

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

Joseph Camp, Joshua Robinson, Christopher Steger, and Edward Knightly
Department of Electrical and Computer Engineering, Rice University, Houston, TX
{camp, jpr, cbs, knightly}@ece.rice.edu

ABSTRACT

Multihop wireless mesh networks can provide Internet access over a wide area with minimal infrastructure expenditure. In this work, we present a measurement driven deployment strategy and a data-driven model to study the impact of design and topology decisions on network-wide performance and cost. We perform extensive measurements in a two-tier urban scenario to characterize the propagation environment and correlate received signal strength with application layer throughput. We find that well-known estimates for pathloss produce either heavily overprovisioned networks resulting in an order of magnitude increase in cost for high pathloss estimates or completely disconnected networks for low pathloss estimates. Modeling throughput with wireless interface manufacturer specifications similarly results in severely underprovisioned networks. 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. In comparison, we find that well-known noncompeting flow models for mesh networks over-estimate network-wide throughput by a factor of 2. Finally, our placement study shows that a regular grid topology achieves up to 50 percent greater throughput than random node placement.

Categories and Subject Descriptors: C.2.1 [Computer-Communication Networks]: Network Architecture and Design-Wireless Communication

General Terms: Measurement, Performance, Reliability, Experimentation, Design, Algorithms

Keywords: Mesh, Measurement-driven, Deployment, Wireless, Access Network, Urban, Pathloss, Fairness, Placement, Two-tier Architecture

This research is supported by NSF ITR Grants ANI-0331620 and ANI-0325971, and by the Cisco ARTI program.

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.

MobiSys'06, June 19–22, 2006, Uppsala, Sweden.

Copyright 2006 ACM 1-59593-195-3/06/0006 ...\$5.00.

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 *mesh nodes* are deployed throughout an area with a small fraction of the nodes featuring *wired* connections. 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 to the nearest wired Internet entry point.

In this paper, we present analysis of extensive field measurements of physical- and application-layer performance for access and backhaul links. We also present application-layer throughput measurements of contending multihop backhaul flows driven by multiple traffic types. Using this data, we develop a measurement-driven deployment methodology for two-tier mesh access networks: We outline the key measurement steps required for mesh access network deployments and characterize the impact of design decisions on both network-wide performance (connectivity, achievable traffic matrices, etc.) and cost (number of nodes and wires). We incorporate inherent variability in factors that cannot be precisely controlled such as the traffic matrix, link quality, and perturbations from *ideal* placement locations due to practical constraints. We also include the impact of traffic *control* strategies such as rate limiting nodes to improve fairness. All measurements are obtained in a Houston urban network that we are deploying in partnership with Technology For All (TFA).¹ The goal of the network is to provide affordable high-speed Internet access to low-income communities [2].

Our contributions are as follows. First, we use extensive measurements at various locations and distances to find our environment's pathloss exponent $\alpha = 3.3$ and shadowing $\sigma_e = 5.9$ (variation in signal strength at a known pathloss). We use linear regression to find the mean throughput as a piecewise linear function of signal strength (in dBm). We then empirically validate determination of link reliability at a given distance for a given minimum throughput threshold by using a Q-function to calculate the Gaussian tail probability. We show that accurate baseline physical layer measurements are essential for an efficient deployment: Using the maximum pathloss exponent of 5 for 2.4 GHz urban environments from [3] would yield networks that have a factor of over 9 times in overprovisioning (i.e., higher cost due to an increased number of nodes) whereas the minimum path loss of 2 would yield a completely disconnected network. Even the average pathloss of 3.5 from [3] has an overprovi-

¹<http://www.techforall.org>

sion factor of 55 percent and suggested urban pathloss of 4 from [4] has an overprovision factor of over 330 percent. We also show that an accurate throughput-signal-strength characterization is critical: a network planned using the manufacturer’s reported values overestimates the link range by approximately three times the appropriate value, and would result in a nearly completely disconnected network. While it is imperative to measure the pathloss of the particular propagation environment, we find that in our case, just 15 random measurement locations yield an 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.

Finally, we study node and wire placement and corresponding network topology issues with a novel mesh placement model. By using the single link and multihop measurement data, we incorporate effects of the physical layer, contention, MAC protocols, the hardware, etc. We explore system performance as a function of factors such as multihop traffic matrices, wire placement and density, mesh node density, and randomness in mesh node placement. Example findings are (i) regular grid structures have an average throughput up to 50% higher than randomly deployed topologies, (ii) adding an additional wired location to our network increases average throughput by a factor of up to 2.75, and (iii) regular grid deployments have no performance degradation with node perturbations up to $\frac{1}{6}$ the inter-node spacing.

Our work contrasts with existing mesh deployments in the following ways. Philadelphia’s planned city-wide mesh deployment depends on exhaustive site surveys [5], and the measurements are devoted exclusively to physical layer measurements of access links. The MIT Roofnet project also employs multihop mesh forwarding [6]. 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




Figure 1: A two-tier network consists of infrastructure nodes which forward packets and client nodes which only source or sink traffic.

are generally heavily obstructed. A complete discussion of related work is presented in Section 6.

