RICE UNIVERSITY

Measurement Driven Deployment of a Two-Tier Urban Mesh Access Network

by

Joseph D. Camp

A Thesis Submitted
in Partial Fulfillment of the
Requirements for the Degree

MASTER OF SCIENCE

Approved, Thesis Committee:

Edward W. Knightly, Chair
Associate Professor of Electrical and
Computer Engineering

David B. Johnson
Associate Professor of Computer Science
and Electrical and Computer Engineering

T. S. Eugene Ng
Assistant Professor of Computer Science
and Electrical and Computer Engineering

Ashu Sabharwal
Faculty Fellow of Electrical and
Computer Engineering

Houston, Texas
December, 2005

ABSTRACT

Measurement Driven Deployment of a Two-Tier Urban Mesh Access Network

by

Joseph D. Camp

Multihop wireless mesh networks can provide Internet access over a wide area with minimal infrastructure expenditure. In this work, we investigate key issues involved in deploying such networks including individual link characteristics, multihop application layer performance, and network-wide reliability and throughput. We perform extensive measurements in a two-tier urban scenario to characterize the propagation environment and correlate received signal strength with application layer throughput. Further, we measure competing, multihop flow traffic matrices to empirically define achievable throughputs of fully backlogged, rate limited, and web-emulated traffic. We find that while fully backlogged flows produce starving nodes, rate-controlling flows to a fixed value yields fairness and high aggregate throughput. Likewise, transmission gaps occurring in statistically multiplexed web traffic, even under high offered load, remove starvation and yield high performance.

Acknowledgments

Countless people have generously contributed academic, economic, physical, and emotional support enabling this achievement. My advisor, Edward Knightly, has given me the opportunity to work on the TFA-Wireless network and been an exemplary model as a researcher. Likewise, others on my committee, Ashutosh Sabharwal, Eugene Ng, and David Johnson, have played critical roles in my development as a graduate student. I am indebted to ELEC 438 students for performing measurements in the summer heat, especially Joshua Robinson, Chris Steger, Jeremy Beasley, and Aditya Nag. Additionally, Joshua and Chris have been instrumental in forming a deployment strategy for the network. Will Reed, CEO of Technology For All (TFA), has found the financial resources fund the project. He and all of the employees of TFA are shining examples of individuals whom use their giftings to help the less fortunate. My beautiful wife, Molly, has gracefully listened to me through the ups and downs of the deployment process and provided unconditional support and a helping hand. She has been and continues to be the greatest blessing in my life. I thank Christ Jesus, my Lord, who has further blessed me by allowing me to use my giftings to serve others in need. Finally, thank you to my parents, grandparents, and other members of my family and friends (from all over Texas) that He has used throughout my life to encourage me to explore the academic desires He placed within me.

Contents

List of Figures

List of Tables

Chapter 1
Introduction

Mesh networks provide high-bandwidth wireless access over large coverage areas with substantially reduced deployment cost as compared to fiber or wireline alternatives [1]. In a mesh network, fixed (i.e., non-mobile) *mesh nodes* are deployed throughout an area with a small fraction of the nodes featuring *wired* backhaul connections. The remaining wireless mesh nodes form a multihop wireless backbone to the nearest wired entry point [2].

Virtually unlimited applications exist for the use of mesh networks, from traffic control, transportation, and parking system inter-connectivity (the original motivation for the Philadelphia wireless project [3]), to Internet connectivity or building automation within a multi-story business or home or a grid of heat sensors that prevent forest fires [4]. The most prominent of these, of course, is providing high speed Internet access to communities and metropolitan areas. Although we focus here on the deployment of a two-tier mesh access network, our results are of merit for the deployment of any type of mesh solution. In a two-tier mesh network, an access tier provides a wireless connection between clients and mesh nodes, and a backhaul tier forwards traffic among mesh nodes.

For our deployment, we collaborate with a non-profit organization, Technology

For All (TFA), interested in providing affordable Internet access to a low-income community in Houston [5]. Thus, our design objectives encompass both performance objectives (pervasive Internet access at Megabits per second) and economic objectives (reduced deployment costs for low-income demographics). For TFA, deployment of a new wireline infrastructure is precluded due to costs of deploying new fiber as high as $200,000 per linear mile [6]. Moreover, monthly fees of existing wireline solutions such as DSL and CATV are beyond financial reach within these demographics.

In contrast, wireless technology provides an economically viable solution for low-cost deployment of broadband networks. Yet, even wireless architectures must inevitably connect to the wireline Internet via costly backhaul, averaging $750/month for T1 1.5 Mb/sec access and $10,500 for T3 45 Mb/sec access.* Consequently, we seek to minimize the number of wireline entries from the wireless access network to the backbone Internet.

Under these constraints, we utilize a two-tier urban mesh access network architecture similar to the network defined in [1]. By using off-the-shelf components in our deployment, we leverage the economies-of-scale that have driven the costs of IEEE 802.11 components to a price point that is feasible for low-income communities. Moreover, we aggregate all traffic to a single wireline backhaul point and use directional antennae to form long-haul links as needed. Here, we focus on how to

*http://www.pathfinderbandwidth.com, May 2005

most effectively deploy the aforementioned criterion to provide broadband Internet access to the Houston neighborhood. In this work, we use measurements in an urban environment to establish mesh performance predictions of single links (both access and backhaul) and multihop throughput falloff (in singly and concurrently active flow scenarios) that drive our deployment. Our contributions are as follows.

By conducting a series of signal strength and throughput measurements, we characterize important aspects of the propagation environment and use them to predict throughput and link-outage performance. We measured received signal strength and throughput over a large number of locations for both backbone and access links. Joining our empirical data with existing theoretical models, we calculate expected signal strength as a function of distance and correlate it with measured throughput. Our empirical link parameterization is novel in that it expresses both the expected throughput and the probability of falling below a minimum acceptable throughput (probability of outage) as a function of link distance. We show from these measurements that elevated nodes (10 meters in elevation) within an urban residential environment receive greater shadowing and multipath effects despite having less pathloss than nodes at the ground level. While we find it emperative to measure the pathloss of the particular propagation environment, we find that in our case, just 15 randomly selected measurement locations yield and average pathloss exponent with a standard deviation of 3 percent about the true value, and 50 measurements reduce the standard

deviation to 1.5 percent.

Next, we perform a broad set of application-layer throughput measurements for competing multihop flows. Existing measurements of single (non-contending) flows capture the basic effect of reduced throughput with increased path length. However, we show that application of such measurements to deployment decisions in a multi-flow environment would yield a large fraction of starving and disconnected nodes. In contrast, by driving the system with many concurrent, fully backlogged flows and concurrent web-emulated flows, we show that (i) starvation occurs for fully backlogged "upload" traffic due to the compounding effects of unequal flow collision probability and equal prioritization of each intermediate node's incoming traffic with *all* forwarded traffic; (ii) proper limiting of each mesh node's maximum rate alleviates starvation and provides near equal throughputs by masking MAC-layer unfairness; and (iii) even under modest to high offered loads, web traffic leaves sufficient free air time via statistical multiplexing and low activity factor to overcome the aforementioned starvation, even without rate limiting. Thus, we use the achievable traffic matrices above to drive placement decisions and show that our empirical definition of the multihop throughput distribution is essential in planning high-performance and cost-effective mesh networks.

Our work contrasts with existing mesh deployments in the following ways. Philadelphia's planned city-wide mesh deployment depends on exhaustive site surveys [3], and

the measurements are devoted exclusively to physical layer measurements of access links. The MIT Roofnet project also employs multihop mesh forwarding [7]. In contrast to our network, Roofnet has randomly placed nodes and a single-tier architecture, i.e., each node serves one in-building client instead of providing access to a large coverage area. Moreover, Roofnet's propagation environment is characterized by its strong Line-of-Sight (LOS) component whereas our links are generally heavily obstructed. A complete discussion of related work is presented in Chapter 5.

The remainder of the work is organized as follows. In Chapter 2, we describe our environment and methodology. Chapters 3 and 4 contain our link and multihop measurement studies. In Chapter 5, we contrast our work with existing literature. Finally, we conclude in Chapter 6.

Chapter 2
Background

2.1 Houston Neighborhood

We perform our measurements in a densely populated, single family residential, urban neighborhood with heavy tree coverage spanning 4.2 square kilometers (see Fig. 2.1). To quantify the density, the lot sizes within the neighborhood are 510 square meters on average.* The overwhelming majority of the homes within the neighborhood are one story wood-frame houses with an approximate height of 5 meters while sparsely placed two story homes have an approximate height of 7 meters. Trees vary in height throughout the neighborhood with heights up to approximately 20 meters. The population of this area is approximately 20,000 residents.

2.1.1 Neighborhood Demographics

The neighborhood of Pecan Park is one of Houston's most economically disadvantaged neighborhoods, e.g., with per-capita income approximately one-third the national average and 36.7% of children (under the age of 18) living below the poverty line according to the 2000 U. S. Census. Access to technology and Internet resources provides critical educational and job-training resources for a community in which 64.2% of the population of adults over the age of 25 are without a high school diploma or

*http://www.har.com, July 2005

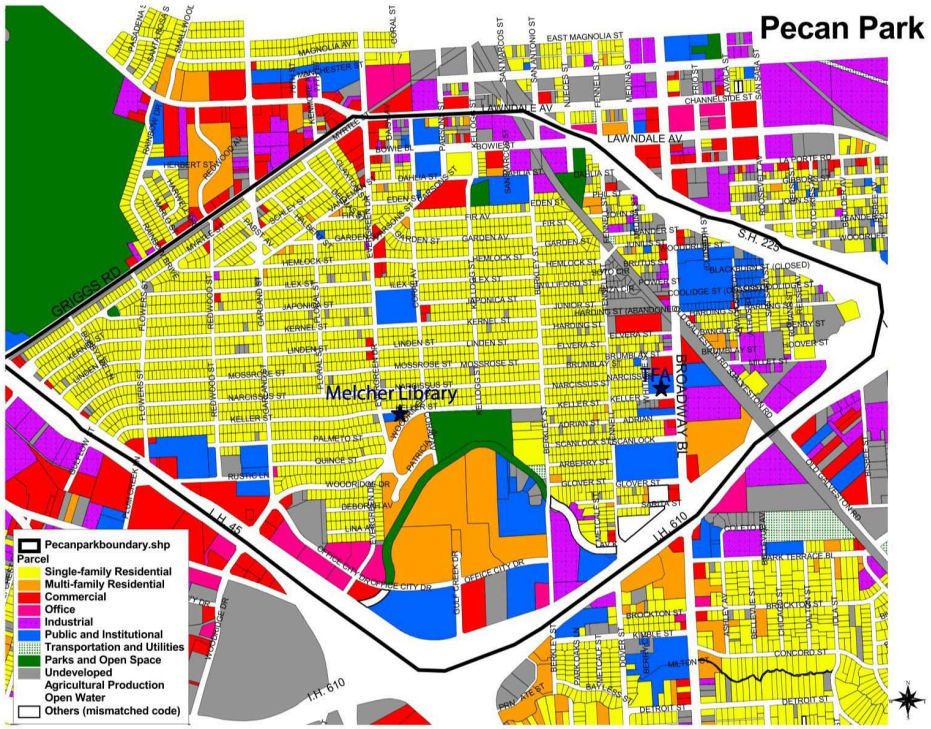

Figure 2.1 Pecan Park is a low-income neighborhood on the Southeast side of Houston covering 4.2 square kilometers.

