

A Measurement Study of a Commercial-grade Urban WiFi Mesh*

Vladimir Brik, Shravan Rayanchu, Sharad Saha, Sayandeep Sen, Vivek Shrivastava, Suman Banerjee

Department of Computer Sciences, University of Wisconsin, Madison, WI 53706, USA
{vladimir, shravan, sharad, sdsen, viveks, suman}@cs.wisc.edu

ABSTRACT

We present a measurement study of a large-scale urban WiFi mesh network consisting of more than 250 Mesh Access Points (MAPs), with paying customers that use it for Internet access. Our study, involved collecting multi-modal data, e.g., through continuous gathering of SNMP logs, syslogs, passive traffic capture, and limited active measurements in different parts of the city. Our study is split into four components — planning and deployment of the mesh, success of mesh routing techniques, likely experience of users, and characterization of how the mesh is utilized. During our data collection process that spanned 8 months, the network changed many times due to hardware and software upgrades. Hence to present a consistent view of the network, the core dataset used in this paper comes from a two week excerpt of our dataset. This part of the dataset had more than 1.7 million SNMP log entries (from 224 MAPs) and more than 100 hours of active measurements. The scale of the study allowed us to make many important observations that are critical in planning and using WiFi meshes as an Internet access technology. For example, our study indicates that the last hop 2.4GHz wireless link between the mesh and the client is the major bottleneck in client performance. Further we observe that deploying the mesh access points on utility poles results in performance degradation for indoor clients that receive poor signal from the access points.

Categories and Subject Descriptors

C.2.3 [Network Operations]: Network monitoring, Public networks; C.2.1 [Network Architecture and Design]: Wireless communication

General Terms

Documentation, Experimentation, Measurement, Performance, Reliability

*All student authors — V. Brik, S. Rayanchu, S. Saha, S. Sen, and V. Shrivastava — are in alphabetical order.

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.

IMC'08, October 20–22, 2008, Vouliagmeni, Greece.

Copyright 2008 ACM 978-1-60558-334-1/08/10 ...\$5.00.

Keywords

WiFi Mesh, Commercial, Measurement, Wireless, Access Network, Urban, Active Measurement, Pathloss, Client Experience

1. INTRODUCTION

The widespread availability of WiFi and the benefits of a low cost tetherless network deployment have spurred a significant amount of interest in wireless mesh networks. As the initial excitement of this new possibility has subsided, realities and challenges of making such a network as a viable technology have set in. The research community has been spending many years and person-hours trying to both understand performance issues of multi-hop wireless communication as an access technology and tackle problems in multiple innovative ways. A few examples of recent progress include new channel-aware routing metrics such as ETX [1] and WCETT [2]; new opportunistic and broadcast-oriented routing strategies such as ExOR [3], COPE [4], MORE [5], and CLONE [6]; channel planning and assignment strategies that combine routing decisions [7, 8], mesh deployment strategies [9, 10]. Additionally, a number of vendors (Cisco Systems, Firetide, Mesh Networks (now part of Motorola), Strix Systems, and Tropos Networks, to name a few) have also spent their efforts in creating and releasing commercial-grade mesh networking solutions that are being used for different purposes, including municipality-wide Internet access, public safety, and commercial use.

A number of detailed and insightful measurement studies in the recent years have characterized performance of various moderate scale, primarily home-grown and organically expanding, mesh networks. Examples include pioneering work on the Roofnet testbed around Cambridge, MA [11], the TFA-Mesh in Houston, TX [9], and the Digital Gangetic Plains project in India [12]. Mesh deployments with a research intent, often, have important limitations. In such cases, we typically use off-the-shelf components for cost reasons, and adapt them to their individual goals. Often, in such deployments we provision additional capability for measurements and experimentation to further the research goals. Additionally, the services offered by these deployments are, often, at no cost to the user. Hence, initial expectations of performance are relatively low. In contrast, users of a commercial network have significantly high expectations of network stability and availability. Therefore, commercial networks are, often, carefully deployed, engineered, and tuned for high quality performance. While multiple studies have documented the experiences of mesh networks, that are off-




Figure 1: Logical view of the MadMesh network.

shoot of research endeavors, in this paper, we attempt to present a first systematic study of a commercial-grade WiFi mesh network¹. The mesh network we study has been operational in Madison, WI, for more than two years now and is managed by a local company called Mad City Broadband (see <http://www.madcitybroadband.com>). We refer to this mesh as *MadMesh*. MadMesh consists of more than 250 Mesh Access Points (MAPs) distributed in the greater downtown area of 10 square miles, currently serves more than 1000 residential customers, provides traffic backhauling capabilities for some other ISPs and small businesses, and provides additional wireless services to different public safety organizations of the city.

1.1 Study goals

Through detailed measurement efforts spanning more than 8 months, we wish to answer a broad categories of questions. In each category, apart from trying to understand existing phenomena, we also attempted to evaluate the relevance of various ongoing research efforts to improving performance. For instance, given that a lot of efforts are being spent in designing effective network coding based routing strategies [4, 6, 5], how applicable are they to common deployment scenarios. These categories are:

- *Mesh planning and deployment*: What are the deployment strategies and their efficacies for a large-scale mesh network spanning a substantial part of an urban area? How effective is such deployment in handling failures?
- *Mesh routing strategies*: What are common routing mechanisms adapted? How well do they perform?

¹At the time of publication, we became aware of two concurrent pieces of work that examine certain characteristics of metro-area WiFi mesh networks [22, 23]. As discussed in Section 7, these efforts are complementary to work reported in this paper.

- *User experience*: What is the client performance in different parts of the mesh network?
- *Usage characterization*: How is the mesh utilized by the users?

While it is always dangerous to generalize observations based on one single deployment, we believe that the relative success of this network makes it a reasonable starting point for other future evaluations along these lines.

1.2 MadMesh Architecture and Use

The MadMesh network is comprised primarily of Cisco 1510 MAPs [13]. The MAPs are typically organized into a tree structure, with the root referred to as a *Root Access Point (RAP)*, i.e., a RAP is a MAP selected to serve as a root (see Figure 1). While a RAP typically has wire-line access to the Internet, in the case of MadMesh, the RAPs use special licensed, wireless frequency bands to communicate to an Internet fiber hub. MadMesh has multiple MAPs configured as RAPs, and hence, there is a separate tree corresponding to each RAP. Each MAP is configured to detect other nearby MAPs, and associate with one tree for all communication. Based on changing channel conditions, a MAP can also change its parent in the tree, or even switch to a different tree, if available. A mesh controller directly configures and manages all the MAPs in the network. In particular, each MAP establishes a Layer 2 tunnel to this mesh controller soon after it boots up and joins an existing tree.

We refer to the links between different MAPs on a tree, as the mesh backbone. The Cisco 1510 MAPs are equipped with two radio interfaces. One interface of each MAP is dedicated for communication on the mesh backbone, and is referred to as the *backbone interface*. The second interface is configured to act as a regular Access Point (AP) for regular clients, and is referred to as the *access interface*. (When we refer to the corresponding wireless links, we refer to them as backbone link and access link respectively.) The backbone interface in MadMesh is configured to operate using the 802.11a standards in the 5 GHz band, while the access interface is configured to operate using the 802.11b/g standards in the 2.4 GHz band. Since each MAP dedicates a single radio interface for backbone communication, all these interfaces of MAPs that form the same tree are made to operate on the same 802.11a channel to establish communication links. Different trees operate on different 802.11a channels. The access interface of different MAPs operate on different 802.11 b/g channels. Clients associate to the access interface of a MAP using common WLAN procedures.