The remainder of this paper is organized as follows. In Section 2, we describe our environment and methodology. Sections 3 and 4 contain our link and multihop measurement studies. We present our placement study in Section 5. In Section 6, we contrast our work with the existing literature. Finally, we conclude in Section 7.

2. TWO-TIER URBAN MESH SCENARIO

2.1 Two-Tier Architecture

In our measurement study and network deployment, we employ a two-tier network architecture as illustrated in Fig. 1. 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 or gateway mesh nodes. 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.

2.2 Houston Neighborhood

We perform our measurements in a densely populated, single family residential, urban neighborhood with heavy tree coverage spanning 4.2 square kilometers. The lot sizes within the neighborhood are 510 square meters on average.² The overwhelming majority of the homes within the neighborhood are one story 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.

2.3 Mesh Hardware Platform

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.

²<http://www.har.com>, July 2005

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. We use an open-source version of the LocustWorld³ mesh networking software that uses AODV routing and HostAP drivers. 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.

2.4 Mesh Antenna

Each mesh node has a 15 dBi omnidirectional 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 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.

3. LINK MEASUREMENTS

In this section, 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. In Section 5, we discuss the impact of our measurements on the performance of a larger system.

3.1 Theoretical Predictions

The multiplicative effects of the wireless channel are divided into three categories: pathloss, shadowing, and multipath fading [4]. 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 [3]. 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.

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 , α is the pathloss exponent, and d_0 is a reference distance for which we have a measured power level [4].

$$P_{dBm}(d) = P_{dBm}(d_0) - 10\alpha \log_{10} \left(\frac{d}{d_0} \right) + \epsilon \quad (1)$$

In the absence of scatterers or other attenuating media, the free space pathloss exponent 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 [3]. The expected shadowing standard deviation, σ_ϵ , is approximately 8 dB [3], [4].

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. Both of the wireless interfaces have *autorate* enabled to determine the physical layer transmission rate.

3.2.2 Results

Fig. 2 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 addition, 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 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

³<http://www.locustworld.com>

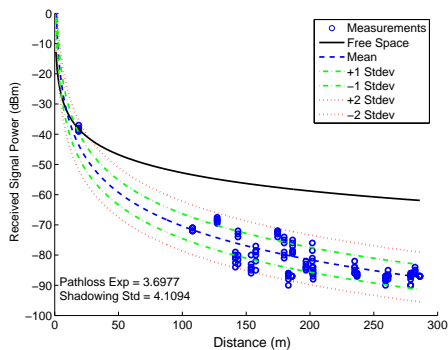


Figure 2: Empirical data and theoretical predictions for signal power received from an access node at 1 meter with a low gain antenna from a transmitter at 10 meters having a 15 dBi antenna.




Figure 3: Empirical distribution of shadowing effects at an access node receiving signal from an elevated mesh node.

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 function of environmental parameters, we measure UDP throughput as a function of signal strength. In Fig. 4, 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.⁴ Hence, our data shows a performance 7 dB below nominal levels, which is likely attributable to multipath effects.

Studies such as [7] show that 802.11 systems originally designed for use indoors suffer significantly diminished performance in the presence of the large multipath delay spreads of urban environments [3], [8], [4]. 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

⁴www.smc.com

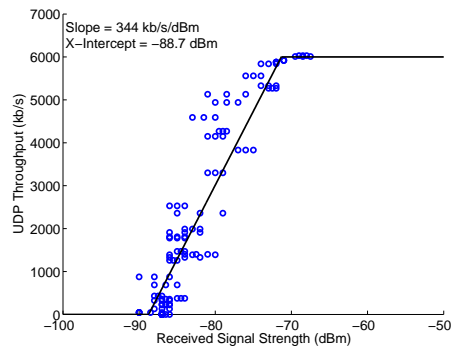


Figure 4: Measured UDP throughput (RTS/CTS off) received by an access node as a function of signal strength with a piecewise linear approximation.

can suffer extreme packet losses when delay spreads exceed expected levels [7].

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. Both of the wireless interfaces have autorate enabled to determine the physical layer transmission rate. 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.

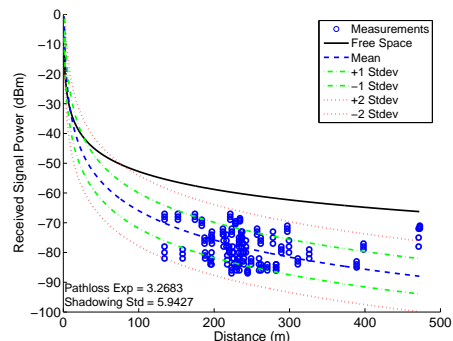


Figure 5: Empirical data and theoretical predictions for received signal power as a function of distance for backhaul links.