GED equivalent.† A recent study indicates that high school graduation rates improve 6-8% with connected computers at home [8].

2.1.2 Technology For All

Technology For All was founded in 1997 with the mission of empowering low-income communities with the tools of technology. TFA creates educational, economic, and personal opportunities for low-income persons that can utilize the benefits of home Internet access. Through a program called "Learn and Earn," students complete

†Houston Independent School District, July 2003

a course and community service to earn a refurbished desktop or laptop. Through the "TFA-JobTech" program with startup funding from the U. S. Department of Commerce Technology Opportunity Program [9], livable-wage jobs are created in the homes of neighborhood residents via document conversion and coding performed via the Internet.* Additionally, the local community college and public library have offered TFA roof storage for antennas on their multi-story buildings.

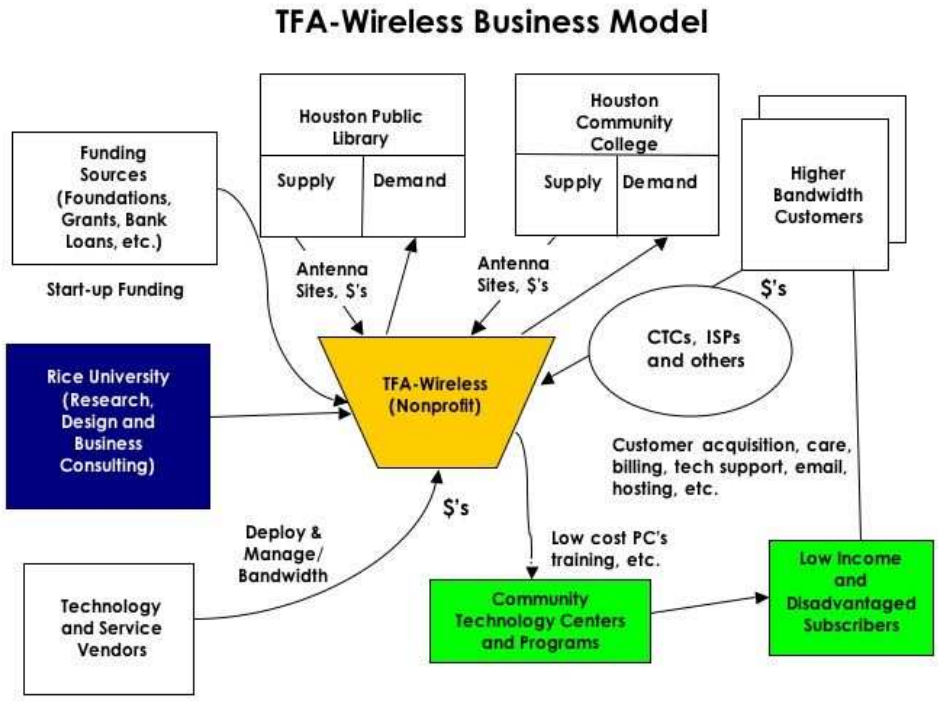

Figure 2.2 TFA's self-sustaining business model to offer free wireless Internet to low-income users

The TFA-Wireless business plan has established a sustainable business model for

*http://www.techforall.org/tfa_jobtech.html

providing broadband connectivity to residents of the neighborhood (see Figure 2.2). The model works by charging competitive market rates to customers needing higher bandwidth and commercial service levels while providing free lower bandwidth service (128 kbps) to low-income customers that have Houston Public Library cards.

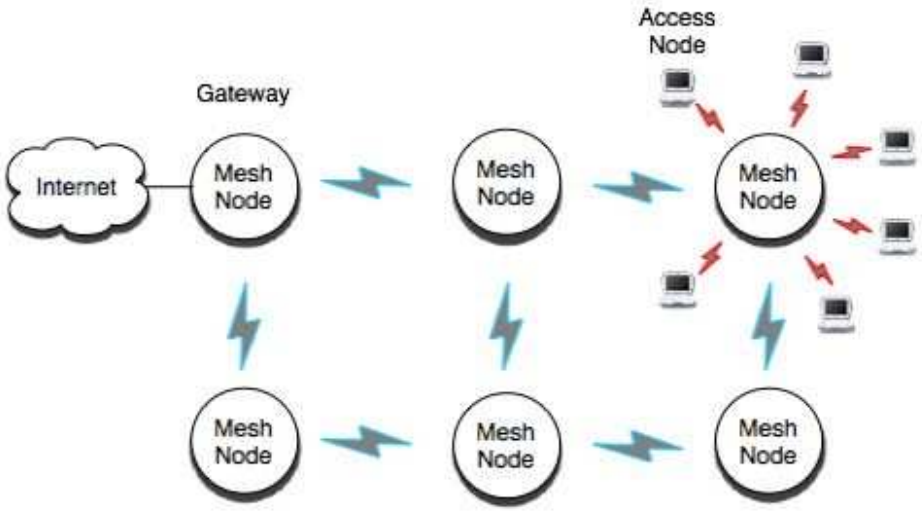

Figure 2.3 A two-tier network consists of infrastructure nodes which forward packets and client nodes which only originate or source traffic.

2.2 Two-Tier Architecture

In our measurement study and network deployment, we employ a two-tier network architecture as illustrated in Fig. 2.3. The *access tier* connects the client wireless device (e.g., a wireless laptop in a home or a wireless Ethernet bridge) to a mesh node. The *backhaul tier* interconnects the mesh nodes to forward traffic to and from wireline Internet entry points. Thus, the network provides coverage to all users within range of

the mesh nodes. In our single-radio deployment, both tiers are realized via the same radio and channel, and we employ traffic management techniques (rate limiting) to ensure proper division of resources between access and backhaul.

Mesh Node Hardware	Specifications
Processor	VIA C3 1 Ghz (x86)
Operating System	LocustWorld Linux 2.4.20
Executable Storage	32 MB Flash
Data Storage	5 GB Hard Drive
Wireless Interface	SMC 2532-B 802.11b card
Transmit Power	200 mW
Max Data Rate	11 Mbps
Wireless Driver	HostAP
Routing Protocol	AODV
Antenna	External 15 dB Omni
Antenna Height	10m

Table 2.1 Off-the-shelf hardware used for the experiments and deployment of our mesh network (approximately $500 per mesh node).

2.3 Mesh Hardware Platform

In choosing an off-the-shelf hardware solution, we desired a platform which was easily programmable yet had minimal startup and recurring cost. Thus, we performed a comparison to multihop wireless hardware solutions such as Motorola Canopy, Proxim, and Tropos Networks (not presented here). Due to the proprietary nature of their protocols they were both far more expensive and considerably less programmable than an Linux-based open-source mesh software solution created by LocustWorld.*

*http://www.locustworld.com

While we do not plan to use all of the protocols in the LocustWorld package indefinitely, we elect this immediate solution to focus on mesh deployment issues.

2.3.1 Backhaul Nodes

Therefore, our hardware platform for both the deployment and our reported measurements is as follows. For each of the mesh nodes, we use a VIA EPIA TC-Series mini-ITX motherboard with a VIA C3 x86-based processor running at 1 GHz (refer to Table 2.1. In the PCMCIA type II slot, we use an SMC 2532-B 802.11b card with 200 mW transmission power. The electrical hardware is housed in a NEMA 4 waterproof enclosure that can be externally mounted on residences, schools, libraries, and other commercial property. The mesh nodes run a minimal version of the Linux operating system which fits in the on-board 32 MB memory chip. As stated previously, we use an open-source version of the LocustWorld mesh networking software that uses AODV routing and HostAP drivers. The total cost of the nodes are approximately $500.

2.3.2 Access Nodes

We quickly found that penetration into the homes would be problematic for average (nominal powered) wireless devices. Thus, the client nodes within our network employ Engenius/Senao CB-3 Ethernet bridges which have 200 mW transmission power and 3 dBi external omnidirectional antennas. Additionally, the desktops and

laptops donated to residents within the neighborhood often lack wireless adapters. Thus, the bridge, costing approximately $100, serves a similar function as a DSL or cable modem in the home.

2.4 Mesh Antennas

In our deployment, we use two different types of antennas. We use omni-directional antennas both to connect to other mesh nodes (backhaul tier) and to provide service into homes (access tier). We use directional antennas for long-haul links between nodes across the neighborhood. By having concentrated beams over long distances, we have essentially created high bandwidth links of only one hop, thereby eliminating multihop throughput falloff at select locations deeper into the network. Additionally, we have set the two antenna types at these locations to use orthogonal frequencies in order to double the wireless bandwidth.

2.4.1 Omni-Directional Antenna

Each mesh node has a 15 dBi omni-directional antenna with an 8 degree vertical beamwidth.* Selecting antenna heights represents a tradeoff that is affected by the each region's particular propagation environment. At one extreme, a very high antenna elevation that clears all rooftops and trees has the advantage of providing strong Line-of-Sight (LOS) links for the backhaul tier. However, several problems arise with high antenna elevations in urban scenarios: (i) high attenuation of access links due to

*http://www.hyperlinktech.com/web/pdf/hg2415u_pro.pdf

tree canopies and buildings, (ii) requirement of multiple antennas or antennas with substantial energy focused both downward (access) and horizontally (backhaul), and (iii) legal and practical restrictions on maximum height. Likewise, while low antenna placement reduces deployment costs, it yields poor propagation paths for both access and backhaul. After completing experiments to balance these issues (not presented here), we selected an antenna height of 10 meters for the deployed nodes within the neighborhood.

2.4.2 Directional Antenna

The long-haul links have 24 dBi panel grid array directional antennas with a 3 degree beamwidth (typically 8 degree vertical beamwidth and 10 degree horizontal beamwidth according to the data sheet).* We used a simple methodology for the aligning of the directional antennas. For the vertical alignment, we used a level to align the pistol parallel to the ground and mounted the antennas at the same height. Houston is approximately flat. Thus, it is a reasonable assumption that two points within a mile of each other will be approximately equal in elevation. We cross-checked this assumption with topographical software. For the horizontal alignment, we pointed the respective antennas in the approximate direction. Then, we used *iwspy*, the standard wireless tool within Linux, to obtain signal strength readings of the attached wireless card while adjusting the antenna horizontally. Lastly, we

*http://www.pacwireless.com/products/GD24_datasheet.pdf

adjusted the remaining directional antenna in a similar manner.