In the MadMesh deployment, the access interface is always configured to use an omni-directional antenna with 5 or 8dBi gain to achieve the desired coverage. The MAPs are mounted on street utility poles and the expected coverage of the AP interface is around 1000 to 1500 feet. Most backbone interfaces of MAPs use a 11 dBi sector antenna for more efficient communication. All of the backhaul traffic is encrypted by the MAPs using hardware-based AES to ensure privacy of the users. The APs support the 802.11i and WPA standard security authentication and encryption mechanisms.

How MadMesh is used? MadMesh is used to provide Internet access to users. Users are typically charged a fixed monthly fee which varies with the quality of service (like

bandwidth limits) promised to them. Overall, the end users of this network use it mostly from the student dormitories, university buildings, cafeterias and other residences.

Figure 2 presents a high-level view of approximately one-third of MadMesh. The centers of the circles mark the positions of the MAP locations. The circle size is proportional to the number of users served. The lines indicate the typical connectivity structure within the mesh.

1.3 Main observations

We now highlight some of the most important lessons and observations learnt about a large-scale, commercial-grade mesh network through our measurement study.

Robustness — local does not mean global

Each MAP in the network has good connectivity with its peers. For example, about 60% of the MAPs had a degree greater than 3 on average, while the top 10% of the MAPs had a degree of 6 or higher. However, surprisingly there were multiple cases where a single link failure could partition the network. Thus although the network planning involved local redundancy, it did not automatically translate to global redundancy.

Bottleneck — it is the access link

The performance of the mesh backbone was fairly robust. The link qualities were usually good. Multiple hops on the backbone, going all the way up to 6-8 hops did not significantly hinder user performance. However, the biggest hindrance to performance is the interference in the access link. We believe that there are two reasons for it. First, the MAPs are on utility poles, and most users are indoors (in brick or other buildings). The access link, therefore, has poor signal quality from such indoor locations. Second, the radio interface in client devices (laptops, PDAs, etc.) often tend to operate in low-power modes than the MAPs. Thus, although clients can ‘hear’ MAP beacons, the uplink communication link is, often, particularly bad. Customer premise equipments (like 802.11 repeaters) can potentially help mitigate some of these performance problems.

Routing paths — flapping is prevalent

The trees that define routing paths have a flapping behavior. While many MAPs had fairly stable paths, about 10% of the MAPs had routing flaps in a regular fashion (more than 4 route changes per hour between the same alternative choices). Often these flaps occur due to availability of multiple equally good or equally bad alternatives, and call for dampening mechanisms to be put into place.

Management — client feedback can really help

The usual management tools at the disposal of network administrators rely on SNMP data collection from MAPs and other infrastructure-based components. Unfortunately, the MAPs, often, do not observe the real performance problems being experienced at clients. Inferring client performance based on observations at the MAPs is harder due to the high variability and complexity of the urban WiFi environment. A limited amount of automated client feedback (client reports) can bring many of these performance problems to light.

Applicability of recent research results — network coding and opportunistic routing can help

In the recent past, new, wireless-specific, routing and MAC mechanisms, such as network coding and opportunistic routing, have been proposed and demonstrated to work through research prototypes in limited settings. However, the question of their real applicability in outdoor mesh deployments have not been answered. Our measurements indicate that certain degree of topology diversity exists in the network that will allow for these mechanisms to lead to performance gains.

User characteristics — night-time peaks and uneven usage

Finally, we have also studied the usual aspects of user behavior on this network. Being primarily a residential access network, we observe that traffic volumes peak in late evenings and the night hours, rather than in the daytime. This is likely to be consistent with traffic patterns of other access networks, but is contrary to observations made in core ISPs (that see daytime peaks). Client distribution between MAPs is also quite uneven.

1.4 Roadmap

The rest of the paper is structured as follows. In the next section we describe our measurement methodology. In Sections 3 to 6, we examine different questions in the four broad categories, namely mesh planning and deployment, mesh routing strategies, user experience, and usage characteristics. In Section 7, we present some related work and place our current effort in perspective. We, finally, conclude in Section 8.

2. MEASUREMENT METHODOLOGY

For our measurement study we have collected data over a period greater than 8 months. The main limitation to our ability to collect data has to do with the commercial nature of the network we study. Although most of our needs were accommodated by the network operator, our access to logs was limited and our experiments had to be conducted in a manner that would not significantly impair the network performance. Combining the passive and active data, however, still allowed us to capture and understand the overall characteristics of the network.

Periodic infrastructure logs

Using our privileged access to the mesh controller we polled SNMP records from all active MAPs, once every three minutes. Each SNMP record had more than 150 parameters that each MAP records about its performance. For example, each MAP maintains statistics about the total number of clients associated to it, the MAC address of the current parent MAP, MAC addresses of its neighbors, the current channel number, the number of failed transmissions, the noise floor level at the MAP etc. In addition, we had access to various management tools and syslogs at the mesh controller, that tracked other global mesh parameters.

Passive monitoring

We strategically placed a few monitoring nodes at different parts of the network to gather wireless traffic passively. We used three forms of passive locations — (i) an outdoor utility

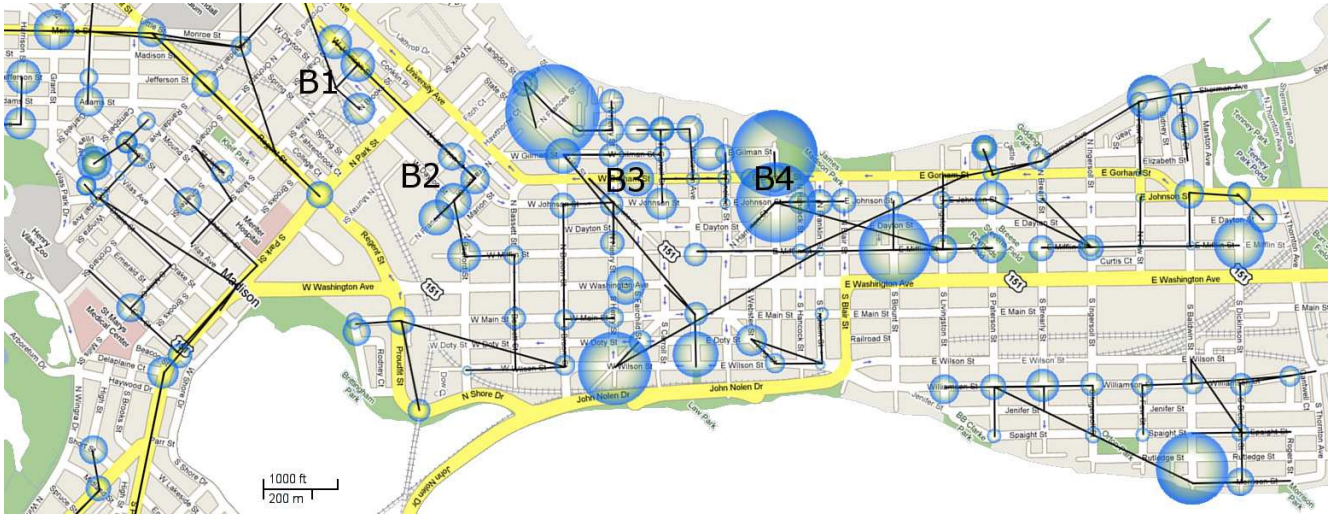


Figure 2: A third of the MadMesh deployment area. The circles represent the MAPs, their relative sizes indicate the relative number of users associated, averaged over three minute intervals over the duration of the study. The largest circle corresponds to an AP which had 6.55 users in average. The lines indicate their connectivity into the tree structure.

pole mounted monitoring node that was close to a MAP, (ii) indoor monitoring nodes co-located with a few residential users, and (iii) a mobile monitoring node mounted on a city-bus traveling all over Madison, WI. Unfortunately, all client traffic on the mesh was encrypted, and hence, it was not feasible for us to do any application-level traffic analysis. However, the MAC-level headers of all wireless frames were available through this method.

