim Wireless Networks have such an access point [5]. The main advantage of having multiple radios is diversity in channel selection.

2.4 Multiple Channel, Multiple Hop networks

Bahl et al. [6] gave a list of standard problems in using multiple channels over a wireless network and experimental results on using multiple channels. Baiocchi et al. [7] simulated a multiple channel wireless network with one control channel and multiple data channels. Raniwala and Chiu [8, 9] designed an architecture for Mesh networks and gave some simulation and experimental data on various issues.

2.5 Signal Quality in Wireless Networks

The wireless medium is non-deterministic. Slight variations of the height, placement or tilt of a wireless antenna can distort the receiving or transmitting signals. Kotz et al. [10] studied the signal strength qualities with respect to wireless networks.

2.6 Wireless Distribution System (WDS)

A Wireless Distribution System is a communications protocol between access points. It allows two access points to forward traffic between each other without a wired connection. By only using WDS, we eliminate the routing protocol overheads of a dynamic approach. We limit the impact of variability by not using routing protocols that introduce routing discovery and periodic updates. Furthermore, since we consider non-mobile access points, the routing paths do not change unlike the mobile ad hoc networks (MANET).

3. EXPERIMENTAL SETUP

We used the dual-radio ORiNOCO AP-2000 Access Point (AP) [5] with the ORiNOCO 802.11b Classic Gold PC Card. Each access point has two of these PC cards, which we will



Figure 1: Photograph of one access point in our experimental setup. It includes an access point, two antennas, one inverter, one car battery and a custom-built wooden stand for the antennas.

refer to as card A and card B. We attached 5.0 dBi gain antennas to each external connector on the cards. On each access point, the two antennas were horizontally separated by four feet and elevated to four feet above the ground. Figure 1 is a picture of our setup for a single access point. Each access point was positioned 100 feet from its neighbor in a linear topology. Access points were linearly associated to one another using WDS links between adjacent neighbors.

To avoid contention for the wireless medium, the experiments were run in an interference-free environment (a parking lot). This location also served to reduce multipath effects because it was a fairly open space. By limiting the external effects, we were able to better control our results for later analysis.

The clients for our experiment were HP/Compaq nc6000 laptops running Fedora Core 3. Each laptop has an Intel Pentium M 1.6GHz processor, 512 MB of RAM and an Atheros AR5212 802.11abg NIC. We ran the experiments using Linux kernel 2.6.11 and MADWIFI wireless drivers CVS version from May 27th, 2005. Traffic was generated and measurements were taken by using the Thrulay [11] software package. The wireless interface rate is fixed at 11Mbps to take out the variability of modulation changes.

3.1 Channel Assignment and Resource Contention

During initial experiments, unfavorable results for non-overlapping channel assignments brought us to question whether radio-to-radio packet processing or wireless channel interference dominated the performance characteristics of a given WLAN. In other words, did it cost more to pass a packet to another radio on the same access point than it did to reuse the same radio on the same channel? To address this question, we ran a single communication flow over five, four, three, and two hop topologies, varying the channel assignment and number of radios used. For each hop test, we

shutdown unused APs to lower equipment interference.

We devised three channel assignments, the first we will call Two Cards Multiple Channels (TCMC), as seen in Figure 2(a). In the figure, the two wireless card interfaces are shown on the top and bottom of the APs. For TCMC, we used channels 1, 6, and 11 and varied them over the five hops. Each access point utilized both of its radios, each operating on a separate channel. The second channel assignment will be referred to as Two Cards One Channel (TCOC), shown in Figure 2(b). In TCOC, each access point utilized both radios, but each radio operated on the same channel. The last channel assignment was One Card One Channel (OCOC), where each access point used a single radio and the entire network ran on the same channel (Figure 2(c)).

By comparing TCOC and OCOC, we could quantify the effect of single card resource contention. Finally, we could compare TCMC to both in order to determine how much non-overlapping channels can improve network performance. See Figure 2 for an illustration of each of the basic topologies.