Chapter 3
Link Measurements

In this chapter, we present the results of our measurements of single-link performance of mesh nodes in our neighborhood. The measurements represent both access and backhaul links and include received signal strength and throughput over a range of distances. We match our data to theoretical models to find a pathloss exponent and shadowing standard deviation so that we can accurately determine the range and reliability of the mesh links. As there are no accepted theoretical models for throughput, we introduce an empirical mapping between signal power and achievable throughput.

3.1 Theoretical Predictions

The multiplicative effects of the wireless channel are divided into three categories: pathloss, shadowing, and multipath fading [10]. In this work, we focus on pathloss and shadowing because they are the most measurable and predictable effects. Multipath fading produces dramatic variations in signal power, but the variations happen on such small scales of time and space that predicting them is prohibitively complex.

Pathloss describes the attenuation experienced by a wireless signal as a function of distance. Extensive prior empirical modeling indicates that signal power decays exponentially with distance according to a *pathloss exponent* that is particular to

the propagation scenario [11]. Pathloss exponents are dependent on the location and composition of objects in the environment and therefore add site-specificity to channel characterization. Pathloss is a very coarse description of a propagation scenario that allows us to generalize between environments that are similar but not identical. The received signal power P_R is expressed in terms of transmitted power P_T, transmit antenna gain G_T, receive antenna gain G_R, carrier wavelength λ, distance d, pathloss exponent α, and miscellaneous system loss L [11] as

$$P_R = \frac{P_T G_T G_R \lambda^2}{16\pi^2 d^\alpha L} \tag{3.1}$$

Shadowing describes the amount of variation in pathloss between similar propagation scenarios. For instance, within a single neighborhood, shadowing represents the difference between the signal power at different points with the same estimated pathloss. Prior measurements show that shadowing manifests as a zero-mean Gaussian random variable with standard deviation σ_ϵ added to the average signal power in dBm. The presence of shadowing makes any prediction of received signal power inherently probabilistic. In the following equation, shadowing is represented by ϵ and d_0 is a reference distance for which we have a measured power level [10].

$$P_{dBm}(d) = P_{dBm}(d_0) - 10\alpha log_{10}\left(\frac{d}{d_0}\right) + \epsilon \tag{3.2}$$

In the absence of scatterers or other attenuating media, the free space pathloss ex-

ponent is 2. With reflective and absorbent materials in the propagation environment, the pathloss exponent will increase. Pathloss exponents in outdoor environments range from 2 to 5 with a rough proportionality between the pathloss exponent and the amount of obstruction between the transmitting and receiving antennas [11]. The expected shadowing standard deviation, σ_ϵ, is approximately 8 dB [11], [10].

3.2 Access Link Measurements

3.2.1 Methodology

The following measurements characterize access links between clients (residences) and mesh nodes. The mesh node antennas are mounted at 10 meters while client nodes are fixed at a height of 1 meter. For a single fixed backbone node installation, we measure throughput and signal strength with the access node at many representative locations in the surrounding neighborhood. We use *iperf* traffic generator to create a fully backlogged UDP flow by connecting a laptop to the access node via Ethernet. We record signal strength measurements provided by the wireless interface in the mesh node.

3.2.2 Results

Fig. 3.1 depicts signal strength measurements as a function of link distance. Using a set of 138 measurements, we calculate an empirical pathloss exponent and shadowing standard deviation. The figure shows the theoretical pathloss curve and curves representing 1 and 2 standard deviations around the mean due to shadowing. In ad-

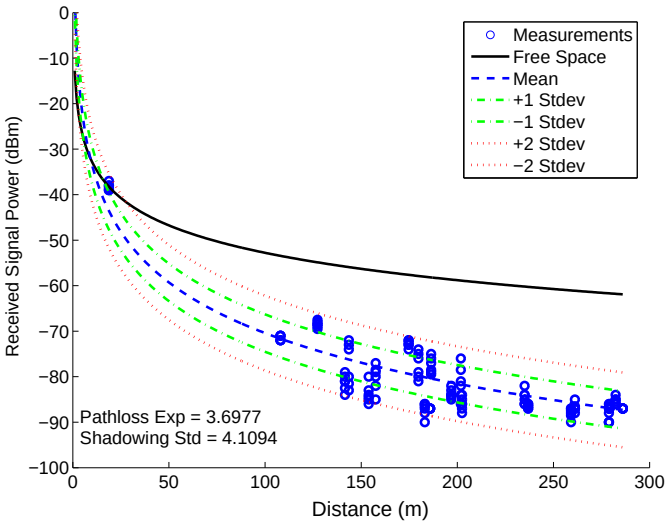

Figure 3.1 Empirical data and theoretical predictions for signal power received from an access node at 1 meter with a low gain antenna from a transmitter at 10 meters having a 15 dBi antenna.

dition, we plot the theoretical free space (unobstructed) pathloss curve for reference. Observe that the measurements taken at under 50 meters appear to be LOS. Our data shows a pathloss exponent of approximately $\alpha = 3.7$ and shadowing with standard deviation $\sigma_\epsilon = 4.1$.

Fig. 3.2 shows the distribution of the shadowing random variable, ϵ, along with a Gaussian random variable of equal mean and variance for comparison. We observe deviation from the predicted density that may be due to the presence of distinct pathloss exponents. However, because the standard deviation remains small, we can confidently predict expected signal levels based on our data.

Since we have no theoretical models to predict throughput performance as a func-

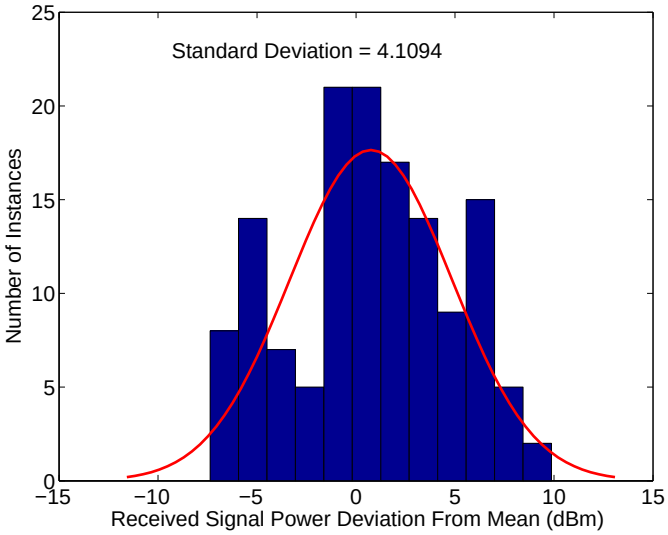

Figure 3.2 Empirical distribution of shadowing effects at an access node receiving signal from an elevated mesh node.

tion of environmental parameters, we measure UDP throughput as a function of signal strength. In Fig. 3.3, we observe that the throughput can be approximated as a piecewise linear function that is zero at all signal powers below −88.7 dBm and reaches a ceiling of approximately 6 Mb/s at −71.3 dBm. The minimum signal power at which we attain a mean throughput of 1 Mb/s is approximately −86 dBm and corresponds to 802.11b's 2 Mb/s modulation. According to the SMC wireless interface datasheet, 2 Mb/s is achievable at −93 dBm.* Hence, our data shows a performance 7 dB below nominal levels, which is likely attributable to multipath effects.

Studies such as [12] show that 802.11 systems originally designed for use indoors suffer significantly diminished performance in the presence of the large multipath

*www.smc.com

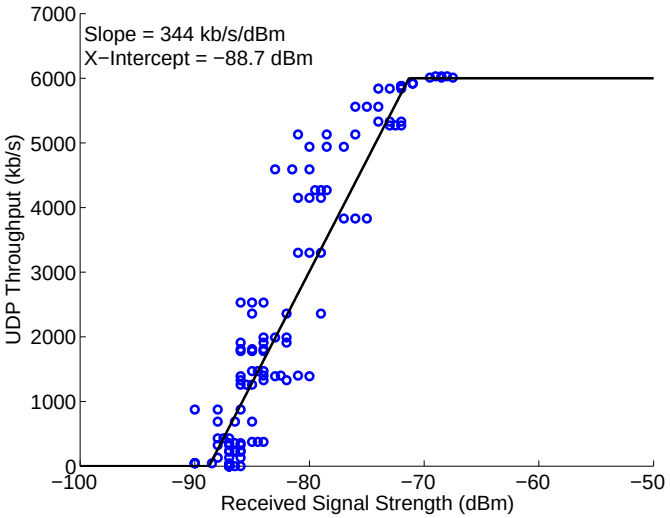

Figure 3.3 Measured UDP throughput (RTS/CTS off) received by an access node as a function of signal strength with a piecewise linear approximation.

delay spreads of urban environments [11], [13], [10]. When the delay spread is larger than the tolerances allowed by the system, we encounter intersymbol interference (ISI) in which a single signal collides with a delayed version of itself. Most wireless systems are designed to tolerate some amount of ISI, but can suffer extreme packet losses when delay spreads exceed expected levels [12].

3.2.3 Discussion: Repeatability

In Fig. 3.4, we show the measurement locations of access clients in respect to two fixed mesh node locations with a total of 138 trials of 60 second intervals. Again, we focus our access link measurements at distances of greater than 100 meters to find the critical range of clients within the network. While fixing both the height of these

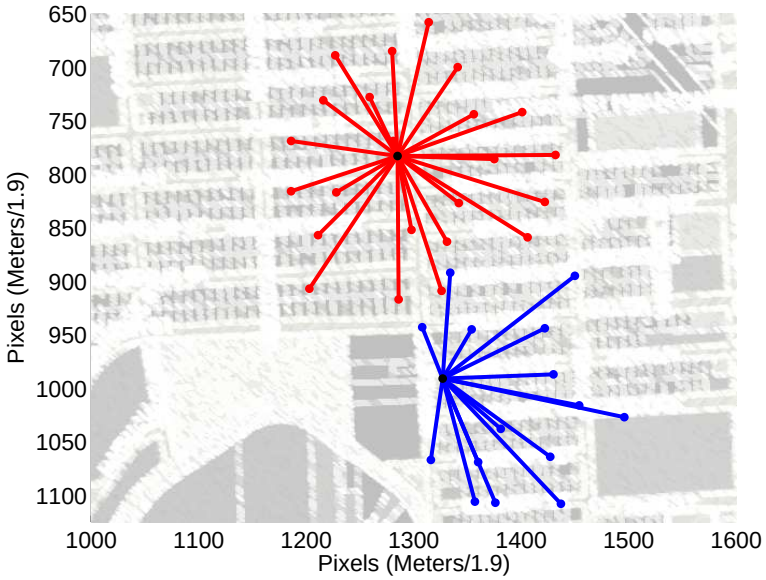

Figure 3.4 Access link measurements taken from two fixed mesh node locations.

access client locations and the alignment of the base of the node to be parallel to the ground, the nodes were randomly rotated about the horizontal axis at each location.