Active measurements

Our log analysis revealed that passively collected data did not adequately describe experience of individual network clients. To address this issue, members of our teams periodically went to different parts of the city to perform limited volumes of active measurement, using tools such as iperf [14]. For these measurements we used laptops equipped with a Cisco AIR-PCM352 PC Card wireless adapter and the Mad-Wifi driver v0.9.3.

Over the duration of these 8 months, the network itself changed many times. For example, MAPs were moved between different utility poles, hardware was changed and upgraded, and so on. So to present a consistent view of the network performance, the core dataset used in this paper comes from a two-week period, between the end November and early December 2007. This part of the dataset had more than 1.7 million SNMP log entries (from 224 MAPs) and more than 100 hours of active measurements.

Based on these data sets, we now present our observations in the four different categories in the following four sections.

3. ON MESH PLANNING & DEPLOYMENT

We begin with our observations on various topological properties of MadMesh. Many of these questions arise when the mesh is being deployed or periodically upgraded. More specifically, we focus on the following questions in this section:

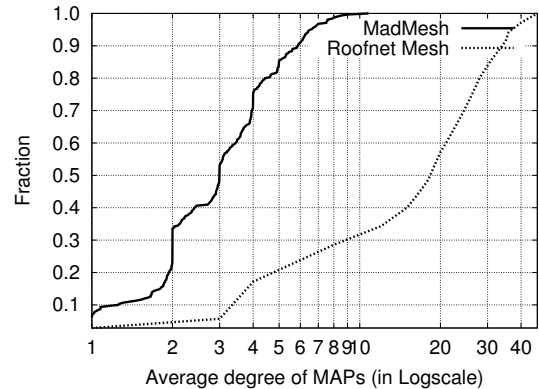


Figure 3: Plot showing the fraction of MAPs with a given average degree for MadMesh and Roofnet

- What does the neighborhood of each MAP look like? What kind of connectivity does each MAP have with its peers?
- How robust is the deployment to failure scenarios?
- What are the link-level error rates and the signal qualities on the backbone and access links? What are their contribution causes?
- Does the network topology lend itself to new routing mechanisms such as network coding and opportunistic routing?

3.1 Average MAP degree

We define the degree of a MAP as the number of neighboring MAPs with link quality above a threshold (14 dB for

this study). The average degree of the MAPs in the network helps us determine the connectivity properties of the deployed mesh network. A low degree of connectivity would imply that the MAPs are constrained in their choices of parent links, which in turn implies limited re-routing choices in presence of losses. Similarly, an extremely high degree of connectivity would imply over-provisioning in the deployment apart from increased possibility of self interference. Figure 3 plots the CDF of the degree for all the MAPs in the mesh. We can observe that 20% of the MAPs have a degree of less than 2 and about 60% of the MAPs have a degree of more than 3. This observation is interesting because more than 70% of the MAPs use directional sector antenna systems with a 45 degree beam for the backbone, and we would have expected a much lower connectedness. Overall, the connectivity of the MAPs within the mesh is fairly good.

However, it is interesting to note that the neighborhood distribution of MadMesh is still much lower than an organically grown mesh, such as the Roofnet. The latter is a network deployed in Cambridge, MA, in and around MIT. Roofnet comprises of a set of wireless nodes that are hosted in homes and apartments of willing volunteers, and hence, there is limited fidelity in controlling its growth and structure. In Figure 3 we also plot the degree of nodes in the Roofnet network, and we believe that the large variation in node density is a consequence of its unplanned growth. In contrast, the deployment of MadMesh is well structured, and was preceded by detailed site surveys. Additionally, positioning of MAPs are continuously changed based on performance requirements. Finally, MAPs in MadMesh are continuously available, and are tightly managed by the network operators, making their uptimes more predictable. Therefore, it is logical to expect that the density of MAPs in MadMesh is much lower than RoofNet.

3.2 Robustness of the deployment

To gauge the quality of the mesh planning one also needs to measure the robustness of the deployment against link failures. A well deployed network should have more than one distinct path to the wired Internet connection. We study the robustness of the mesh topology by looking at the min-cut of each MAP – the minimum number of edges, whose removal would disconnect the MAP from the graph. To understand this, we build a graph out of connectivity reports obtained through SNMP logs and calculate the min-cut of each MAP from the different RAPs. Figure 4, shows the scatter-plot of the average min-cut of all the MAPs as it varied against average degree. As can be seen from the plot, around 8% of the MAPs have a mincut less than 2. This implies that the MAPs would get disconnected from the rest of the network if less than 2 other specific MAPs fail. The figure also shows that MAPs with neighbor degree as high as 7 can still have a min-cut lower than 2. This can be the case, if the neighbors of a MAP have a common ancestor in the path to the RAP. Figure 4 shows a specific instance of this phenomenon where a group of MAPs are connected to the rest of the network via a single path (at a specific time instant). In this case, failure of the common MAP would result in a disconnected topology.

The analysis presented above assumes that all MAPs in the mesh have equal failure probability. However, in reality the failure of a MAP can be triggered by many independent factors including hardware failures, channel fluctuations or




Figure 4: Scatter plot of the mincut and average degree of MAPs

external interference in the medium. Thus the exact failure probability of each MAP may be different depending on aforementioned factors and this information needs to be coupled with the mincut values to provide a better estimate of the robustness of a MAP. One way of estimating the failure probability of a MAP is to track the uptimes of all the MAPs in the system. Unfortunately, we currently do not have access to this information for Madmesh and we hope to provide more details on node failure in our future work. But we still believe that tracking the mincut of the MAPs in the mesh provides a reasonable estimate of the robustness of the mesh and network planners should examine the path diversity in their mesh deployments. More specifically, they need to ensure that multiple paths do exist between each MAP and the different RAPs, that can help tide over individual failures, and a high neighbor degree of each MAP does not automatically guarantee robustness.