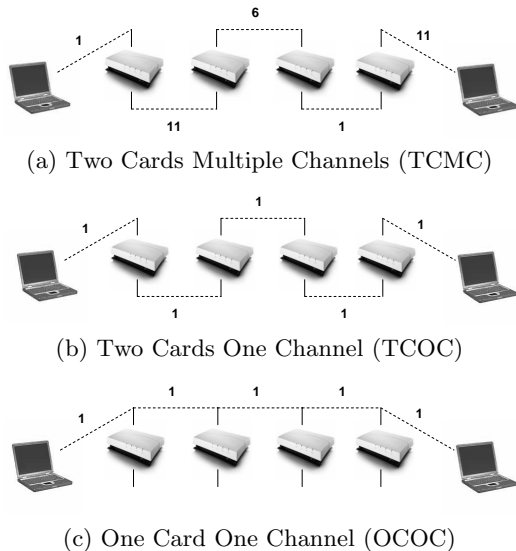


Figure 2: Different topologies used in our experiment with channel labeling.

3.2 Effect of Multiple Communication Flows

This experiment was designed to quantify the effects of multiple communication flows over a multi-hop wireless network. We define a communication flow in this paper as a TCP bulk data transfer between two laptops over any number of wireless hops. We used the TCMC channel assignment with non-overlapping channels. Three servers were positioned at the end of the linear topology (Figure 2(a)); at the other end, three clients communicated with a unique corresponding server. TCP bulk data transfers were run over the network using one, two, and three communication flows. The results of this experiment would help us determine the effect of multiple flows on aggregate goodput, single station goodput, and fairness in bandwidth allocation.

3.3 Effect of RTS/CTS

With a congested network as we see in the multiple flow tests, we questioned whether RTS/CTS would improve overall performance. Using the TCMC channel assignment, we ran one, two, and three communication flows with RTS/CTS turned on for all nodes (access points and laptops). The goal for this experiment is to quantify the performance drop when using RTS/CTS as opposed to not using it.

3.4 Antenna Proximity with Non-Overlapping Channels

As mentioned above, initial experiments using non-overlapping channels yielded unfavorable results. Our setup at that time did not include external antennas, but instead relied on the internal PC card antenna. For the ORiNOCO AP-2000, these PC cards are positioned less than one inch from one another on the board. In an attempt to account for our unfavorable results, we attached external antennas to the closely placed PC cards in order to vary the antenna proximity. We ran a single communication flow over two hops with the two antennas placed at zero, one, two, three and four feet from one another. One antenna ran on channel one while the other antenna utilized channel eleven. Being that both channels used were non-overlapping, one might expect goodput to remain constant despite the difference in space between the antennas. In fact, goodput varied greatly depending on that distance. These results would help us determine at what distance the antennas must be separated in order to avoid this kind of interference.

4. EXPERIMENTAL RESULTS

The experimental setup process also helped us learn and qualify some obvious and some not-so-obvious issues about WMNs. Initially, the APs were laid on the ground which yielded dismayingly poor performance because of ground interference. Thus, we used a custom-built wooden stand to provide elevation to the antennas. Although the improvement in link quality was expected, we were surprised at the magnitude of the difference. Similarly, the separation of the antennas, discussed below was also an interesting observation. In this section we graphically represent the data obtained for various configurations as outlined in Section 3.

Several experiments were performed and reconfirmed on separate days. For a given experiment, goodput and round trip time (RTT) were measured five times by using 20 second TCP bulk data transfers over the wireless medium. The results were averaged to get the data we provide in our graphs. Goodput is defined as the total data successfully transmitted (not including headers) over time. We chose to use goodput instead of throughput because it is an accurate measure of performance consistent with what is perceived by the user. RTT is defined as the time in between when the packet is sent and when the acknowledgement is received. We varied the number of radios used, number of hops, the channel assignment, the number of communication flows, and the antenna proximity.