3.3 Backhaul Link Measurements

3.3.1 Methodology

We study the links among mesh nodes that form the backhaul network with measurements of links between a pair of identical 15 dBi antennas, both at 10 meters elevation. We use several fixed installations while moving a portable installation to numerous locations in the surrounding area. At each location, we generate UDP traffic over the link and record signal levels reported by the mesh node's wireless interface. Through initial experiments, we found that signal strength reached critical

levels at link distances of 200-275 meters. We choose a wide variety of measurement locations while focusing on this critical range.

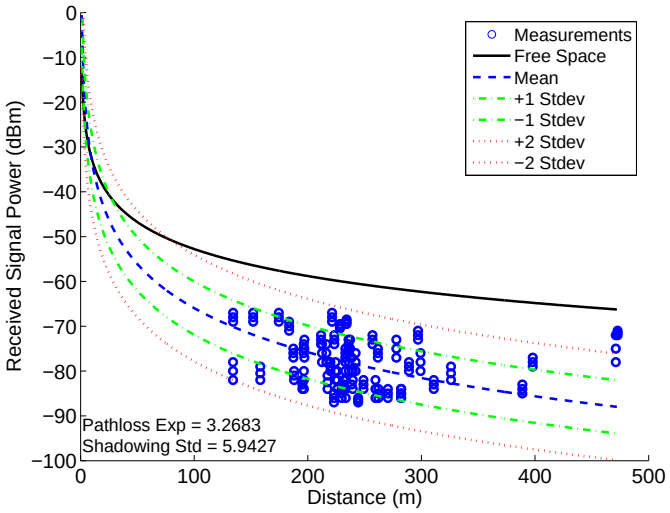

Figure 3.5 Empirical data and theoretical predictions for received signal power as a function of distance for backhaul links.

3.3.2 Results

Fig. 3.5 shows measured signal strength as a function of link distance along with the associated theoretical curves. Our measurements indicate a pathloss exponent of approximately $\alpha = 3.3$ and a shadowing standard deviation $\sigma_\epsilon = 5.9$ dBm. For purposes of comparison, note that [10] predicts a pathloss exponent of 4 and a standard deviation of 8 dBm for urban environments consisting of concrete and steel high-rise buildings, whereas our measurement environment features small, wood-frame houses and dense foliage. In comparison to our access link measurements, we observe that the

pathloss exponent improves with antenna elevation, which indicates that antennas at 10 meters are above most rooftops in the measurement neighborhood. In comparison to our access link measurements, we observe that the pathloss exponent improves with both antennas at the same elevation, which indicates that there is greater signal obstruction between the 10 meter antenna and the ground-level antenna, namely wooden house frames and rooftops. We also encounter measurement locations at a range of almost 500 meters that were nearly LOS, but they were highly atypical in our environment.

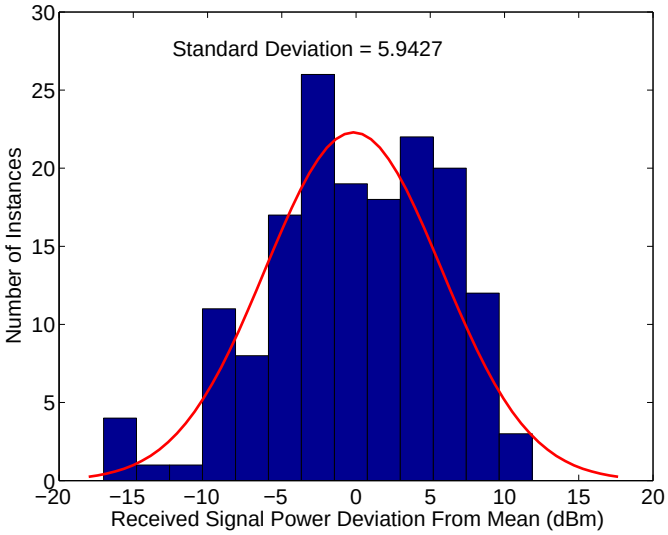

Figure 3.6 Empirical distribution of shadowing effects for backhaul links.

Fig. 3.6 plots the distribution of the shadowing variable, ϵ. We observe that it conforms well to the predicted Gaussian distribution which indicates that our average pathloss exponent value is representative of the neighborhood as a whole. In other

words, we have not measured two different sections of our neighborhood with unique pathloss exponents, which would have given us a bimodal distribution not conforming to the theoretical model. Thus, we can safely assert that the neighborhood is characterized by a single pathloss exponent and we can extrapolate from our measurements to make predictions for the entire neighborhood.

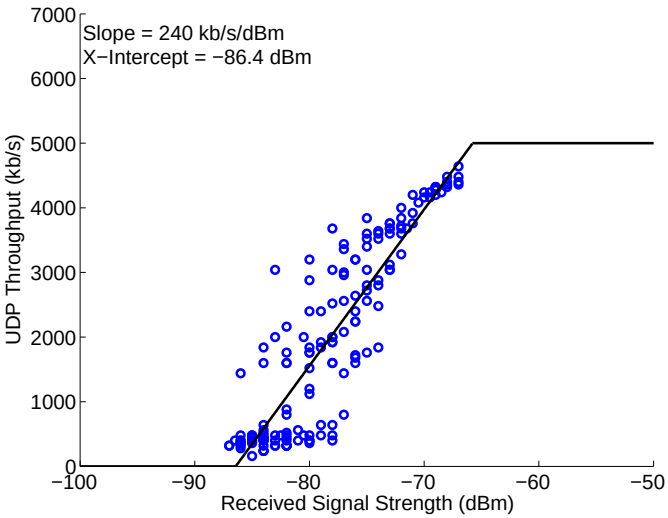

Figure 3.7 Measured UDP throughput (RTS/CTS on) over backhaul links as a function of signal strength with first order approximation.

Fig. 3.7 shows the relationship between UDP throughput and signal strength as well as a piecewise linear approximation of the data. With this approximation, the throughput in kbps at signal strength x is given as $\min(5000, \max(240x - 56.4, 0))$. We observe noticeable separation between our data and the linear approximation, and we find that the standard deviation is approximately 3 dB. The discrepancies are

most likely due to the wide range of delay spreads in our urban environment. Large delay spreads do not affect the signal power reported by the wireless interface, but they can increase packet loss rates dramatically. As a result, our mean throughput measurements fall approximately 11 dB below the nominal level. The multipath effects observed in our access and backhaul links are similar to those in [14]. This is significant because this means that using a signal-strength-to-throughput mapping obtained from the card's datasheet would result in throughput predictions an order-of-magnitude too high.

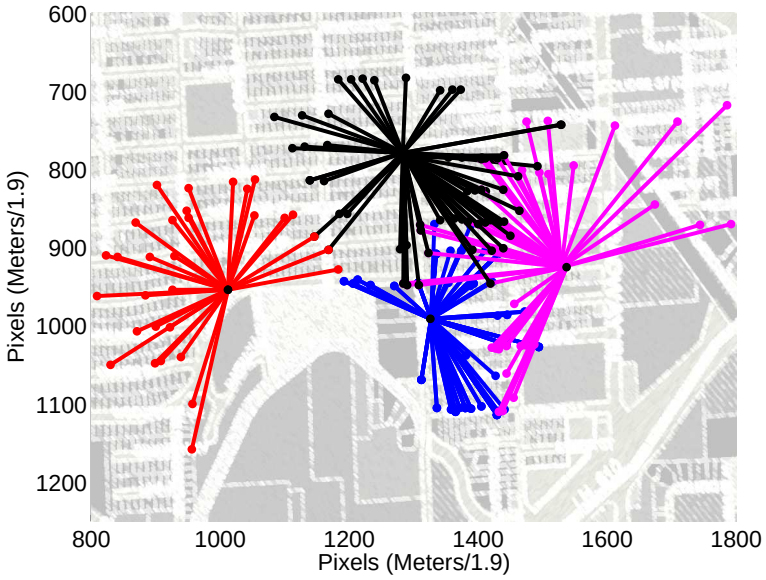

Figure 3.8 Backhaul link measurements from four fixed locations.

3.3.3 Discussion: Repeatability

In Fig. 3.8, we plot all of our backhaul measurement locations for a total of 235 trials of 60 second intervals around four fixed mesh nodes. We have biased our measurements slightly by avoiding locations with ranges less than 150 meters to focus on the outer limit of our mesh nodes' range. The irregular measurement patterns in Fig. 3.8 arise because we avoided measurement locations that we deemed highly atypical or unrealistic in our deployment scenario. Many of these locations were either inaccessible or located in large fields or parking lots. While the latter often presented good performance due to a strong LOS signal, they were not representative of the dense development and heavy foliage that dominates the deployment area, and they would have made our results overly optimistic. Both our access and backhaul measurements show high consistency and indicate homogeneity of our propagation environment. Further, our parameterization of the shadowing distribution accurately characterizes the deviation about our expected link performance. When combined, our performance functions thoroughly describe the important features of each link. Thus, since we found the links to be generally consistent throughout our environment, we can extend our measurements of individual link performance to describe multiple links in series and parallel. In the next two sections, we discuss extrapolating our understanding of single link reliability and throughput performance to complex multihop networks.

Chapter 4
Multihop Backhaul Experiments

In this chapter, we empirically determine achievable traffic matrices within a linear topology of nodes containing competing multihop flows. We show that (i) with no fairness mechanism and fully backlogged traffic, nodes with greater hop count starve; (ii) the RTS/CTS collision avoidance mechanism has an overall negative effect on per node throughput despite minimal gains in fairness; (iii) a simple static rate limiting scheme yields a fair multihop throughput distribution even with heavily loaded traffic patterns; (iv) web traffic yields sufficient idle times to significantly improve fairness and aggregate throughput in comparison to fully backlogged traffic.

4.1 Methodology

4.1.1 Parking Lot Traffic Matrix

We refer to the linear chain of nodes with traffic sourcing and sinking at the gateway node as the parking lot traffic matrix because it is analogous to the unfairness characteristics of multiple lines of vehicles leaving from one exit of a crowded parking lot. In an 802.11 parking lot, simulations indicate severe "spatial bias" in which nodes closest to the gateway obtain the highest throughput [15]. Ideally, there would be an equal per-node bandwidth distribution, i.e., a bandwidth share independent of location relative to the wire.