3.3.2 Results

Fig. 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 $\sigma_e = 5.9$ dBm. For purposes of comparison, note that [4] 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 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.

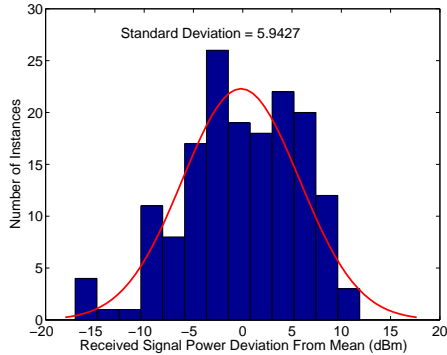


Figure 6: Empirical distribution of shadowing effects for backhaul links.

Fig. 6 plots the distribution of the shadowing variable, ϵ . We observe that our measurements of the pathloss exponent value within the neighborhood are Gaussian distributed about a single mean. Thus, the average pathloss exponent 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.

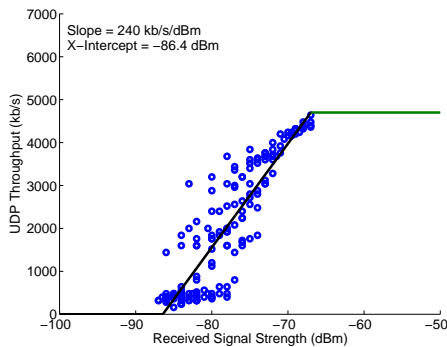


Figure 7: Measured UDP throughput (RTS/CTS on) over backhaul links as a function of signal strength with first order approximation.

Fig. 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 $\min(5000, \max(240x -$

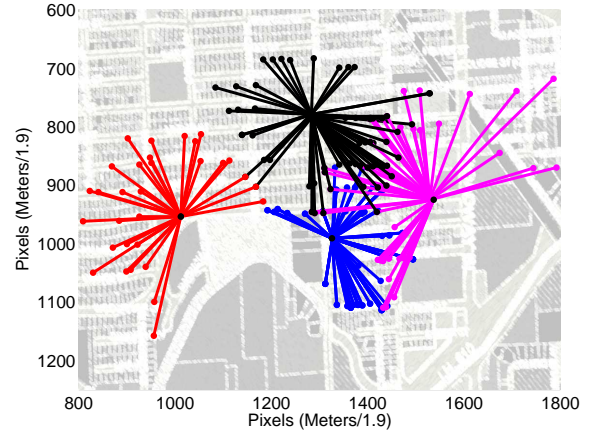


Figure 8: Backhaul link measurements from four fixed locations.

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 [9]. 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.

3.4 Discussion: Repeatability

In Fig. 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. 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. We similarly conducted 138 trials of 60 second intervals for the access link measurements at distances of greater than 100 meters (not pictured here). 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 sin-

gle link reliability and throughput performance to complex multihop networks.

4. MULTIHOP BACKHAUL EXPERIMENTS

In this section, 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 [10]. Ideally, there would be an equal per-node bandwidth distribution, i.e., a bandwidth share independent of location relative to the wire.




Figure 9: 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 Section 3 and using our mesh node hardware (refer to Section 2). To reduce the physical size of the parking lot, we opt for low gain (3 dBi) omnidirectional antennas mounted at approximately 2 meters high. 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. 9, 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. 9) 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. We additionally measure multihop, single active flow traffic (not presented here as it is well studied in the literature) to sufficiently plan the parking lot experiments and have a baseline for comparison within our placement study.

4.2 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 Fig. 9). 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.

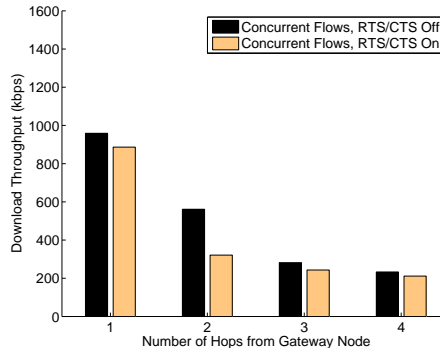


Figure 10: Concurrently active download TCP flows (parking lot) sourced at the gateway node and destined for each mesh node.

4.2.1 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. 10, 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.

4.2.2 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




Figure 11: Concurrently active upload TCP flows (parking lot) sourced at each mesh node and destined for the gateway node.

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. 11, 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.

4.2.3 Effect of RTS/CTS

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. 9), 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.

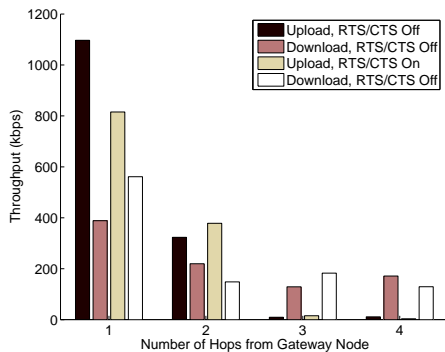


Figure 12: Parking lot scenario in both directions with and without RTS/CTS enabled.