3.3 Error rates of backbone and client access links

The mesh network utilizes two different spectral bands (2.4 GHz for access links and 5 GHz for backbone links) for communication. We now characterize the relative performance of access and backbone links. In order to do this, we first compare the packet error rates (PER) for the access and backbone links. We define PER as the fraction of unicast wireless frames for which no corresponding acknowledgment was received. For the calculation of PER on the backbone links, we utilize two SNMP counters which report the number of packets for which the MAP did not receive an acknowledgment (F) and the number of transmitted packets for which an acknowledgment was successfully received (T). We then calculate PER for the backbone links using $F/(T + F)$. However, the SNMP data for calculating the PER on access links was not available to us. We therefore carried out a set of directed active measurements to estimate the PER of the access side. As part of our active measurement experiments, we randomly selected a set of 35 locations in the coverage area, and at each location we connected with the MAP with the strongest signal and performed three sets of TCP *iperf* sessions, each session lasting 120 seconds. Our *iperf* server was running on the mesh controller. We were also capturing packets in the monitor mode (on a different interface), from which we determine the number of retransmissions (and hence the loss rates). We observed a total of 15 distinct MAP's in this active measurement procedure.




Figure 5: Plot showing the CDF of (a) PER of Access (2.4 GHz) and (b) Backbone (5.2 GHz) band.

The PER is averaged for each MAP over all the runs at different locations.

We have plotted our results in Figure 5. The figure shows the CDF of the PER for all the backbone (5.2 GHz) and the access (2.4 GHz) bands. We observe that the PER on the access side is much higher than that on the backbone side. This result indicates that the errors rate in 2.4 GHz access band is much higher than the 5.2 GHz backbone band, which can be attributed to much higher level of interference and noise in the 2.4 GHz band, used widely by most wireless access points to serve end clients. Further, on the access side, the client to MAP link is more vulnerable due to the lower transmit power of the client devices as compared to the MAP. This link asymmetry which can lead to packet losses from client to MAP, even though the signal strength from the MAP to the client is high.

We validated this assumption using a client device with a higher transmit power, and as expected the PER on the access band was significantly lower for that client.

Since client devices typically have lower transmit power, the performance on an end-to-end mesh path will be severely impacted by the interference and errors observed on the access link. Even if the mesh backbone is of high quality (which is the case for MadMesh), the performance observed by clients will be limited by the interference and error effects on its direct connection to the first MAP.

3.4 Channel selection in backbone and access links

As described before, the packet error rates on the access links are significantly higher than those of the backbone links. In order to understand this contrast, we first examine the signal-to-noise ratio (SNR) of the backbone links. For a link operating on a higher SNR, there is a higher probability of successful packet transmissions. Similarly, links with low SNR values can result in high packet error rates. Such low SNR values could be because of presence of high levels of ambient interference in the network or due to a very low received signal strength. In Figure 6(a) we plot the CDF of the SNR values reported for all the backbone links across the duration of study. We observe that nearly 99% of the links have an SNR above 15 dB. These high values of SNR explain the minimal values of PER seen on the backbone links.

We next turn our attention to the quality of the access links. Similar to the backbone links, high SNR values in the access would imply that the clients would experience relatively low losses on the access links. To study the characteristics of the access links ideally we would like to plot the SNR for these links. However, different client radios

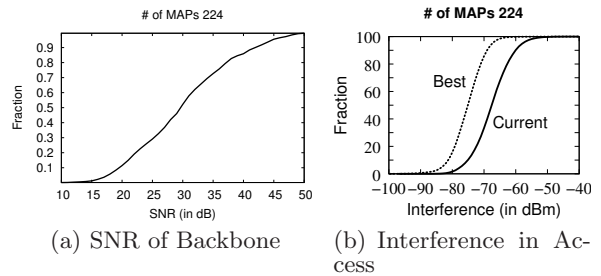


Figure 6: Plot showing the (a) CDF of SNR for all backbone links and (b) CDF of Interference of all MAPs for access, along with the least interference possible (in another channel).

use different radios (quality/brand) and different transmit powers. This makes the reported SNR a function of the client device used and also its proximity to the MAP, which are difficult to quantify. We therefore analyzed the amount of ambient noise (the raw power received from non-802.11 sources operating in the same channel) and the amount of interference (raw power received from 802.11 sources other than the client in communication). These values are reported periodically in the SNMP records for the duration of study. We observed that the ambient noise floor was at an acceptable level of -90 dBm, for nearly 90% of the access links. In contrast, we observed that the amount of interference was very high. Figure 6(b) shows the CDF of the amount of interference for all the access links. As can be seen from the plot, more than 20% of the access links experience a high interference of -70 dBm. We attribute such high values of interference to be one of the main reasons for high PER seen on the access links. Another possible reason could be the low RSS of the client radios, however we could not verify this as we did not have access to the MAP software.

A possible remedy for the losses would involve a) forcing the clients to use a better RSSI while talking with the MAPs and b) ensure that the access link operates in the channel with the least amount of ambient noise and interference. While, the signal strengths from the clients can not be controlled by the MAPs, they can ensure that they operate in the best possible channel (channel with least amount of interference and ambient noise). To find out whether the MAPs are indeed working in the best channel, we plot the CDF of the least amount of interference present in any given channel for all the SNMP snapshots in Figure 6(b). The plot shows the MAPs are not working in the best channel available. We believe that adoption of a channel of operation selection algorithm would result in better performance in the mesh. Designing such channel selection algorithm is nontrivial, since while selecting the locally best channel of operation the MAPs have to ensure that they operate in independent channels to avoid interfering with each other. However given the huge loss rates observed in the current settings, we believe that this optimization would result in improvement of overall mesh performance.

3.5 Feasibility of Network Coding

A good degree of connectivity in the mesh network has an implication in context of current ongoing research in the




Figure 7: Plot showing the maximum coding gain achievable at each of the MAPs.

field of opportunistic routing and coding based approaches for mesh networks [3, 4, 5, 15]. These mechanisms exploit the broadcast nature of the wireless medium and are based on the possibility of overhearing (of data packets) in the network. As shown in Figure 3 nearly 50% of the MAPs have a degree of more than 3. This indicates that there might be a good possibility of overhearing (of data packets) in the mesh network.

In order to understand the achievable gains from overhearing-based mechanisms in the MadMesh deployment, we take the example of COPE [4], a network coding scheme based on opportunistic overhearing of data packets, and calculate an estimate of throughput improvements achievable at each MAP. COPE achieves in-network data compression by XORing multiple data packets together and transmitting a single coded packet. The number of packets which can be coded together at each MAP is determined by the coding rule in [4], which indicates that n packets (destined to n neighboring MAPs) can be coded together only if the structure of the network permits each of the n neighboring MAPs to overhear the other $n - 1$ packets. In such a case, a *coding gain* of n is said to be achieved at this MAP. In order to estimate the throughput improvements possible at each MAP, we derive the maximum coding gain at each of the MAPs using the above coding rule. Note that, the coding gain achievable in practice also depends on the underlying routing mechanism, the number of flows and the direction of these flows. Here, we are interested in finding out the network coding opportunities the deployment inherently supports and we therefore look at the best case scenario i.e., we assume that each of the MAPs always has backlogged traffic to send to each of the other neighboring MAPs.

Figure 7 shows the maximum coding gain for each of the MAPs in the network, if it were to act as the relay node (i.e., the node where the packets were being coded). We observe that around 10% of the MAPs have no coding opportunities available as they were the leaf MAPs (MAPs with only one neighboring MAP). For about 66% of the MAPs in the network the maximum coding gain achievable was only 2 i.e., even though the number of neighboring MAPs for each of these MAPs were more than 1, the structure of the network did not permit coding more than 2 packets. However, there are around 24% of the MAPs where coding gains of more than 2 were possible with the maximum coding gain reaching 6 for some of the MAPs.