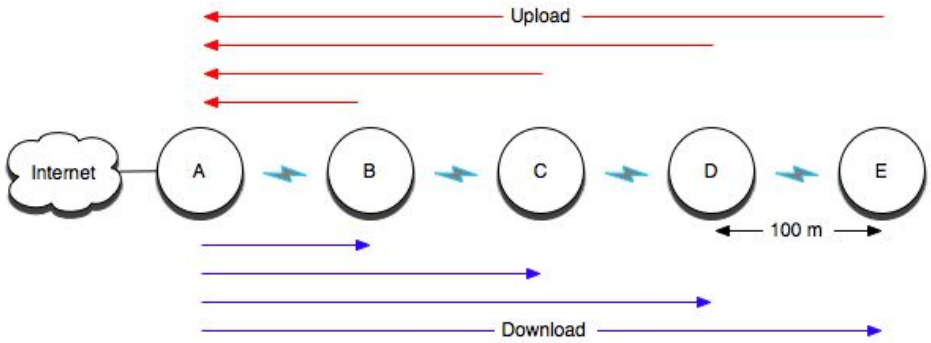

Figure 4.1 Nodes A through E are in a chain topology with A being the wired gateway node. All flows in the parking lot experiments are pictured here.

4.1.2 Experimental Set-Up

We construct a parking lot traffic matrix consisting of five wireless nodes contending for bandwidth in a single branch of the backhaul tree. We perform the experiments outdoors in the same physical environment as the measurements in Chapter 3 and using our mesh node hardware (refer to Chapter 2). To reduce the physical size of the parking lot, we opt for low gain (3 dBi) omnidirectional antennas mounted at approximately 2 meters. We space the mesh nodes to achieve a target signal strength of -75 dBm which is typical of a link between deployed mesh nodes. We perform each experimental trial at fixed locations, spaced approximately 100 meters apart. In Fig. 4.1, we illustrate the topology with every traffic flow of the parking lot experiments. In each test, we ensure that each mesh node will route data to its nearest neighbor only. That is, no node will send traffic directly to nodes that are two nodes away

in the chain. There is one wired gateway mesh node (node A in Fig. 4.1) for the topology. We use *iperf* sessions on the gateway and each of the nodes to generate TCP traffic for test intervals of 120 seconds.

4.1.3 Preliminary Experiments

Independently, we run *iperf* server-client applications from each node to its nearest with a fully backlogged queue to find the single hop link capacities. The physical layer rate is set to 11 Mbps on the wireless interface to remove autorate fallback effects. We find that the effective link capacity between nodes is 4 Mbps on each link along the chain.

	UDP		TCP	
Hop	Upload	Download	Upload	Download
1	3223	3255	4251	4047
2	1960	1652	1908	1618
3	1232	1001	877	1020
4	1120	989	774	902

Table 4.1 Absolute throughput for single active flows with UDP and TCP traffic (RTS/CTS disabled) for comparison with concurrent flow scenarios. Note that UDP and TCP experiments were performed on different test days and had different channel conditions.

4.2 Singly Active, Multihop Flows

We empirically obtain the achievable throughput of each multihop flow singly active both uplink and downlink to have a baseline for comparison with the competing multihop flow scenarios presented in the following sections of the chapter. For

completeness, we include the absolute values of the mean of all single active flow experiments in Table 4.1.3. The UDP experiments and TCP experiments were performed on different test days with different channel conditions. Thus, we normalize the end-to-end throughput of each flow with respect to the single hop throughput measured across the first link (AB in Fig. 4.1) from the wire to account for changing channel conditions across different test days.

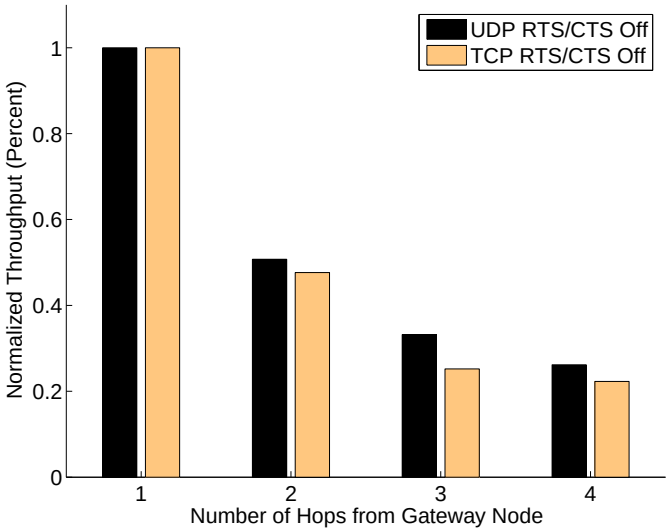

Figure 4.2 Single active UDP and TCP download flows (with RTS/CTS turned off) originating at the gateway node and destined for each of the mesh nodes normalized to the throughput of the first link.

4.2.1 Download

For the results in Fig. 4.2, each mesh node downloads a single active flow from the gateway node. The effect of the half-duplex link is evident as flows two hops from the wire have approximately half the bandwidth of one-hop flows. Within a

single contention neighborhood, one would expect that the three-hop flow would be one-third of the capacity, which is supported by our measurements. However, for the four-hop flow, spatial reuse (node E can send to D while node B sends to A; refer to Fig. 4.1) mitigates the throughput loss between three-hop and four-hop flows.

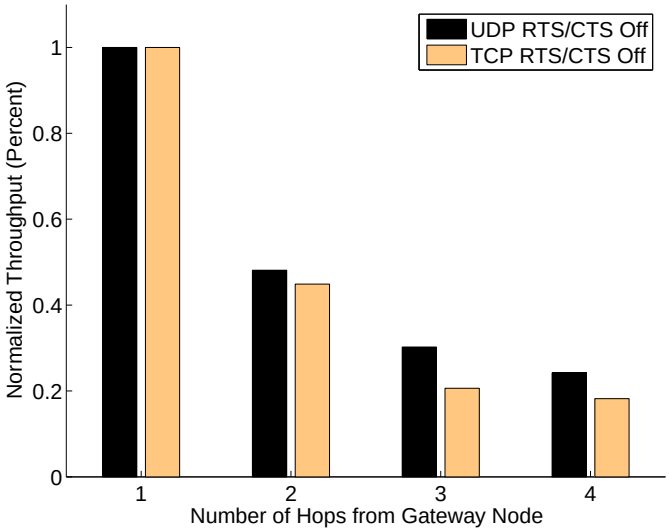

Figure 4.3 Single active UDP and TCP upload flows (with RTS/CTS turned off) originating at each of the mesh nodes and destined for the gateway node normalized to the throughput of the first link.

4.2.2 Upload

In Fig. 4.3, each flow is independently sourced at the mesh nodes and destined for the gateway node. We see patterns similar to the download measurements; the per-hop falloff for the upload direction is an average of 5% more severe than the download direction for UDP traffic and 11% for TCP traffic. We expect the observed similarity because directionality of flows is irrelevant when flows are singly active.

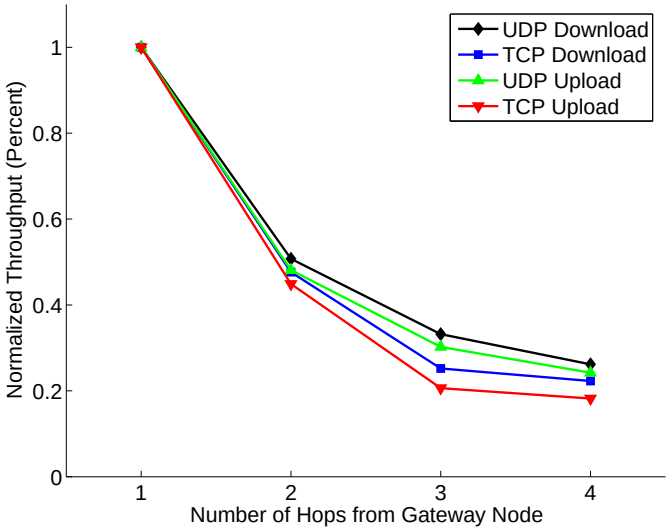

Figure 4.4 Single active UDP and TCP flows (with RTS/CTS disabled) to and from each of the mesh nodes and destined for the gateway node normalized to the throughput of the first link.

4.2.3 UDP vs. TCP Falloff

In both the upload and download case, the throughput falloff is not as pronounced for UDP traffic as TCP traffic (refer to 4.4. The reasons for this are twofold. Primarily, the additive increase, multiplicative decrease nature of TCP ensures that it will be sending below peak bandwidth most of the time. In contrast, our UDP experiments transmit traffic at the peak rate for the duration of the experiment regardless of packet loss. Additionally, TCP consumes bandwidth for acknowledgments in the reverse direction.

4.3 Fully Backlogged Parking Lot Experiments

We now investigate fairness trends of the fully backlogged parking lot traffic matrix where each node always has a packet to send along a linear topology (refer to 4.1). We show the unfairness of the download direction, upload direction, and both directions concurrently and compare each scenario with and without the RTS/CTS mechanism enabled to find its effect on fairness. Then we employ static rate limiting to eliminate starvation in the unidirectional parking lot traffic matrices.

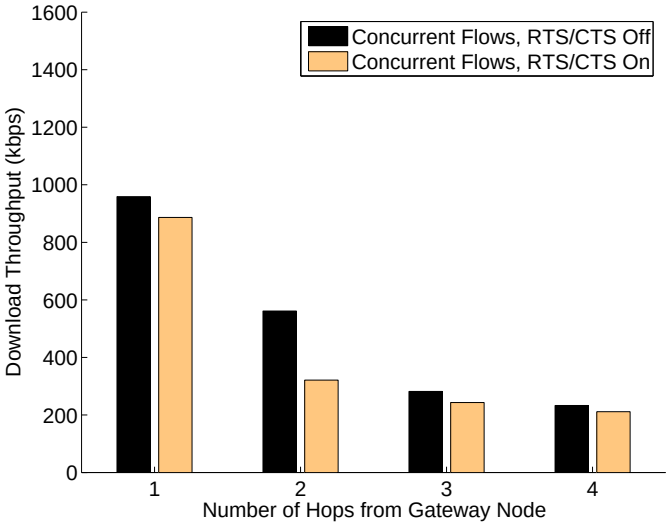

Figure 4.5 Concurrently active download flows (parking lot) originating at the gateway node and destined for each of the mesh nodes.

4.3.1 Download Traffic

Since the gateway transmits received packets in first-come-first-serve order, it fairly schedules the first wireless link among all four flows. However, as illustrated

in Fig. 4.5, unfairness occurs due to the forwarding overhead of multihop flows proportional to the hop count from the gateway. Even with fully backlogged queues, the last node in the chain receives only about 200 kbps with and without RTS/CTS.