4.2.4 Bidirectional Traffic

In Fig. 12, 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.

4.3 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. 9, 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.

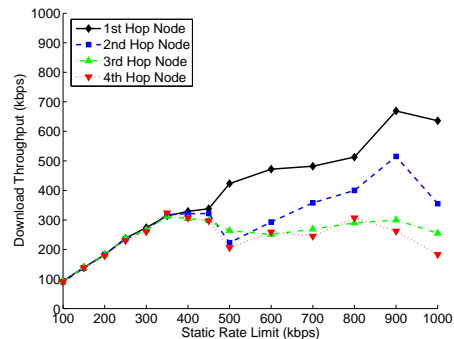


Figure 13: The fully backlogged parking lot traffic matrix downstream with each flow equally rate limited at the source.

4.3.1 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. 13). 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.

4.3.2 Upload Traffic

In Fig. 14, 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 14: The fully backlogged parking lot traffic matrix upstream with each flow equally rate limited at the source.

4.4 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 (e.g., web traffic) to enable greater fairness due to (i) increased transmission gaps that can be 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: the fair share of the bottleneck link and the activity factor of flows (defined in 4.4.2).

4.4.1 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 offered load of each node relative to its fair share of the bottleneck link, $\frac{C}{m}$, where C is the effective wireless link capacity and m 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, $\frac{C}{m}$,

$$\phi = \frac{Nsm}{TC} \quad (2)$$

4.4.2 Activity Factor

We define the activity factor as the mean number of actively contending flows at a given instant in time within a parking lot. We can expect far greater channel utilization if there are no contending flows within the topology. The round trip time, rtt_i for flows being requested at node i , is the time the flow is active on the wireless medium less the delay over the wired medium to negate propagation differences of Internet servers. Denoting 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

$$f = \frac{Nm\tau}{T} \sum_{i=1}^m rtt_i \quad (3)$$

4.4.3 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 [11]. 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 account 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.

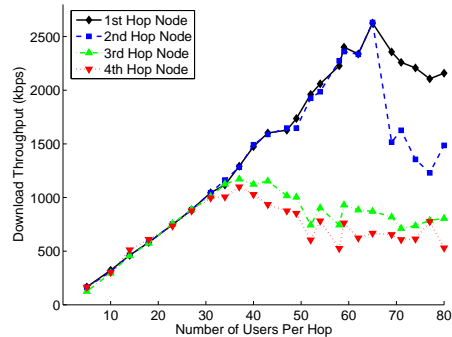


Figure 15: Each mesh node has an equal number of web-emulated users in the same linear topology.

4.4.4 Results

From the traffic load on each of the mesh nodes in our experiment (solving Eqn. 2 for $\phi \geq 1$ with the traffic characteristics of 4.4.3), we expect spatial bias with more than 30 users. Indeed in Fig. 15, 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. 10). From Eqn. (3), we find that on average there is less than one flow contending for the channel in the download direction (upload traffic is less than 1 percent of the total bandwidth in our experiments) for up to 25 users where the total number of round trips are 5 and the rtt_i are 3, 6, 9, and 12 ms respectively. Thus, we attribute the throughput and fairness gain

i	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		Down	Up	Down
	Off	On	Off	On	Off	On	Off	On			
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 1: Empirically Defined Multihop Throughput Distribution, $\vec{\beta}$ (bold values used in Section 5)

to the small web flows (25 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 through smaller idle time, larger download files, or an increase in the number of users, the behavior eventually approaches that of the fully backlogged scenario.

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

We now empirically quantify the falloff of multihop throughput from the aforementioned parking lot experiments which we use to develop a placement model in Section 5. We also compare our measurements with simulations.

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.

4.5.1 Fully Backlogged Parking Lot

Table 1 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. 9). 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 [10], 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 [10] 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.

4.5.2 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. We use these values in our placement study to compare against fully backlogged scenarios without rate limiting and cases where flows are singly active without competing flows.

4.5.3 Web-Emulated Parking Lot

In Table 1, we use $\vec{\beta}$ for web traffic at 30 users, the point at which the Parking Lot Load (Eqn. 2) 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.

5. PLACEMENT STUDY

The measurement results in Sections 3 and 4 provide an empirical basis for our placement study. In this section, we develop a computational model for network reliability and throughput based on our probabilistic link model and explore the impact of several important factors in mesh network deployment. We then investigate mesh node density, wired node density, $\vec{\beta}$ values and random wired entry point location.