This shows that techniques like network coding can potentially improve the performance of such densely deployed outdoor mesh deployments.

4. ON MESH ROUTING STRATEGIES

Routing in multihop wireless mesh networks has been a field of significant research in recent times. Algorithms proposed in [3, 1, 4] describe routing algorithms designed to improve the performance of the network. Studying the functioning of a routing algorithm in a mesh network spanning a city is an exciting problem in its own right. Ideally, such characterization would involve large scale experimentation. However, the commercial nature of the mesh deployment constrains the amount of experimentation feasible. In particular, we could not change the parameters of the routing algorithm to observe its characteristics. Instead, we studied the performance of the routing algorithm in terms of the routing paths created and the relative stability of the routing paths. To reason about the quality and stability (or lack thereof) of the routing paths, one needs to know about the factors which affect the routing decision. In this section, we attempted to answer the following questions:

- How often do routes change and what specific events (from the routing algorithm’s perspective) trigger these changes?
- What is the consequence of the routing algorithms used on the structure of data trees?
- What are potential inefficiencies in the routing mechanisms?

4.1 Understanding behavior of mesh routing algorithm

In this section we present a study of the mesh routing decision algorithm. Our goal is to correlate each possible routing changes with its root cause. The current mesh deployment uses *ease* metric for route creation. Details of this metric are presented in [17]. The *ease* is based on a weighted sum of the SNR and hop count of the potential MAPs. The MAP chooses a neighbor which has the best value for the metric. On comparing this metric with ETX [1], we find that ETX uses (expected transmission count) over a link as an indicator of the quality of the link, in contrast the current metric uses the SNR value as a predictor of the same. Both of them sum the metric over the entire path.

On studying the SNMP logs to identify root cause of a route change we found that both hop count and SNR changes were involved in 0.9 of the entire parent changes ($1 - 0.1 = 0.9$). For the rest 0.1 of the cases both SNR and hopcount worsened due to the parent change. On closer inspection we found that in another .06 of the (total) cases the one of the ancestors increased its hop in the routing tree which made it a bad parent option and hence forced a routing tree change. We could not account for the rest 0.04% by looking at the logs.

4.2 Implication of the mesh routing metrics

A routing metric which is a weighted sum of the link SNRs and hop count, has some non-obvious implications on the final routing paths being used by the MAPs. We comment on them below.

Implication on hop count of the MAPs

The SNMP data contains periodic updates about the hop count of each MAP. We utilize this information to plot Figure 8 which depicts the distribution of average number of MAPs on different hops in the network. We observe that around 15% of the MAPs in the network are RAPs. Also, the average number of MAPs decrease with increase in the hop count i.e., a higher number of MAPs are present at the lower hops. Thus, the network is well deployed and the routing algorithm performs well for most of the time.

However, we also observe that around 8% of the MAPs have a hop count of more than 5. Conventional wisdom suggests that the achievable throughput of the in a multi-hop network degrades drastically with increasing hop count. Presence of longer paths between the MAP and a RAP might




Figure 8: Plot showing fraction of MAPs at various hops, and the number of users at each hop

be because of two reasons: (a) the network was not planned well and therefore sufficient routing choices were not available (b) the routing algorithm is not choosing the right path.

To investigate, whether this high hop-count is a result of bad deployment, we looked into the the best available neighbors in terms of hop-count (i.e., neighbor MAPs with lowest hop-counts) for MAPs with a hop-count greater than 4. We have plotted the CDF of the hopcounts of such neighbors with the best available hops and reasonable link quality (link SNR higher than 14 dB) in Figure 9. As can be seen from the plot, the MAPs at hop-counts higher than 4 always have a neighbor with better hopcount available. This leads to the conclusion that the phenomenon is not an artifact of the deployment i.e., there were other (shorter) paths available in the network, but the routing algorithm did not choose to use it. This behavior is an implication of using a routing metric which is a weighted sum of SNR and hop-count. A neighbor with which the MAP has a better link (in terms of SNR) is given priority as potential parent over another MAP with lower hopcount and a relatively lower SNR. Such long paths can be avoided if the route selection algorithm uses a threshold on SNR for selecting the potential parents and then decides amongst the potential parents based on hop-count.




Figure 9: Plot showing the fraction of current and best hop-counts of MAPs with a hop-count greater than 4

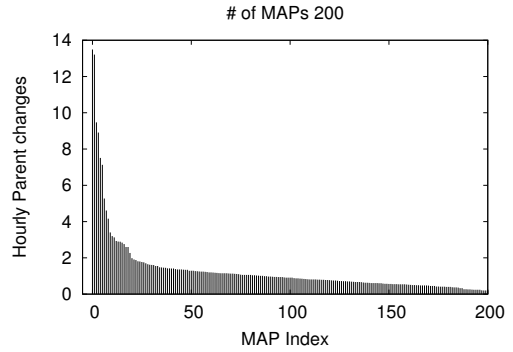


Figure 10: Plot showing the average hourly rate of parent changes for all the MAPs (Note: Rest of the MAPS did not show any parent change)

Implication on route stability

The MAPs in the network can form links to more than one MAP in their neighborhood. A parameter of interest to gauge the performance of the routing algorithm is the relative stability of the routing paths. Large fluctuations in the path imply that the network conditions are unstable. These might be caused by the appearance of another potential link with better characteristics, disappearance of the current parent link or occurrence of certain other events such as increased interference etc., which make a previously unattractive parent link more attractive.

We show the stability characteristics of the mesh routing algorithm in Figure 10. This figure shows the relative frequency of the parent changes for all the MAPs occurring in a given duration. The data was collected by analyzing a parent change counter present in the SNMP log for each MAP over the span of passive data collection. This plot shows that some MAPs have a faster rate of changing their parent than other MAPs. We observed that the high parent change frequency for some of the MAPs was due to: (a) presence of multiple parent choices with similar quality (in terms of SNR and hop count) which coupled with momen-

tary fluctuations in the wireless characteristics makes one of the links momentarily better than the rest forcing the MAP to flap its route. (b) the link to the parent for some intermediate MAP (ancestor) has very bad SNR (less than 10 dB), this causes the intermediate MAP to choose a different parent frequently, thus causing a route flap. A possible remedy for reducing the amount of route flapping would be to keep a threshold on the number of times a MAP can change its route in a given quantum.

5. ON MESH USER EXPERIENCE

Robust client performance is important in commercial, pay per use mesh network. Although careful evaluation of SNMP logs provide us with valuable insight into the dynamics of mesh infrastructure, client performance can be better understood by carrying out targeted active measurements at different locations in the mesh. Specifically, we want to know the following regarding client performance in commercial mesh deployment under study:

- How good is the quality of client to mesh connectivity in MadMesh ? Are coverage holes prevalent ? What is the impact of client mobility on coverage holes?
- What is the maximum achievable throughput by a mesh client ? What is the impact of hop-count, RSSI, channel congestion on the client throughput?
- Is there any issue of starvation at higher hops when clients are present at lower hops as well ?

Broadly, the aforementioned questions relate to two main issues - how easily can a client connect to the network and once connected, what is the observed performance. We first describe our measurements for characterizing client connectivity in the mesh deployment, followed by a detailed analysis of client performance.