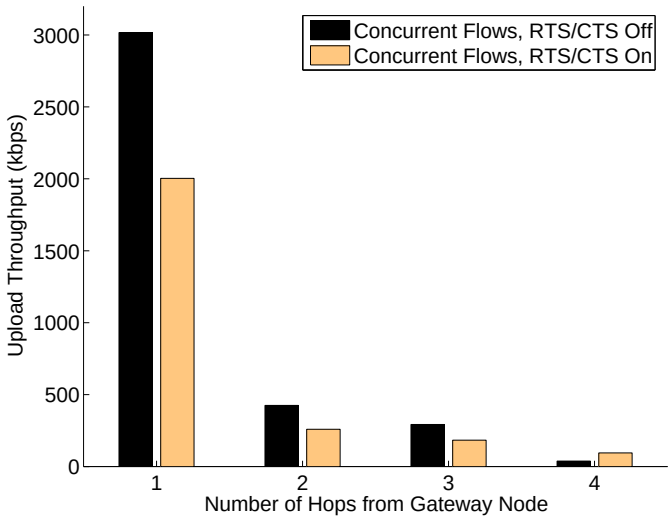

Figure 4.6 Concurrently active upload flows (parking lot) originating at each mesh node and destined for the gateway node.

4.3.2 Upload Traffic

We expect the upload traffic to have much worse fairness characteristics than download traffic. As flows are forwarded to the gateway, they capture less of the share of the upstream links. Each time the flow is forwarded there is a probability of loss due to collision which is compounded with increased hop count until finally there is a dissimilar distribution of the first link shares. Also, the MAC of each intermediate node, gives equal priority to its own traffic and forwarded traffic on the upstream

wireless link. Thus, the upload direction has the most exaggerated parking lot effect resulting in the greatest degree of spatial bias between the two traffic matrices. In Fig. 4.6, we encounter a pronounced falloff in bandwidth with increased hop count. In fact, the last node is starved in the scenario with RTS/CTS off.

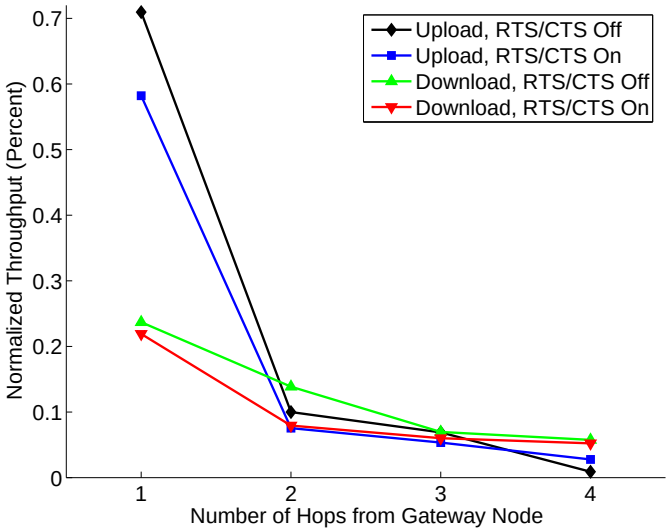

Figure 4.7 Comparison of normalized throughput in respect to the achieved throughput of a TCP flow across the first hop with RTS/CTS disabled.

4.3.3 Effect of RTS/CTS

In Fig. 4.7, we find in the fully backlogged upload scenario that the starvation of the last node in the chain is slightly mitigated – 3% of the bottleneck link as opposed to 1%. If node E sends an RTS to node D (see Fig. 4.1), D will notify E of the upstream channel conditions with or without a CTS packet in return. Thus, E is more aware of contention with RTS/CTS enabled. The same effect is not experienced

within the fully backlogged download scenario as node E (the node least likely to know channel conditions) simply receives data packets and does not have to contend for the channel. In fact, we find that the non-starved nodes within both traffic matrices receive approximately one-third less bandwidth on average due to overhead losses with the RTS/CTS mechanism. Therefore, we conclude that RTS/CTS overhead losses in non-starved nodes outweigh gains in starved nodes.

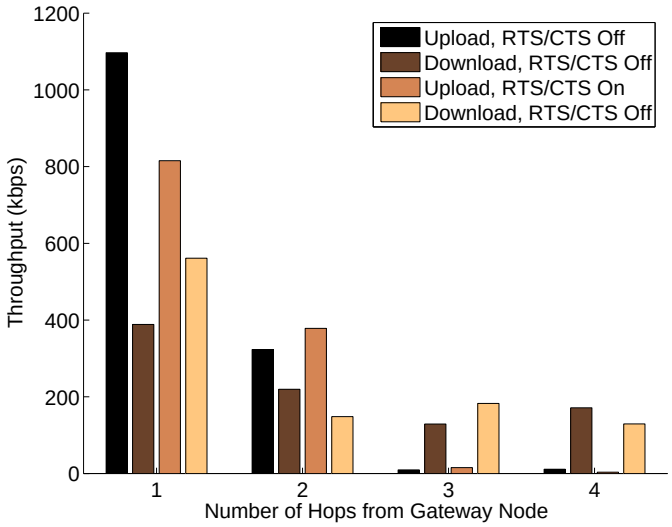

Figure 4.8 Parking lot scenario in both directions with and without RTS/CTS enabled.

4.3.4 Bidirectional Traffic

In Fig. 4.8, even in the highly bandwidth constrained scenarios, download traffic maintains some of the fairness patterns. However, the upload traffic obtains the majority of the bandwidth in the first two hops and starves the remaining downstream

nodes. Note that the fourth hop node has greater throughput than the third hop node for download traffic and RTS/CTS disabled. We found that if one of these two flows began first, the other flow would starve. Thus, we fairly alternate which flow started first, yielding unusual results.

4.4 Rate Limited Parking Lot Experiments

Our objective here is to supply each of the mesh nodes with an equal distribution of achieved throughput from the gateway node using static rate limiting. If we consider each of the single hop parts of one direction of the multihop flows pictured in Fig. 4.1, we find that there are 9 subflows. Thus, we expect the static rate limit that achieves our objective to be $\frac{1}{9}$ of the effective capacity of our links (approximately 4 Mbps). We use TCP traffic with the RTS/CTS mechanism disabled since we have shown this to be optimal for the fully backlogged parking lot.

4.4.1 Download Traffic

When we statically rate limit all nodes to equal shares of the medium, we find that the download fully backlogged parking lot scenario to be fair for values less than or equal to our expected $\frac{1}{9}$ of the effective capacity or 450 kbps (refer to Fig. 4.9). Similar to the fully backlogged experiment with no rate limiting, we expect this to have greater fairness than the upload scenario because each flow is equally shared over the first hop link.

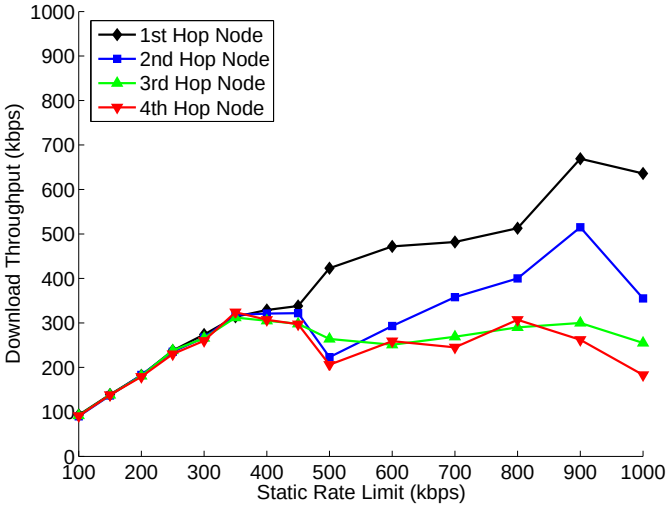

Figure 4.9 The fully backlogged parking lot traffic matrix downstream with each flow equally rate limited at the source.

4.4.2 Upload Traffic

In Fig. 4.10, we find that the same rate limiting scheme of each flow being equally limited at the source is less effective for upload fully backlogged traffic. Even at 450 kbps, the last node in the chain achieves only 60 percent of the throughput of the second node in the chain from the gateway. While each flow has equal rate at the source, as the flows are forwarded the two aforementioned effects remain: (i) with each forwarded subflow there is some probability of loss due to collision and (ii) nodes forwarding traffic give equal priority to their own traffic as all forwarded traffic. Thus, the more times a flow is forwarded the lower the share that flow will obtain on the first hop link, resulting in spatial bias even when the rate is equal at the source.

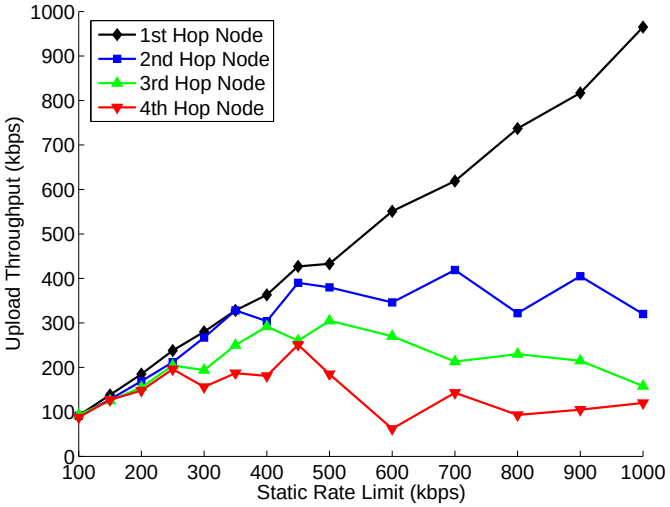

Figure 4.10 The fully backlogged parking lot traffic matrix upstream with each flow equally rate limited at the source.

4.5 Web-emulated Parking Lot Experiments

While fully backlogged traffic without rate limiting produced dismal results for multi-hop flows, here we study the ability of variable rate traffic with idle times (i.e., web traffic) to enable greater fairness due to (i) increased transmission gaps that can exploited by under-served flows and (ii) statistical multiplexing gains even with no external rate control. Two factors will influence the spatial bias within the web parking lot traffic matrix: the fair share of the bottleneck link and the activity factor of flows within a clique.