5.1 Methodology

For our baseline topology scenario, we consider a square (Manhattan) grid in an infinite plane. We also examine two variants: grid placement with random perturbations and unplanned (random) topologies (Fig. 16). The ratio of wired mesh nodes to wireless mesh nodes is given by w , where $w = 1$ is a wireless access network with a wire connecting every mesh node to the Internet. We study regular wire placements before exploring the effects of random wire placement. Additionally, we do not consider edge effects in our network which arise due to the use of a finite plane in computation. Therefore, we allow the edge nodes to potentially be used in routes to a wired gateway but an edge




Figure 16: Three general topologies considered here are regular grid, regular grid with perturbation, and random placement.

node does not report its performance results. For estimating mesh node throughput, we implicitly assume that there is a uniform traffic demand throughout the coverage area. Our findings in Section 3 indicate that access links are not a limiting factor in backbone deployment because the effective range of an access link is comparable to that of a backhaul link. Thus, for an access node to be disconnected, our regular grid must also be disconnected. Consequently, we do not include the access link in our mesh node placement study.

The signal strength distribution of the link between two backbone nodes is given by our empirical model in Section 3. Using our measured parameters, we model the distribution of the signal strength as a Gaussian random variable with a mean determined by link distance. We consider a link usable if its average signal strength is greater than a certain threshold T_s . We then find an expected link throughput by mapping signal strengths to our measured mean throughputs for each signal level. While we consider links as symmetric, as is nearly true in many scenarios, our performance measures can easily be extended to the asymmetric case by considering upload and download paths separately.

5.2 Reliability

We define *reliability* as the probability that a node has at least one path to a wired node in which each link satisfies an average minimum signal level, T_{min} . The value of T_{min} is dependent on the physical layer technology and the minimum level of service we wish to provide. In the remainder of this work, we use a threshold of -75 dBm which provides an expected throughput of approximately 2 Mbps based on our empirical measurements (see Fig. 7).

To calculate the reliability of service at a mesh node, we first evaluate the probability that a route from mesh node A to mesh node B , R_{AB} , exists (i.e. is usable at a desired performance level) as the probability that each link along the multihop path from A to B exceeds the minimum signal level:

$$Pr[R_{AB} \text{ exists}] = \prod_{\forall i \in R_{AB}} Pr[S_i > T_{min}] \quad (4)$$

where S_i is a random variable representing the signal strength of link i and T_{min} is the minimum acceptable signal strength.

The *reliability* of A is defined as the probability that A is connected to a wired mesh node. Therefore, reliability is the probability of having at least one successful route among the set R_{AW} of all possible routes from A to a wired entry point W as given by:

$$Pr[A \text{ connected}] = Pr[\exists R_i \in \{R_{AW}\} : R_i \text{ exists}] \quad (5)$$

Observe that A may have any number of routes to a wired

node and many routes may share individual links. Also, note that a wired node is connected with probability 1. Table 2 presents a pseudocode description of our algorithm to find average reliability.

```

Let  $\mathcal{M}$  be set of all mesh nodes.
Let  $N$  be the total number of mesh nodes.
Let  $L_r$  be the length of route  $r$ 
Let  $S_k$  be the expected signal strength of link  $k$ 
Let  $T_s$  be minimum acceptable signal strength

Foreach wireless mesh node  $m_i \in \mathcal{M}$ 
  Find  $\mathcal{R}$ , the set of all routes from  $m_i$  to a gateway
   $R(m_i) = Pr[\exists r \in \mathcal{R} \text{ such that } r \text{ is connected}]$ 
    where  $Pr[r \text{ is connected}] = \prod_{k=1}^{L_r} Pr[S_k \geq T_s]$ 
Foreach wired mesh node  $m_j \in \mathcal{M}$ 
   $R(m_j) = 1.0$ 

Avg Reliability =  $\frac{1}{N} \sum_{i=1}^N R(m_i)$ 

```

Table 2: Pseudocode for finding average mesh node reliability.

We first consider a regular grid network and examine the average reliability of the wireless nodes. Our objective is to minimize the network cost by finding the minimal node density at which a threshold of performance is achieved. Fig. 17 depicts average reliability as a function of mesh node density and indicates three regions in the graph. At high node density (low inter-node spacing) the system is not sensitive to small changes in density, whereas at inter-node distances between 200 and 300 meters the average reliability declines quickly for all wire ratios. At distances greater than 300 meters, the system slowly converges to zero reliability as nodes become completely disconnected. The results suggest a desirable operating point at node spacings of approximately 200 meters in our scenario.




Figure 17: Average reliability as a function of the distance between nodes in a regular grid topology.

Another important factor affecting network cost is the density of wires in the network coverage area. Fig. 18 plots the increase in average reliability as the wire ratio increases. We observe a tradeoff between mesh node density and wire density. When the network permits a higher wire density, the node density can be relaxed while still achieving high reliability. But if the availability of wired connections is

tightly constrained, then the mesh node density must be high in order to also achieve high average reliability.




Figure 18: Average reliability as a function of wire density with three different node densities.




Figure 19: Mean and standard deviation of reliability as a function of node density.

Fig. 19 presents the mean and standard deviation of reliability as a function of mesh node spacing. As the node density decreases, the reliability of the mesh nodes in the network varies significantly based on a node's proximity to a wired entry point. This suggests that even as mean reliability declines, the network will still feature regions of higher reliability as well as more disconnected regions.