5.1 Client connectivity

Ubiquitous client connectivity is one of the most important goals of large scale mesh deployments. In order to maintain client connectivity, it is important to have a monitoring infrastructure in place which can identify ‘coverage holes’ created due to obstacles, weather and temporary interference sources. Once such coverage holes are detected, corrective action can be taken by adjusting power levels of different MAPs or by deploying new ones. Existing approaches commonly used by leading vendors, employ pathloss models to estimate the expected area of coverage. Such models describe the attenuation experienced by wireless signal as a function of distance. In order to assess the efficacy of such pathloss models, we first perform detailed experiments to characterize the pathloss exponent in our urban environment.

Characterizing pathloss exponent

In the following equation, α is the pathloss exponent, and ϵ is the shadowing component that describes the variation in pathloss exponent. $P_{dBm}(d)$ is the signal strength measured at a given distance d , while $P_{dBm}(d_0)$ is the signal strength at the reference distance d_0 [16].

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

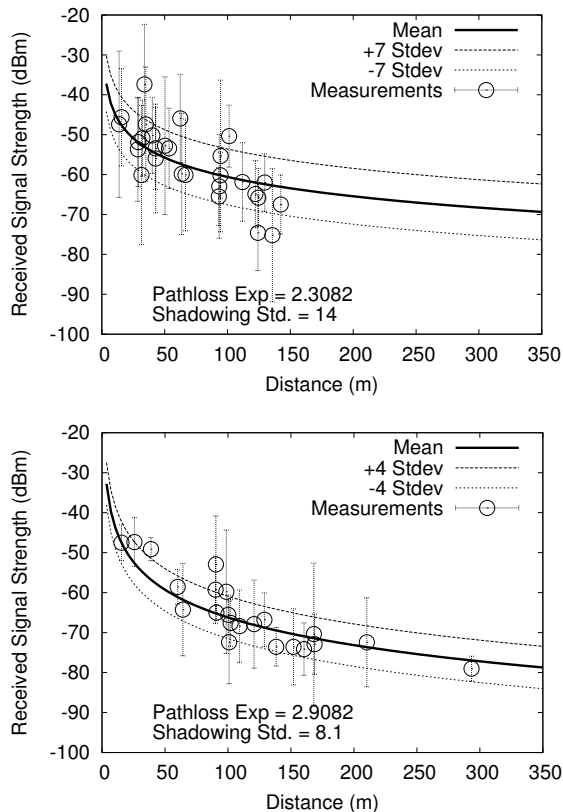


Figure 11: Propagation pathloss for two different MAPs in diverse settings. The MAP in (a) is located in campus and has a pathloss of 2.3, while (b) is located in downtown and its measured pathloss is 2.9.

We follow the measurement methodology reported in prior research work on propagation modeling [9] and collect signal strength information at 25 different locations for each MAP. We compute the pathloss exponent for different MAPs in MadMesh network. Our results indicate significant variations in pathloss exponent between different MAPs. Figure 11(a) and (b) show the signal strength measurements as a function of link distance for two MAPs, located in downtown and campus regions of the city. As shown in the figure, the pathloss exponent for the downtown MAP is 2.9 while the pathloss exponent for the campus MAP is 2.3. This is in contrast to the pathloss of 3.3 shown in [9], who also report that their pathloss exponent is stable across different access points in their network. We attribute this significant variation in pathloss exponent to diverse set of obstacles and external interference, which also vary significantly from one location in the city to another. Our results show that generalizing a pathloss exponent for a city wide mesh deployment may be inaccurate, and targeted experiments must be performed to determine the pathloss in different parts of the city. This observation further reinforces the inefficacy of pathloss models in determining coverage holes. Next we describe a simple monitoring tool that can detect such coverage holes efficiently.

Characterizing coverage holes

In order to assess the prevalence of coverage holes in the mesh deployment under study, we perform extensive client measurements. We report on some sample results in a 6×6 block area of the mesh deployment. In our experiments, a few clients (IBM laptops with Cisco Aironet wireless card) were equipped with a module which continuously records the information about the location, current state of association and received signal strength. Periodically, the clients upload this information to a central server. Over a period of time (seven days, in our case) information aggregated from these clients is used to detect coverage holes in the network. Figure 12 shows the average client connectivity in our target area. Although, the propagation model based radio map generated by the mesh controller shows this entire area to be ‘covered’, we found additional coverage holes were observed by the clients.

Vehicular client connectivity

Wireless access from mobile devices has been an active area of research recently [18]. In that context, we wanted to evaluate MadMesh in terms of providing client connectivity from moving vehicles. Towards this end, we repeat our measurements from moving vehicles that makes round of the same 6×6 block area that we targeted for our walking experiments. The average speed of the vehicle was 25 miles/hr. The coverage holes detected at such vehicular speeds is shown in Figure 12. As shown in the figure, the holes detected by the clients at vehicular speeds are much larger then the holes detected during earlier client measurements. In fact, we observe that about 65 % of the total path falls under the category of coverage holes at vehicular speeds. It is important to note that the observations would have been very difficult to make without the help of actual measurements on the client side. We believe that such measurements can provide significant corrective feedback to the operator regarding coverage holes, which are much more accurate then the propagation models used in current mesh controllers.

5.2 Client Performance

In order to assess the performance of end users in the mesh deployment, we undertake targeted active measurements, where we randomly sample 100 locations in the mesh coverage area and perform bandwidth tests to determine the achievable throughput at that location. At each sampled location, we associate to the MAP with strongest signal strength and run TCP *iperf*[14] from the client to the mesh controller. We use TCP as it is the dominant traffic type in mesh networks, and secondly it is less intrusive then a UDP test, which can completely saturate the link and negatively impact other client in the mesh. We perform three iterations of 100 seconds each. Figure 15 shows the distribution of TCP throughput at the sampled locations. As shown in the figure, the measured throughput closely matches a uniform distribution, with about 10% of the clients achieving less then 0.2 Mbps and 80% of the client achieve throughput less than 1 Mbps. This upper limit on client throughput is expected in view of the Service Level Agreement(SLA) of MadMesh, which advertises a 1 Mbps service to the clients.

Further, to understand the impact of hop count, channel congestion and RSSI on client’s throughput, we perform targeted experiments on one stable 6 hop tree (shown in Figure 13), comprising of eight MAPs and a RAP. MAPs one, two




Figure 13: Tree for our targeted experiments to understand the impact of RSSI, hop count, channel congestion on client performance.

and three shown in Figure 13 are located on a busy main road of the city that has substantial interference from other wireless hotspots in the area. On the other hand, MAPs four to eight are located inside the neighborhood areas, experiencing relatively less interference. We choose a minimally loaded tree, so that our experiments are not impacted by the presence of other MadMesh users on the same tree. Our results from the active measurements on chosen tree is summarized in Table 1. The main observations are as follows:

- In all experiments, maximum client throughput is limited to 1 Mbps, which indicates that bandwidth shaping may be performed by the mesh operator for meeting the SLA.
- Client throughput remains stable with RSSI to a point, beyond which it drops quickly. Since per client throughput is limited by the operator, higher RSSI, which can sustain higher data rate, does not improve client throughput.
- External interference from other wireless sources has a significant impact on the client throughput beyond the first hop.
- Throughput unfairness is observed when clients at different hops of the same tree are activated simultaneously, with clients at higher hops achieving low throughput share.

We now describe each observation in detail.