4.1 Channel Assignment and Resource Contention

The goal of this experiment was to address whether the access point would suffer significantly by passing packets to another radio or if that processing time was negligible in

comparison to the multiple channel advantage. As seen in Figure 3, our results indicate that utilizing non-overlapping channels on multiple radios gives us a significant performance advantage to a single-radio implementation. Our TCMC channel assignment, which utilized three non-overlapping channels, obtained the best results with a five hop goodput of over 2 Mbps (see Figure 3(a)). In contrast, the goodput for a single radio implementation (OCOC) was around 1 Mbps, less than half of TCMC.

For fewer hops, the goodput of TCMC is much larger than that of TCOC or OCOC. For TCMC, there is a significantly larger drop in goodput seen between hops three and four. As shown in Figure 2, no channels are reused until the fourth hop. Once channel one is reintroduced into the network, goodput is roughly halved. This affect might be less if the two mesh routers had been in separate collision domains, as they would be for a real-world network. This pattern can also be observed in TCOC and OCOC. Between hops two and three, throughput is significantly reduced due to the first instance of channel reuse by the second AP.

The average RTT values are shown in Figure 3(b). The RTT with TCMC is significantly lower than that of TCOC or OCOC. Furthermore, with the increase in number of hops, rate of increase in RTT is relatively slower for TCMC compared to TCOC or OCOC.

To quantify the effect of switching radios, we ran two tests using the same channel assignment and one or two radios (TCOC and OCOC respectively). We found that the processing time to switch between the two cards on the access point was negligible in comparison to the wireless medium interference encountered in a single channel assignment. TCOC and OCOC yield roughly the same goodput and RTT (see Figure 3).

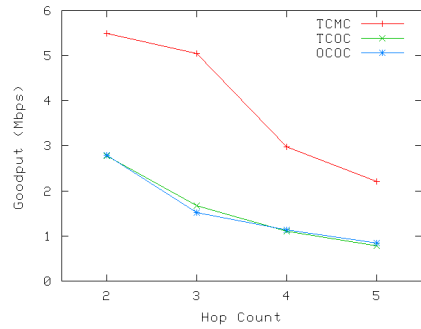
4.2 Effect of Multiple Communication Flows

In this experiment, we tried to identify the effects of multiple communication flows over a multi-hop wireless network. We questioned whether aggregate goodput might drop on a more congested network or if all clients would be given a fair portion of bandwidth. We used the TCMC channel assignment for this experiment as it yielded optimal goodput. As seen in Figure 4(a), aggregate goodput for multiple flows remained roughly the same as a single flow for a more congested network. In fact, there was even a slight overall increase in aggregate goodput when more than one communication flow utilized the network. With multiple clients contending for the channel, the channel will be less idle than if only one client was using it. The aggregate goodput drops with the increase in the number of hops because of the increase in likelihood of collision with multiple hops.

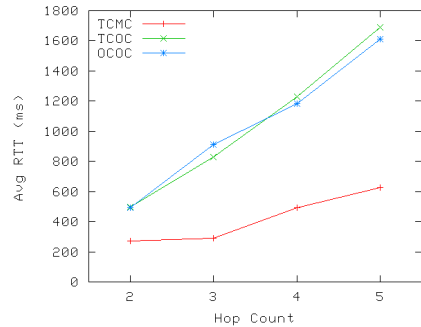
The average goodput for a given station was a function of the number of communication flows, as shown in Figure 4(b). The rate of drop in average goodput is also proportional to the number of flows.

We wanted to explore the fairness in bandwidth allocation across multiple flows in multi-hop networks. Figure 5 demonstrates bandwidth allocation across each communication flow. It can be observed that the bandwidth allocation across all flows is fair with respect to varying number of hop counts. We have shown the results for the two and three flow cases.

4.3 Effect of RTS/CTS



(a) Goodput



(b) Average RTT

Figure 3: Results of the varying topologies in Figure 2 with 1 flow.