5.3 Throughput

As our second performance metric, we define *throughput* as the maximum of the expected rates along every path to a wired node. We do not consider multi-path routing in our scenario; throughput is determined by the single best route, which is not necessarily the shortest route.

Using $\vec{\beta}$ introduced in Section 4 (see Table 1), we define the expected achievable rate for a multihop flow across heterogeneous links. We designate ρ_i as the achievable rate for a flow of length i hops and define ρ_i as follows

$$\rho_0 = C_w \quad (6)$$

$$\rho_i = \beta_i \min_{1 \leq j \leq i} (C_j) \quad (7)$$

where C_j is the capacity of link j and C_w is the capacity of a wired mesh node. We assume the wired backhaul connection does not limit system performance and therefore the capacity of the wired mesh node is upper-bounded by the

Let \mathcal{M} be the set of all mesh nodes.
 Let N be the total number of mesh nodes.
 Let C_{acc} be the capacity of the access tier
 Choose $\vec{\beta}$ according to desired traffic matrix

ForEach wireless mesh node $m_i \in \mathcal{M}$
 Find \mathcal{R} , the set of all routes from m_i to a gateway
 ForEach route $r \in \mathcal{R}$
 Let L be length in hops of route r
 Let C_r be smallest capacity link along route r
 Calculate throughput of route r as $T_r = \beta_L C_r$
 Mesh node throughput $T(m_i) = \max(T_r), \forall r \in \mathcal{R}$
 ForEach wired mesh node $m_j \in \mathcal{M}$
 $T(m_j) = C_{acc}$

$$\text{Avg Throughput} = \frac{1}{N} \sum_{i=1}^N T(m_i)$$

Table 3: Pseudocode for finding average mesh node throughput.

access tier capacity. By taking the *min* of the link capacities we find the slowest link and scale it by the appropriate throughput falloff β_i .

We then compute ρ for each wireless node in our topology based on the best path to a wired mesh node (see pseudocode in Table 3). By choosing only the best throughput path, we decouple our throughput metric from reliability. Also, the reliability metric is independent of throughput falloff and multi-flow contention.




Figure 20: Average throughput as a function of distance between mesh nodes in a regular grid.

Fig. 20 presents the average throughput for different wire ratios using $\vec{\beta}$ from the upload parking lot scenario. With this metric, there are two regions unlike the three regions in Fig. 17. The throughput falls off much faster than reliability due to the harsh penalties reflected in $\vec{\beta}$ as mesh nodes begin to connect via more than one wireless hop.

When the wire ratio approaches 1 : 1, the throughput becomes less sensitive to increased node spacing due to fewer long, slow flows. This scenario is similar to a cellular network in which each base station is connected via wired backhaul and there are no longer any multihop flows. In this situation, the capacity of the access tier could potentially become the bottleneck.

Fig. 21 presents the average throughput for measured val-

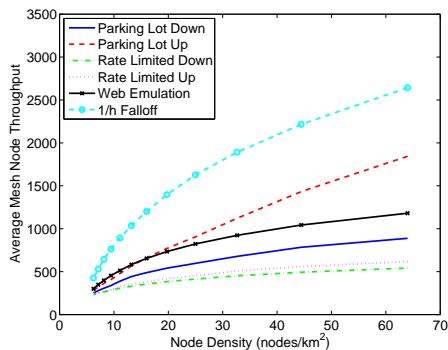


Figure 21: Several empirical and theoretical $\vec{\beta}$ values and their impact on average throughput where $w = \frac{1}{25}$.

ues of $\vec{\beta}$ associated with different traffic matrices (see Table 1). Included is the $\vec{\beta}$ value associated with the commonly accepted single-flow throughput falloff model wherein the throughput of an h -hop flow is proportional to $\frac{1}{h}$. This does not model any contention effects and as a result, the throughput estimates using this choice of $\vec{\beta}$ are at least twice that of our empirical values, which do account for contention. As flows in a deployed mesh network will most likely experience significant contention effects, using the commonly accepted single-flow $\vec{\beta}$ can result in an under provisioned network.

Considering the empirical $\vec{\beta}$ values, the rate limited flows lead to approximately half the average throughput of the non-rate limited flows. The fairness imposed by rate limiting prevents the one-hop flows from capturing a large fraction of the channel time. Statistical multiplexing of the web emulation traffic matrix with 30 users per mesh node allows for fairness and higher performance. These results indicate that given a realistic web traffic workload and rate limiting, a mesh network can achieve both fairness and performance.

5.4 Accurate Link Parameters

In Section 3 we argued for performing physical layer measurements to determine pathloss exponent, shadowing, and achievable throughputs for the target environment and wireless interface. Here we study the sensitivity of mesh network design to these parameters which demonstrates the importance of their accurate measurement. To achieve this, we use our placement model and a particular pathloss parameter to choose a mesh node density which will give us an average mesh node throughput of 1 Mbps.