Impact of hop count

As shown in Table 1, there is no strong correlation between the throughput and the hop count. Different throughput is seen at clients associated to MAPs at the same hop count in the tree. As shown in table 1, throughput of 0.3Mbps and 0.92Mbps is observed on two different MAPs at a hop count of three. Similarly, a throughput of 0.91Mbps and 0.6 Mbps is observed on MAPs at a hop count of five. This variation in throughput at different MAPs with same hop count can be attributed to channel congestion on their access sides which we discuss next.

Impact of Channel Congestion

Presence of traffic due to other 802.11b/g sources can have a great impact on the throughput observed at each hop. We




Figure 12: Actual network coverage as observed by clients, in areas estimated to be perfectly covered by infrastructure-side management tools, that rely on propagation models.

estimate the channel congestion by monitoring the traffic on each hop while doing the TCP iperf experiments. Table 1 shows that the throughput achieved on MAPs at the same hops is well correlated with the channel congestion at their respective access sides. We further observe that channel congestion does not have any impact on client connected to the first hop. This is because of the lesser number of links these clients have to contend for on the backbone. However if the channel congestion is relatively high, this observation might not hold true.

Impact of Shared Congestion

In another set of experiments we study the impact on throughput of clients at a lower hop count in the presence of other clients in the tree. We first associate only one client to a MAP at a higher hop count in the tree and calculate its TCP throughput using iperf as shown in figure 14(a). In order to see the effect on throughput due to other clients in the tree, we associate another client at a lower hop count on the same tree and start running TCP iperf, shown in figure 14(b). As shown in the figure, on running the second client the throughput of the first client suddenly drops from 0.98 Mbps to 0.43 Mbps. This is due to shared channel congestion as discussed in [19], when multiple clients try to contend for the same backbone path. This can have a great impact on the clients connected to MAPs at a higher hop count, which can suffer from increased throughput degradation with the increase in clients at lower hops (closer to RAP).

Summary: In Madcity mesh network, because of bandwidth shaping policies enforced, hop count did not really seem to be the bottleneck for performance. However, this is mostly true in absence of shared congestion; that is in presence of multiple flows sharing the same backbone path, the throughput of higher hop-count routes would be lower. Hence although the penalty of using higher hop counts is diminished due to bandwidth shaping, choosing a lower hop path is still better due to the possibility of shared congestion in the path.

6. ON MESH USAGE CHARACTERIZATION

We now answer one of the most basic questions about the mesh network – *how is the network being used?* Specifically, we want to know the following:




Figure 14: Effect of shared congestion on a client associated to a MAP at hop 4. In isolation it achieves close to 1 Mbps, but when another client at hop 3 is activated, its throughput drops to 0.43 Mbps.

- How many clients are using the network? How does their number vary across time?
- How are the clients distributed across the coverage area?
- What is average number of clients connected to each MAP? Are there any popular MAPs?
- How does client distribution vary across different hops?

Client distribution across time

Figure 16 shows the average number of clients per hour connected to the network over this 2 week period. The error bars show the 95% confidence limits. We observe that the average number of clients varied considerably across the duration of the day, with most number of clients being connected at around 10 PM and the least number of clients at 5 AM. We note that the observed usage pattern is unique to this mesh network as it is mostly accessed by the users from their residences. This is apparent from the fact that

MAP Index	Hop Count	Avg. RSSI. (MAP to client)	Avg. Chnl. Util.	TCP Thrpt.	TCP loss rate (Mbps)	TCP RTT (msec)
1	1	34	0.28	.96	0.021	111±98.4
2	2	33	0.27	0.4	0.092	158±115
3	3	35	0.20	0.3	0.087	258.2±168.7
4	3	40	0.09	.92	0.007	192.5±91.9
5	4	33	0.10	0.7	0.021	252.2 ± 126.4
6	5	32	0.05	.91	0.007	215 ± 73
7	6	37	0.09	0.5	0.030	208.2 ± 117.5
8	5	33	0.11	0.6	0.015	278.7± 127.4

Table 1: Experimental results for the tree under study. Confidence intervals for RSSI and throughput is small and omitted for brevity




Figure 15: CDF of throughput measured through active measurements at 100 random locations in the coverage area of MadMesh. Throughput is almost uniformly distributed, with maximum throughput being clipped around 1.2 Mbps

number of clients starts increasing from around 6 AM, remains steady throughout the afternoon and then again increases from around 6 PM as the users start returning to their homes. It reaches its peak around 10 PM when most of users are their homes and starts tailing off as the night progresses. We observe that during the busiest hour around 627 clients were connected to the network with around 498 being connected to the network on average.

Client distribution across the MAPs

The average number of clients connected to a MAP gives a measure of the amount of load experienced by that MAP. The client distribution across the MAPs also helps us identify ‘client hotspots’ and accordingly deploy more MAPs in that region to evenly distribute the load across the access points. In Figure 17 we plot the number of clients connected to each MAP averaged over the period of study. The MAPs are sorted in the decreasing order of the average number of clients connected to them. We observe that certain MAPs are much more popular when compared to the others, with the average varying from around 7 to less than 1. The figure also shows the 25, 50, 75 and 95 percentiles of the number of the clients and the corresponding number of MAPs to which these clients are connected. For example, one can see that around 50% of the clients are connected to 40 most popu-




Figure 16: Distribution of the average number of clients connected to the network across time.

lar MAPs which account for only 20% of total the deployed MAPs. We also find that many of the MAPs are lightly loaded in the mesh network with around 110 MAPs having on an average less than one user connected to them.

Spatial distribution of clients

In Figure 2, we show the spatial distribution of the clients using the mesh network where each MAP is represented by a circle. The size of the circle represents the average number of clients connected to the MAP. We can clearly observe the uneven distribution of the clients across the coverage area. More importantly, we note that the most of the popular MAPs (the MAPs with higher number of clients connected to them on an average) are concentrated in the area depicted which is a popular area near the downtown. Further, in this region we can observe the formation of a small number of clusters in the areas depicted by B2, B3 and B4. We note that there are several student dormitories in the area B2 which can explain its popularity. While B3 is very popular among the people with a high number of coffee shops and restaurants concentrated in that area, B4 is popular because of an open park where student activity is prominent.

Hop count and number of clients

The distribution of the number of clients across the different hops of a network informs us about how good the deployment is. In a well planned deployment, one can expect to see most of the clients connected to network to be within a few hops. If there are popular MAPs at a higher hop count, network planners might deploy a RAP in the area in order

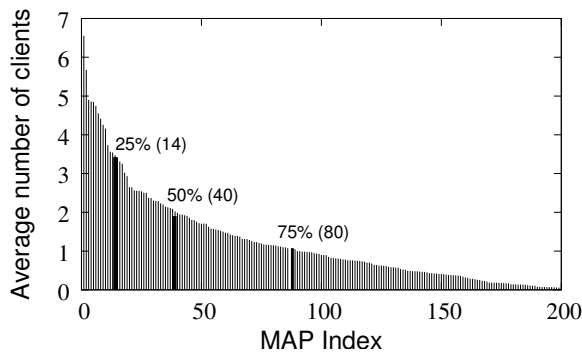


Figure 17: Distribution of the average number of clients connected to the network across different MAPs.

to reduce the hop count. Figure 9 shows the distribution of the clients across different hops of the deployed mesh network. We observe that around 15% of the clients are directly connected to the RAPs and amongst the other MAPs, the number of clients decrease with increase in the hopcount. In particular, we observe that around 85% of the clients are connected to the network within 3 hops. We note that this distribution is similar to the distribution of MAPs (Figure 8) shown earlier because a large fraction of the deployed MAPs are also at a lower hop count.