Parking Lot Load

If each link in a linear topology has equal capacity (and the wired connection to the Internet is not the bottleneck), the wireless link from the gateway node to its first neighbor will always be the bottleneck in a parking lot traffic matrix. Thus, we define the parking lot load, ϕ, as the utilization of each node relative to its fair share of the bottleneck link, $\frac{C}{i}$, where C is the effective wireless link capacity and i is the total hop count of the chain. If ϕ is less than 1 for each of the nodes, each of the offered loads are satisfied and fairness is achieved. Conversely, if ϕ is greater than 1 for each of the nodes, the system will experience spatial bias as nodes closer to the gateway capture greater bandwidth than the fair share of the bottleneck link. We define offered load in terms of the mean request size in bits, s, the number of users accessing each mesh node, N, and the mean inter-arrival time between requests, T. We express the parking lot load of a node as the offered load, $\frac{Ns}{T}$, over the fair share, $\frac{C}{i}$,

$$\phi = \frac{Nsi}{TC} \tag{4.1}$$

Activity Factor

We find the activity factor within a web parking lot traffic matrix to identify the mean number of flows that are contending for the channel. We can expect far greater channel utilization if there are no contending flows within the topology. Denoting the

round-trip times as rtt_i, the total hop count of the chain m, and the total number of round trips necessary to fulfill a request as τ, we define the activity factor, f, as

$$f = \frac{Nm\tau}{T} \sum_{i=1}^{m} rtt_i \tag{4.2}$$

4.5.1 Experimental Set-Up

In our experiment, we use the previously described physical layout of the chain topology. Each mesh node uses a web-emulation program to download web pages from a server on the Internet at exponentially distributed inter-arrival times between requests. Each session will request a 28 KB web page to download with a mean delay between requests of 7 seconds, the average think time between clicks [16]. The number of users will be varied from 5 to 80 users per node, but held constant for the two minute duration of each trial. We have accounted for emulation processing delays within our measurements to ensure behavior matching Poisson distributed requests of the number of mesh users associated to each mesh node. We have chosen to emulate web traffic on nodes B through E because we assume that the wired Internet is not the bottleneck. Thus, since the access tier bandwidth is not accounted for in our multihop experiments (that is, there is no access tier contention), if node A were also emulating web users, the only effect would be an increase in load over the wired connection to the Internet.

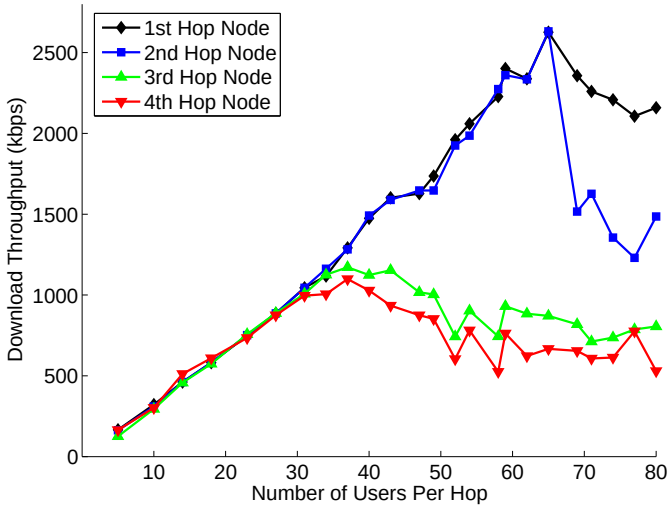

Figure 4.11 Each mesh node has an equal number of web-emulated users in the same linear topology.

4.5.2 Results

From the traffic load on each of the mesh nodes in our experiment (refer to Eqn. 4.1), we expect spatial bias with more than 30 users. Indeed in Fig. 4.11, we find that for traffic loads of up to 30 web-emulated users per mesh node (approximately 1 Mbps aggregate throughput) the mesh nodes maintain fairness between one another. However, as the number of users increases beyond this point, nodes closer to the gateway take an increasingly larger share of the available bandwidth from their downstream neighbors.

Surprisingly, the aggregate throughput is more than double that of the fully backlogged download scenario, and yet has much greater fairness (refer to Fig. 4.5). From Eqn. (4.2), we know that on average there is less than one flow contending for the

	Fully Backlogged Parking Lot TCP (Unidirectional)				Fully Backlogged Parking Lot TCP (Bidirectional)				Rate Limited Fully Backlogged Parking Lot		Web-Emulated Parking Lot (30 Users)
	Download RTS		Upload RTS		Download RTS		Upload RTS				
i	Off	On	Off	On	Off	On	Off	On	Down	Up	Down
1	0.237	0.219	0.710	0.582	0.096	0.139	0.258	0.192	0.085	0.107	0.232
2	0.139	0.079	0.100	0.075	0.054	0.037	0.076	0.089	0.081	0.098	0.232
3	0.070	0.060	0.069	0.054	0.032	0.045	0.002	0.004	0.075	0.065	0.224
4	0.058	0.052	0.009	0.028	0.042	0.032	0.003	0.001	0.074	0.063	0.221

Table 4.2 Empirically Defined Multihop Throughput Distribution, $\vec{\beta}$

channel in the download direction (upload traffic is less than 4 percent of the total bandwidth in our experiments) for up to 25 users. Thus, we attribute the throughput and fairness gain to the small web flows (approximately 20 packets in length) occurring on such small time scales that they are essentially singly active and suffer minimal contention from competing flows. However, as the web traffic demand increases, the behavior resembles that of the fully backlogged scenario.

4.6 Multihop Throughput Distribution, $\vec{\beta}$

We now empirically quantify the falloff of multihop throughput from the aforementioned parking lot experiments. We define the multihop throughput distribution, $\vec{\beta}$, which accounts for throughput loss at each hop when forwarding traffic over multiple wireless nodes due the effects of protocol overhead, contention, and half-duplex links as follows: *β_i is the fraction of end-to-end throughput an i-hop flow achieves with respect to the effective capacity of the first wireless link where i is the number of hops away from the wired node.*

For example, if a 2-hop flow achieves 400 kbps and the effective capacity of the wireless link is 4 Mbps, then $\beta_2 = 0.1$.

In IEEE 802.11 CSMA/CA networks, $\vec{\beta}$ depends on factors such as carrier sense range, RTS/CTS usage, the exponential backoff window, and type of traffic. With increased hop count, $\vec{\beta}$ should be monotonically decreasing as each node further from the wired gateway node achieves at most the throughput of its upstream neighbor. We expect the multihop throughput distribution to more gradually decrease at $i > 4$ because bandwidth from the gateway has fallen off to the point where the capacity of a single clique is no longer the bottleneck and flows can have spatial reuse.

4.6.1 Fully Backlogged Parking Lot

Table I shows $\vec{\beta}$ for upload, download, and bidirectional parking lot scenarios for TCP traffic with and without the RTS/CTS mechanism. Some of the β_4 values are not monotonically decreasing, which indicates that the flow from the fourth hop receives greater bandwidth than the flow from the third hop. This occurs because when flows from node E begin even slightly before flows from node D, node D is temporarily starved (refer to Fig. 4.1). Conversely, if the flow from node D begins before the flow from E, often E is starved. In order to avoid an artificial bias in our experiments, in half of the tests, we began the three-hop flow first, and in half, the four-hop flow.

In [15], upload parking lot traffic matrices are simulated in ns-2 with only the

immediate neighbor nodes in transmission range, and nodes two hops away are in carrier sense range. With TCP traffic and RTS/CTS disabled, the reported results are $\beta_1 = 0.382$, $\beta_2 = 0.200$, and $\beta_3 = 0.135$. Thus, we find a more extreme unfairness with our empirical measurements ($\beta_1 = 0.71$, $\beta_2 = 0.100$, $\beta_3 = 0.069$, and $\beta_4 = 0.009$). indicating that [15] is overly optimistic with respect to the fairness of 802.11 in parking lot traffic matrices. This can be attributed to the binary carrier sense range within ns-2 whereas our measurements indicate fluctuations in and out of carrier sense range due to fading channels.

4.6.2 Static Rate Limiting Parking Lot

We present $\vec{\beta}$ in the table for the upload and download scenario both statically rate limited at 450 kbps. Again, we see that rate limiting to equal rates works well in the download scenario. In the upload scenario, there is still some slight spatial bias as there is loss as traffic is forwarded upstream and the flows are unable to equally share the wireless link to the gateway node.

4.6.3 Web-Emulated Parking Lot

In Table I, we use $\vec{\beta}$ for web traffic at 30 users, the point at which the Parking Lot Load (Eqn. 4.1) is equal to 1. Again, the burstiness of web traffic reduces the number of competing flows within the parking lot traffic matrix. Thus, the throughput of the fair shares of web traffic is over twice that of the statically rate limited, fully

backlogged download scenario.

Chapter 5
Related Work

5.1 Physical Layer

Propagation and link-layer studies in the 2.4 GHz ISM band are generally concerned with indoor environments [17] [18] [19]. However, based on measurements at other frequencies in [11] and [10] and the limited literature available on outdoor 2.4 GHz channels [13], we know that outdoor propagation characteristics are very different from indoor scenarios. As a result, the nominal performance of hardware designed for indoor use may be severely degraded. In particular, the comparatively large delay spread values encountered outdoors can cause significant numbers of dropped packets even while SNR levels indicate a strong link [12]. In addition, propagation studies are generally strongly influenced by local topography and development such that even apparently similar locations can exhibit very different RF behavior [11]. Thus, it is difficult to employ any existing empirical results without first performing measurements in order to verify their applicability.

Much simulation and modeling work has been done on directional antennas in mesh networks such as neighbor discovery algorithms [20], deafness [21], scheduling [22] and MAC design [23]. Yet measurement studies with directional antennas have yet to be explored. We avoid these known problems by using orthogonal frequencies for our directional antennas in respect to the omni-directional backhaul network to

create stable, high throughput connections at two additional locations.

5.2 Capacity Models and Fairness

Many works explore capacity within mesh, WLAN, and ad hoc networks, most notably [24]. In [25], they expand the work from [24] for ad hoc networks with simulation and experiments in chain and lattice topologies. Capacity has also been considered for multi-radio, multi-channel mesh networks [26]. [27] presents an analytical study of ad hoc networks with base stations and finds that the number of base stations must increase as $\sqrt{n}$ where n is the number of wireless nodes in order to achieve an appreciable capacity increase. However, their results are not applicable to our deployment because they analyze a peer-to-peer traffic matrix.

Several simulation and analytical works examine fairness in mesh networks [28, 29, 30]. More specifically, Gambiroza et al. [15] use simulation to show unfair CSMA/CA parking lot scenarios and validate a multihop wireless backhaul fairness model. In contrast, we explore the unfairness of 802.11 by performing outdoor experiments on a parking lot topology.