Table 4 shows the results obtained for a range of pathloss exponents, α , between 2 and 5, the range predicted for urban environments from [3]. We find that even slight deviations from the actual pathloss either severely hinders the connectivity or conversely, overprovisions the network. At pathloss exponents of 2.5 or less, the network is completely disconnected. Even the average of the range, 3.5, produces a network that has 55% more nodes than necessary to achieve the target rate yet has only 7% error in the pathloss exponent. Most notably, at a pathloss of 5, we find that the network would have nearly 10 times the amount of nodes appropriate for the target rate, and thus, would proportionally increase the cost of the network.

$\alpha = 2.0$	Network is disconnected
$\alpha = 2.5$	Network is disconnected
$\alpha = 3.0$	Average throughput is 601Kbps 40 % lower than desired
$\alpha = 3.5$	Network is 1.55 times overprovisioned
$\alpha = 4.0$	Network is 3.39 times overprovisioned
$\alpha = 4.5$	Network is 5.86 times overprovisioned
$\alpha = 5.0$	Network is 9.42 times overprovisioned

Table 4: Mesh nodes required for average throughput of 1Mbps with pathloss exponent (α) estimates. The actual $\alpha = 3.27$.

In Section 3, we observed an 11 dBm shift from the nominal throughput levels for the mesh node wireless interface due to multipath fading effects within our environment. Thus, if we model the expected throughput for a given signal strength from the manufacturer’s datasheet, we find that the network should need only 2.05 nodes per km^2 whereas in actuality, we know that the network needs approximately 20 nodes per km^2 to achieve a target rate of 1 Mbps per mesh node network-wide. Again, deploying a mesh network based on this consideration would lead to a disconnected network. We ran similar tests using commonly accepted estimates for shadowing standard deviation of 8 dB in comparison to the measured values and found this does not lead to significant error in network provisioning.

5.5 Random Perturbations

We next add random perturbations to our regular grid while keeping the wired entry points regular throughout the network. This reflects the realistic scenario of having only partial control over deployment locations. Each node is perturbed from the regular grid by choosing a random angle and radius from a uniform distribution.




Figure 22: Random perturbations from regular grid placement with node spacing of 225 meters.

Fig. 22 plots of the average throughput for increasing node perturbations. Surprisingly, small perturbations lead to slightly improved average throughput, though with higher standard deviation. In this scenario, the maximum throughput increase of approximately 6% occurs with perturbations of one sixth the distance between mesh nodes. The small increase in throughput is due to the routing protocol being able to take advantage of links that improve due to perturbations and avoid links that worsen by using other routes. As we continue to increase the perturbation, coverage gaps

begin to appear and the average throughput decreases and the standard deviation increases.

5.6 Random Placements

We now study the effects of random wire and node placement in the network because the network designer often does not have control of the wired entry point locations or the mesh node locations. For random topologies, we generate a spatial Poisson process of intensity equal to the node density in the regular topologies to which we compare.

We display reliability (Fig. 23) and throughput (Fig. 24) for three scenarios: the baseline regular grid placement, regular grid placement with random wired entry points, and random node and wire placement. We observe at least a 20% improvement in average throughput and reliability with a regular wire placement over a random wire placement in a regular grid.

The reliability of the random node placement scenario improves very little as the number of wires increases, whereas the reliability continues to improve with the random wire placement scenario. The throughput flattens out for the regular placement for the previously stated reason that our metric only captures the wireless backhaul throughput and does not react to additional nodes becoming wired. Finally, note that the region where regular grid placement most outperforms the random topologies is at a wire ratio of approximately 1 : 10 which is a likely operating regime.




Figure 23: Effect of random wire and node placement on reliability with 20 nodes per square kilometer (225 meter spacing).

5.7 Case Study Network

We now examine the specific size and wire placement constraints of our case study network. Based on our target service level, we choose a regular grid with 225 meter spacing.

The dominant constraint of this network is the wire placement. The only existing wired connection is in a corner of the approximately square neighborhood. Topologies with only this wired entry point indicate an average throughput less than 200 kbps. We have deployed two more virtual wired entry points at cooperating institutions, one in the center of the neighborhood and one at the opposite corner, by installing additional radios connected to directional antennas at the wired node. These directional links are on a separate channel from the rest of the network and form point-to-point links that act as additional wired entry points.