With the addition of RTS/CTS into the network, we found that performance dropped by about 25% (see Figure 6(a)). RTS/CTS did not affect the network performance any differently based on the number of users in the network. In our experiment, each AP was in the same collision domain, thus performance did not suffer from hidden and exposed terminal problems. In a highly congested network with widely spaced mesh routers, one might expect the hidden and exposed terminal issues to affect performance. However, they would have to be a very significant problem in the network to warrant the potential 25% drop in goodput that its remedy incurs. As we can see in Figure 6(b) and Figure 6(c), using multiple flows degrades the average goodput and RTT because of TCP fairness.

4.4 Antenna Proximity with Non-overlapping Channels

The goal of this experiment was to measure the interference between adjacent antennas on non-overlapping channels. Our results indicate that even though the antennas operated on non-interfering channels, they still affect each other when placed too closely. We found that our 5.0 dBi antennas needed a distance of at least three feet between them before they no longer affected the performance of the other (see Figure 7). When separated by one foot or less, goodput was significantly impacted, dropping from roughly 5.5 Mbps to 1 Mbps.

This result is highly important when considering the design of a multi-hop wireless network, such as a mesh network. To improve performance, these networks might utilize multiple channels and multi-radio nodes, such as our

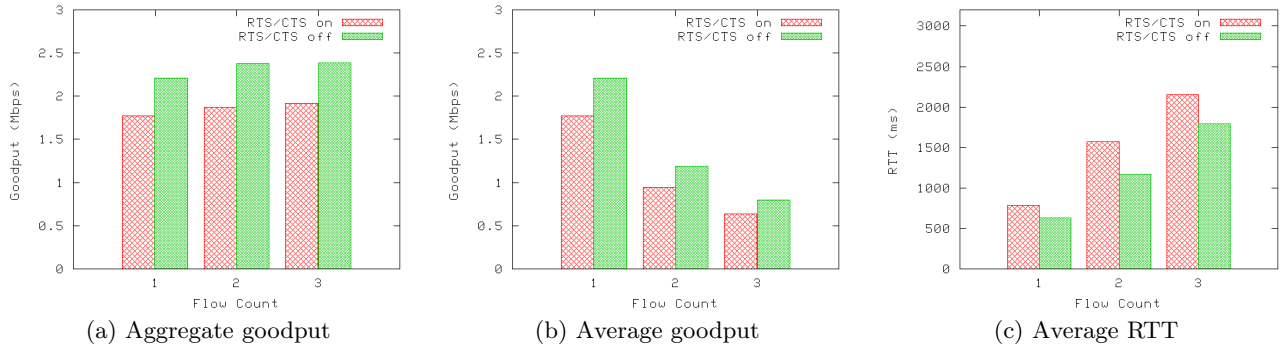
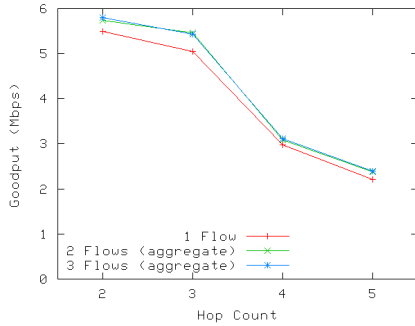
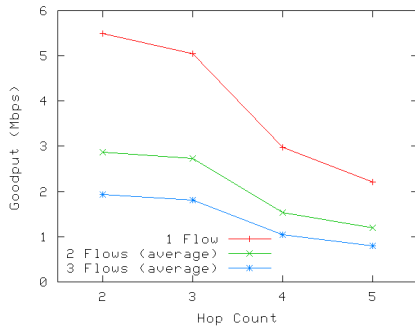


Figure 6: Effect of RTS/CTS over Five Hops.

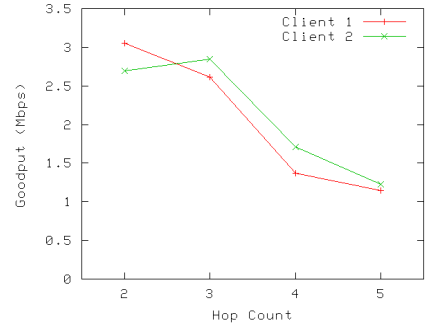


(a) Aggregate goodput of multiple flows.