5.3 Routing

Though routing is beyond the scope of this work as we have focused on measurement driven deployment, there have been extinsive work in this area on mesh networks. Our mesh nodes use AODV routing [31], a hop-by-hop routing scheme

which as the number of TCP connections increase, scales worse [**?**] than another on-demanding ad hoc routing protocol, DSR [32]. Additionally, AODV is hard to cache and fails to support unidirectional links well which are prominent in wireless environments [33]. A path metric called expected transmission count (ETX) enhances DSR and DSDV to find the route with the fewest expected number of transmissions (including retransmissions) in [34]. Results from the 29-node testbed show that longer paths experience a two-fold improvement with ETX. In [35], the Weighted Cumulative Expected Transmission Time (WCETT) expands upon the ETX metric to include multi-radio mesh nodes and multiple transmission bandwidths. We can draw from all of these findings in our future work as the network scales.

5.4 Mesh Measurements

There are several existing measurement studies on multihop wireless networks. As in [14], we measure link level performance, but for a significantly different environment (see Section 1). Other works focus on the evaluation of route metrics [34], mobility and route repair [36], and building ad hoc multihop wireless testbeds [37]. We differ from ad hoc multihop wireless in that our infrastructure is static and the traffic matrix is not arbitrary.

In [38], multi-radio interference effects are measured on a 2-hop mesh testbed. They found that due to RLC crosstalk interference between cards, no more than two cards are reasonable within one mesh node, agreeing with findings from the testbed

in [35]. Also, the channels of the respective cards should have a 35 dB spacing. Thus, we reserve the dual radio configuration for mesh nodes needing long haul links only (i.e. single radio for the majority of the mesh nodes) and use sufficient frequency spacing.

Our deployment strategy differs from the unplanned topology in [7]. Further, the authors consider only single flows active independently whereas we consider contention-based measurements to evaluate a mesh topology. [39] formulates the mesh planning problem in terms of placing wired gateway nodes within a fixed network and assumes a dense deployment of wireless mesh nodes. Further, their work is purely analytical in nature whereas we extensively measure single and multihop experiments in an urban propagation environment.

Chapter 6
Conclusions

In this work, we develop a measurement driven deployment strategy for a two-tier urban mesh access network. With our link measurements, we parameterize the wireless channel and correlate channel effects with throughput performance. We then generate a probabilistic model for link throughput and reliability as a function of distance. We find that our backhaul links will be the bottleneck within our urban mesh deployment. Broadly, we find that elevated or backhaul nodes within an urban residential environment receive greater shadowing and multipath effects despite having less pathloss than nodes at the ground level. In our multihop throughput experiments, we parameterize the multihop throughput distribution of competing, multihop flows. We find that in the fully backlogged parking lot scenario, simulations have less extreme spatial bias than CSMA hadware in an identical scenario [15]. Further, we find that rate-control mechanisms and web traffic eliminate starvation and yield high performance. For our deployment, we find that we can expect mesh nodes support up to 35 simultaneous web users fairly with the ability to burst beyond this with reasonable fairness characteristics. Using the empirical data gained from our single and multihop measurements, we can gain intuition on the performance of a broad class of CSMA mesh networks.

In ongoing work, we are using our measurement driven placement techniques for the Technology For All mesh network deployment in Houston. We will refine our models and study real traffic as the network expands and evolves to serve the immediate neighborhood and surrounding areas.

References

1. R. Karrer, A. Sabharwal, and E. Knightly, "Enabling large-scale wireless broadband: the case for TAPs," in *Proceedings of HotNets-II*, Cambridge, MA, Nov. 2003.

2. S. Lee, S. Banerjee, and B. Bhattacharjee, "The case for a multi-hop wireless local area network," in *Proceedings of IEEE INFOCOM '04*, March 2004.

3. D. Neff, "Wireless Philadelphia business plan: Wireless broadband as the foundation for a digital city," 9 February 2005.

4. I. F. Akyildiz, X. Wang, and W. Wang, "Wireless mesh networks: A survey," *Computer Networks (Elsevier)*, March 2005.

5. J. Camp, E. Knightly, and W. Reed, "Developing and deploying multihop wireless networks for low-income communities," in *Proceedings of Digital Communities*, Napoli, Italy, June 2005.

6. M. Weingarten and B. Stuck, "Is fiber to the home affordable?" *Business Communications Review*, June 2004.

7. J. Bicket, S. Biswas, D. Aguayo, and R. Morris, "Architecture and evaluation of the MIT Roofnet mesh network," in *Proceedings of ACM MobiCom*, Cologne, Germany, August 2005.

8. R. Fairlie, D. Beltran, and K. Das, "Do home computers improve educational outcomes? evidence from matched current population surveys and the national longitudinal survey of youth 1997," May 27 2005. [Online]. Available: http://econ.ucsc.edu/ fairlie/papers

9. D. of Commerce, "The Technology Opportunities Program," http://www.ntia.doc.gov/top.

10. G. Stuber, *Principles of Mobile Communication*, 4th ed. Boston: Kluwer, 2000.

11. T. Rappaport, *Wireless Communications, Principles & Practice*, ser. Emerging Technologies Series, T. Rappaport, Ed. Upper Saddle River, New Jersey: Prentice Hall, 1996.

12. C. Steger, P. Radosavljevic, and J. Frantz, "Performance of IEEE 802.11b wireless LAN in an emulated mobile channel," in *Proceedings of IEEE Vehicular Technology Conference*, vol. 2, April 2003, pp. 1479–1483.

13. G. Woodward, I. Oppermann, and J. Talvitie, "Outdoor-indoor temporal and spatial wideband channel model for ISM bands," in *Proceedings of IEEE Vehicular Technology Conference*, vol. 1, Fall 1999, pp. 136–140.

14. D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris, "Link-level measurements from an 802.11 mesh network," in *Proceedings of ACM SIGCOMM*, Portland, OR, 2004.

15. V. Gambiroza, B. Sadeghi, and E. Knightly, "End-to-end performance and fairness in multihop wireless backhaul networks," in *Proceedings of ACM MobiCom*, Philadephia, PA, September 2004.

16. S. Ranjan, R. Karrer, and E. W. Knightly, "Wide area redirection of dynamic content in internet data centers," in *Proceedings of IEEE INFOCOM '04*, March 2004.

17. H. Zepernick and T. Wysocki, "Multipath channel parameters for the indoor radio at 2.4 GHz ISM band," in *Proceedings of IEEE Vehicular Technology Conference*, vol. 1, Spring 1999, pp. 190–193.

18. C. Huang and R. Khayata, "Delay spreads and channel dynamics measurements at ISM bands," in *Proceedings of IEEE International Conference on Communications*, vol. 3, June 1992, pp. 1222–1226.

19. M. Yarvis, K. Papagiannaki, and W. S. Connor, "Characterization of 802.11 wireless networks in the home," in *Proceedings of Wireless Network Measurements (WiNMee)*, Riva del Garda, Italy, April 2005.

20. S. Vasudevan, J. Kurose, and D. Towsley, "On neighbor discovery in wireless networks with directional antennas," in *Proceedings of IEEE INFOCOM '05*, Miami, FL, 2005.

21. R. R. Choudhury and N. H. Vaidya, "Deafness: A mac problem in ad hoc networks when using directional antennas," in *Proceedings of IEEE ICNP '04: 12th IEEE International Conference on Network Protocols*, Berlin, Germany, 2004.

22. L. Bao and J. Garcia-Luna-Aceves, "Transmission scheduling in ad hoc networks with directional antennas," in *Proceedings of ACM MobiCom*, September 2002.

23. R. R. Choudhury and N. H. Vaidya, "On designing mac protocols for wireless networks using directional antennas," *IEEE Transactions of Mobile Computing (TMC)*, 2005.

24. P. Gupta and P. R. Kumar, "The capacity of wireless networks," *IEEE Transactions on Information Theory*, vol. 46, no. 2, Mar. 2000.

25. J. Li, C. Blake, D. S. J. De Couto, H. I. Lee, and R. Morris, "Capacity of ad hoc wireless networks," in *Proceedings of ACM MobiCom*, Rome, Italy, July 2001.

26. M. Kodialam and T. Nandagopal, "Characterizing the capacity region in multi-radio multi-channel wireless mesh networks," in *Proceedings of ACM MobiCom*, 2005.

27. B. Liu, Z. Liu, and D. Towsley, "On the capacity of hybrid wireless networks," in *Proceedings of IEEE Infocom*, San Fransisco, CA, April 2003.

28. Z. F. Li, S. Nandi, and A. K. Gupta, "Study of IEEE 802.11 fairness and its interaction with routing mechanism," in *IFIP MWCN 2003*, Singapore, May 2003.

29. X. Huang and B. Bensaou, "On max-min fairness and scheduling in wireless ad-hoc networks: Analytical framework and implementation," in *Proceedings of ACM MobiHoc*, Long Beach, CA, Oct. 2001.

30. T. Nandagopal, T. Kim, X. Gao, and V. Bharghavan, "Achieving MAC layer fairness in wireless packet networks," in *Proceedings of ACM MobiCom*, Boston, MA, Aug. 2000.

31. C. E. Perkins and E. M. Royer, "Ad-hoc on-demand distance vector routing," in *Proceedings of the 2nd IEEE Workshop on Mobile Computing Systems and Applications*, 1999.

32. D. B. Johnson and D. A. Maltz, "Dynamic source routing in ad hoc wireless networks," in *Mobile Computing*, Imielinski and Korth, Eds. Kluwer Academic Publishers, 1996, vol. 353.

33. D. Kotz, C. Newport, and C. Elliott, "The mistaken axioms of wireless-network research," Dartmouth College Computer Science, Hanover, NH, Tech. Rep. TR2003-467, July 2003.

34. D. De Couto, D. Aguayo, J. Bicket, and R. Morris, "A high-throughput path metric for multi-hop wireless routing," in *Proceedings of ACM MobiCom*, September 2003.

35. R. Draves, J. Padhye, and B. Zill, "Routing in multi-radio, multi-hop wireless networks," in *Proceedings of ACM MobiCom*, September 2004.

36. D. Johnson, "Routing in ad hoc networks of mobile hosts," in *Proceedings of the IEEE Workshop on Mobile Computing Systems and Applications*, December 1994, pp. 158–163.

37. D. Maltz, J. Broch, and D. Johnson, "Quantitive lessons from a full-scale multi-hop wireless ad hoc network testbed," in *Proceedings of the IEEE Wireless Communications and Networking Conference*, September 2000.

38. J. Robinson, K. Papagiannaki, C. Diot, X. Guo, and L. Krishnamurthy, "Experimenting with a multi-radio mesh networking testbed," in *Proceedings of Wireless Network Measurements (WiNMee)*, Riva del Garda, Italy, April 2005.

39. R. Chandra, L. Qiu, K. Jain, and M. Mahdian, "Optimizing the placement of integration points in multi-hop wireless networks," in *Proceedings of ICNP*, Berlin, Germany, October 2004.