In Fig. 25, we compare the average throughput of the case study network’s wired node locations with random wire

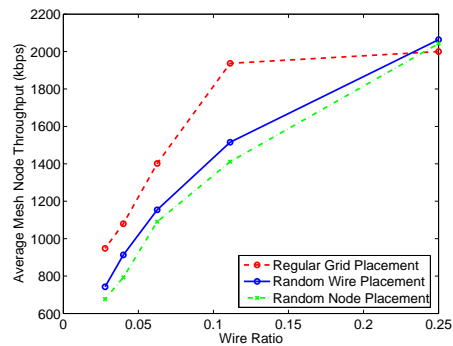


Figure 24: Effect of random wire and node placement on throughput with 20 nodes per square kilometer (225 meter spacing).

placements. We find that our wire placement is 2% better than the median random placement, although still 12% (130 kbps) less than the best placement of three wires. Through careful placement of our second and third wired locations, we overcome the suboptimal placement of the first wire to achieve five times the average throughput of the single-wire case with the addition of two wired entry points. In fact, with only the addition of the second wired gateway near the geographic center of our network, the average throughput increases by a factor of 2.75.

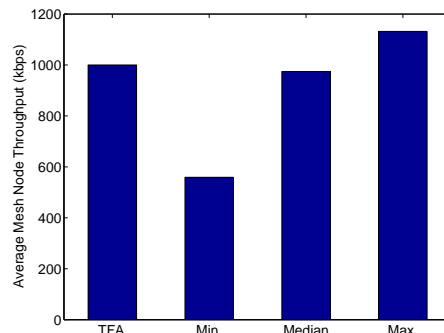


Figure 25: Average throughput in TFA case study network with the constrained wire placement and bounds of random placement. The mesh node spacing is 225 meters.

6. RELATED WORK

Propagation and link-layer studies in the 2.4 GHz ISM band are generally concerned with indoor environments [12] [13] [14]. However, based on measurements at other frequencies in [3] and [4] and the limited literature available on outdoor 2.4 GHz channels [8], 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 [7]. 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 [3]. Thus, it is difficult to employ any existing empirical results without first performing measurements in order to verify their applicability.

There are several existing measurement studies on multihop wireless networks. As in [9], we measure link level performance, but for a significantly different environment (see Section 1). Other works focus on the evaluation of route metrics [15], mobility and route repair [16], and building ad hoc multihop wireless testbeds [17]. We differ from ad hoc multihop wireless in that our infrastructure is static and the traffic matrix is not arbitrary.

Several simulation and analytical works examine capacity [18] and fairness in mesh networks [19, 20, 21]. More specifically, Gambiroza et al. [10] 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 traffic matrix.

Our planned deployment strategy differs from the unplanned topology in [6]. Further, the authors consider only single flows active independently whereas we consider contention-based measurements to evaluate a mesh topology. [22] 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 while we consider the joint problem of placing wireless and wired mesh nodes. Further, their work is purely analytical in nature whereas our model is driven by our measurements. Finally, [23] 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.

7. CONCLUSIONS

In this work, we develop a measurement driven deployment strategy for a two-tier urban mesh access network. With our link measurements, we generate a probabilistic model for link throughput and reliability as a function of distance. We then empirically define the CSMA multihop throughput distribution for competing, multihop flows of different traffic types, including fully backlogged, rate limited, and web-emulated traffic. Using the empirical data gained from our single and multihop measurements, we compute the performance of a broad class of mesh networks. We show that well-known theoretical and practical assumptions about physical environments, throughput models, and traffic matrices each can lead to mesh networks that are either completely disconnected or heavily overprovisioned, increasing cost by an order of magnitude. Further, we find that rate-control mechanisms and web traffic eliminate starvation and yield high performance. Finally, we demonstrate that network performance improves up to 50 percent with respect to both throughput and reliability by careful placement of both wired and wireless nodes.

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.

8. 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] 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.
- [3] T. Rappaport, *Wireless Communications, Principles & Practice*, ser. Emerging Technologies Series, T. Rappaport, Ed. Upper Saddle River, New Jersey: Prentice Hall, 1996.
- [4] G. Stuber, *Principles of Mobile Communication*, 4th ed. Boston: Kluwer, 2000.
- [5] D. Neff, "Wireless Philadelphia business plan: Wireless broadband as the foundation for a digital city," 9 February 2005.
- [6] 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.
- [7] 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.
- [8] 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.
- [9] 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.
- [10] V. Gambiroza, B. Sadeghi, and E. Knightly, "End-to-end performance and fairness in multihop wireless backhaul networks," in *Proceedings of ACM MobiCom*, Philadelphia, PA, September 2004.
- [11] 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.
- [12] 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.
- [13] 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.
- [14] 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.
- [15] 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.
- [16] 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.
- [17] D. Maltz, J. Broch, and D. Johnson, "Quantitative lessons from a full-scale multi-hop wireless ad hoc network testbed," in *Proceedings of the IEEE Wireless Communications and Networking Conference*, September 2000.
- [18] P. Gupta and P. R. Kumar, "The capacity of wireless networks," *IEEE Transactions on Information Theory*, vol. 46, no. 2, Mar. 2000.
- [19] 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.
- [20] 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.
- [21] 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.
- [22] 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.
- [23] B. Liu, Z. Liu, and D. Towsley, "On the capacity of hybrid wireless networks," in *Proceedings of IEEE Infocom*, San Francisco, CA, April 2003.