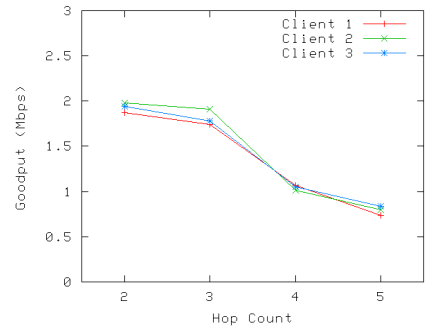


(b) Average goodput of multiple flows.

Figure 4: Multiple Flows in TCMC topology.



(a) 2 Clients



(b) 3 Clients

Figure 5: Individual Goodput for Multiple Flows in TCMC topology.

ORiNOCO AP-2000. If the nodes are deployed out-of-box or if consideration isn't taken into antenna proximity, the performance of these networks will be only a fraction of their potential. System administrators must be aware of the fact that non-overlapping channels are only non-interfering when the radios are sufficiently set apart. This may also affect the performance of wireless sensor networks, where nodes are so small that multiple antennas may not be able to avoid interfering with one another. Transmission power might need to be unusually low in order to combat this problem on very small nodes.

5. DISCUSSION AND INFERENCES

In this section, we discuss the inferences and design guidelines that can be derived from the experimental configura-

tions and the measurement data depicted in Sections 3 and 4.

Use multiple non-overlapping channels whenever possible. By measuring the affects of channel assignment, multiple radios, and resource contention, we have found that the processing overhead of multiple PC cards and channel conversion is negligible in comparison to the doubling of goodput seen by utilizing non-overlapping channels. Mesh networks should be designed with multiple channels in mind, as our results have shown. If limited to single radios, Quality of Service (QoS) may suffer but a reasonable goodput can still be obtained with a sufficiently powerful antenna.

Aggregate goodput remains roughly the same even when the number of users increase. In addition, individual user goodput is a function of the number of users in the network. Bandwidth is divided fairly evenly between

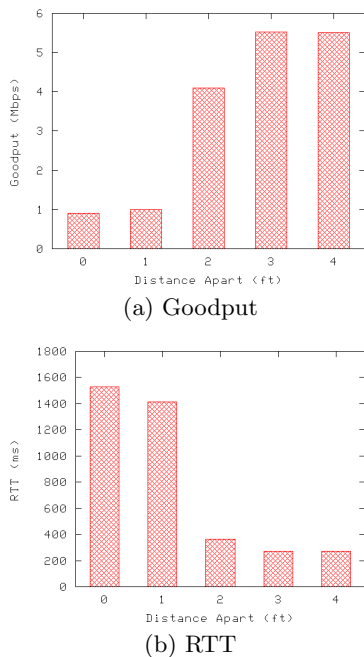


Figure 7: Effect of Antenna Proximity on a single AP

each user, given equal demands and equal distance from the source. Thus, network designers must weigh the achievable goodput with user QoS. Additional access points must be incrementally added on non-overlapping channels when user demands exceed achievable performance in a given location.

RTS/CTS decreases the goodput and increases RTT of the network. RTS/CTS severely handicaps any wireless network, particularly a mesh network where the wireless medium is so extensively used. The introduction of RTS/CTS into a mesh network must be warranted by a significant impact of hidden and exposed terminal issues in the network. If hidden and exposed terminal issues are not prevalent, RTS/CTS can degrade network performance up to 25%.

Multiple hops need to use multiple channels. When considering using multiple hops, if the adjacent links use the same wireless medium (channel), goodput will decrease because of the contention. By using different channels in adjacent links, we can essentially forward data without contending with the previous hop.

Limit the number of hops needed for each flow in the wireless network. When most traffic is destined to the Internet, the number of hops between a client and a gateway node must be limited. Additionally efficient flow control techniques could be deployed.

Antenna placement is very very important. Antennas on non-overlapping channels can still significantly impact one another's performance. With a gain of 5.0 dBi, at least three feet is needed between two antennas. For lower gains, smaller distances will likely be sufficient. Mesh network designers will not benefit from non-overlapping channels if the antennas are too closely placed. Proximity of the antennas to the ground is very important also. Too close to the ground, and the signal strength becomes degraded.

Round Trip Time (RTT) generally decreases when using multiple channels. RTT for the TCMC case is relatively better than the other topologies shown in Figure 3(b). Its important for RTT to be relatively low so the end-users can have a better perceived QoS. For mesh networks that are deployed in residential areas, customers would like near-instant access to Internet resources rather than wait 1 or 2 seconds for the information to show up on their screen.

TCP fairness does not account for far away clients. In a mesh network, the advantage was to be able to connect to the Internet or similar backbone through wireless. However, this advantage may be eclipsed by the fact that users who need to multi-hop through many mesh routers see degraded performance. This degradation of performance comes from the number of hops and the number of competing flows. With more users in a wireless network, there will be more contention and the users that are closer to the gateway to the Internet get an overall advantage than the users that are farther away. With this in mind, routes over a wireless network (linear or otherwise) must be chosen dynamically to bypass competing flows or else congestion will occur. Please refer to our technical report [12] for specific data regarding this observation.

6. CONCLUSION

In this paper, we have characterized an 802.11b wireless mesh network. We setup a 4 access point wireless mesh to gain some insight into its performance. The experiments included varying number of hops and flows in the network, as well as the number of channels and configurations. We quantified the goodput when multiple clients transmit their data to a single server. We have experimented with antenna proximity and have some interesting results. The inferences from our work will help the designers of multi-hop mesh networks to set various parameters and deploy nodes in various configurations.

7. REFERENCES

- [1] I. F. Akyildiz, X. Wang, and W. Wang, "Wireless mesh networks: A survey," 2005.
- [2] A. Vasani and A. U. Shankar, "An empirical characterization of instantaneous throughput in 802.11b wlans," University of Maryland, College Park, Tech. Rep., 2002.
- [3] V. Gambiroza, B. Sadeghi, and E. W. Knightly, "End-to-end performance and fairness in multihop wireless backhaul networks," in *MobiCom '04: Proceedings of the 10th annual international conference on Mobile computing and networking*. New York, NY, USA: ACM Press, 2004, pp. 287–301.
- [4] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris, "Link-level measurements from an 802.11b mesh network," in *SIGCOMM '04: Proceedings of the 2004 conference on Applications, technologies, architectures, and protocols for computer communications*. New York, NY, USA: ACM Press, 2004, pp. 121–132.
- [5] P. W. Networks, "Orinoco ap-2000 access point," 2005.
- [6] P. Bahl, A. Adya, J. Padhye, and A. Walman, "Reconsidering wireless systems with multiple radios," *SIGCOMM Comput. Commun. Rev.*, vol. 34, no. 5, pp. 39–46, 2004.
- [7] A. Baiocchi, A. Todini, and A. Valtetta, "Why a multichannel protocol can boost ieee 802.11 performance," in *MSWiM '04: Proceedings of the 7th ACM international symposium on Modeling, analysis and simulation of wireless and mobile systems*. New York, NY, USA: ACM Press, 2004, pp. 143–148.
- [8] A. Raniwala and T.-c. Chiueh, "Evaluation of a wireless enterprise backbone network architecture," in *High Performance Interconnects*, 2004.

- [9] —, “Architecture and algorithms for an ieee 802.11-based multi-channel wireless mesh network,” in *IEEE Infocom*, 2005.
- [10] D. Kotz, C. Newport, R. S. Gray, J. Liu, Y. Yuan, and C. Elliott, “Experimental evaluation of wireless simulation assumptions,” in *MSWiM’04*, 2004.
- [11] [Online]. Available: <http://thrulay-hd.sourceforge.net/>
- [12] S. Liese, D. Wu, and P. Mohapatra, “Experimental characterization of an 802.11b wireless mesh network,” University of California, Davis, Tech. Rep., July 2005. [Online]. Available: <http://www.cs.ucdavis.edu/research/tech-reports/2005/CSE-2005-16.pdf>