Distribution of MAPs with high packet losses

We observed that some of the MAPs had experienced a packet error rate of more than 35%. On further investigation, we observed that most of these MAPs had a very low number of clients connected to them. We also observed that the MAPs with similar packet error rates are clustered together representing the regions of high and low interference on the access links. For example, we observe that the MAPs in the area B4 experience very low packet losses as they are deployed in open park. On the other hand, the MAPs with higher packet error rates are clustered in near B3 where there are other WiFi networks causing interference.

7. RELATED WORK

In this section we present a summary of previous work being done on the study of wireless network deployments.

A substantial body of research has already been conducted on evaluation of the performance characteristics of wireless networks. Studies reported in [20, 21] utilize SNMP traces to understand the performance of wireless networks. Specifically, Kotz et. al. [21, 20] present a comprehensive study of the usage patterns (application popularity, temporal variation in utilization etc.), of a campus wide WLAN network. Aguayo et. al report their findings on the link level characteristics of an 802.11b rooftop based mesh network in [11]. The network is deployed in a urban city. Their study focuses on the link level characteristics of the deployment. In contrast, we present results on the quality of the deployment and the application level performance of our network along with link level characteristics of the network.

Chebrolu et. al. [12] and Sheth et. al [9] also study the link level characteristics of outdoor mesh networks, however

their work is applicable to rural settings. Our study was done on a commercial mesh while all of the above mentioned studies were conducted on custom testbeds built explicitly for experimentation.

The work by Knightly et. al reports a measurement study of a mesh network deployment in [9] and highlights the importance of measurements in accurately planning and provisioning mesh networks. While their deployment is has a two-tier architecture as well, their deployment operates exclusively in 2.4 GHz settings while ours operates in both 2.4 GHz and 5 GHz. Also, the span of MadMesh network (250 nodes) is far bigger than their deployment (18 nodes). We summarize and contrast the our measurement study with prior work on mesh network deployments in Table 2. As can be seen from the table, the unique features of our study are, a) our deployment has a far bigger scale in terms of nodes deployed b) use of two type of RF bands for network operation (802.11 a & b), and c) the commercial nature of the MadMesh deployment.

At the time of publication, we became aware of two independent pieces of work performed concurrently with ours, that evaluated different aspects of a metropolitan-area mesh network whose scale was similar to this study. The first of them is work by Knightly et. al [22] that estimated the coverage properties of the Google WiFi mesh network deployed in and around downtown Mountain View, CA. The other is work by Afanasyev et. al [23] that observes the usage characteristics for different user device classes (such as smartphones, stationary modems, etc.) in the same Google WiFi mesh, focusing on application workloads, mobility patterns, and device popularity at different locations. Both these efforts complement our measurement study, and together help provide a greater understanding of different aspects of a metro-area WiFi mesh.

8. CONCLUSION

This paper presents the first systematic study of a commercial grade wireless mesh network deployed in an urban setting. We find that the planned part of the network (backbone) is performing far better than the access side. This disparity in performance is mostly a result of unmitigated interference in 2.4 GHz spectrum in urban settings. The study also presents a set of interesting statistics on the actual usage of the mesh network, which would help in customization of future deployments to make them more profitable. We also present a set of lessons which if followed would result in more robust deployments and stabler routing algorithms in future.

Finally, the study throws open a set of immediate next steps that need to be carefully addressed in real deployments to make mesh networks viable. Some examples include (i) a better architectural design to mitigate interference on the client access link, (ii) design of mechanisms to detect topology robustness in a global sense, (iii) strategies to mitigate route flapping, as common metrics that determine routing changes frequently, and (iv) utilization of client feedback in management of these networks.

9. ACKNOWLEDGMENTS

We thank Patrick Luell, Jordan Walker and Su Yong Lee for their help in performing outdoor experiments. We would like thank Seng-How Liaw, Todd Anderson, Derek Lipp,

	No. of Nodes	Coverage (Sq. Km)	Typical link length	Intended usage	Architecture	Env.	Hw/Sw technology
Roofnet [11] (Rooftop mesh)	~50	4	Mostly < 500 m	Testbed	Omni antennas on rooftops	Urban	Single Tier 802.11b
DGP [12] (Long-distance Mesh)	17	80 km (point-to- to-point)	Up to few tens of kms	Testbed	High gain dirnl. antennas on tall towers	Rural	Single Tier 802.11b
TFA@Rice [9] (Sub-Urban Mesh)	18	3	Mostly < 500 m	Non commercial	Omni antennas on poles	Urban	Two tier 802.11b
MadMesh (Commercial Mesh)	250	26	Mostly < 500 m	Commercial	Omni & dirnl. antennas on poles	Urban	Two Tier 802.11a 802.11b

Table 2: Table comparing this study with prior published measurement studies reported in literature.

Laura Gottfried and Shana Valenta from Mad City Broadband for facilitating this work in multiple ways. We would also like to acknowledge our shepherd, Dina Papagiannaki, whose guidance bought this paper together into its final form. All authors were supported in part by the US National Science Foundation through awards CNS-0639434, CNS-0627589, CNS-0627102, CNS-0520152, and CNS-0747177.

10. REFERENCES

- [1] D. De Couto, D. Aguayo, J. Bicket, and R. Morris. A high-throughput path metric for multi-hop wireless routing. In *MOBICOM*, 2003.
- [2] Liang Ma and Mieso K. Denko. A routing metric for load-balancing in wireless mesh networks. volume 2, pages 409–414, Los Alamitos, CA, USA, 2007. IEEE Computer Society.
- [3] Sanjit Biswas and Robert Morris. Exor: opportunistic multi-hop routing for wireless networks. In *ACM Sigcomm*, 2005.
- [4] S. Katti, H. Rahul, W. Hu, D. Katabi, M. Medard, and J. Crowcroft. Xors in the air: Practical wireless network coding. In *ACM SIGCOMM*, 2006.
- [5] Szymon Chachulski, Michael Jennings, S. Katti, and D. Katabi. Trading structure for randomness in wireless opportunistic routing. In *ACM SIGCOMM*, 2007.
- [6] Shравan Rayanchu, Sayandeep Sen, Jianming Wu, Sudipta Sengupta, and Suman Banerjee. Loss-aware network coding for unicast wireless sessions: Design, implementation, and performance evaluation. In *ACM SIGMETRICS*, 2008.
- [7] M. Alicherry, R. Bhatia, and L. Li. Joint channel assignment and routing for throughput optimization in multi-radio wireless mesh networks. In *ACM MobiCom*, 2005.
- [8] A. Mishra, E. Rozner, S. Banerjee, and W. Arbaugh. Exploiting partially overlapping channels in wireless networks: Turning a peril into an advantage. In *ACM/USENIX IMC*, 2005.
- [9] J. Camp, J. Robinson, C. Steger, and E. Knightly. Measurement driven deployment of a two-tier urban mesh access network. In *ACM MobiSys*, 2006.
- [10] J. Robinson and E. W. Knightly. A performance study of deployment factors in wireless mesh networks. In *IEEE INFOCOM*, 2007.
- [11] Daniel Aguayo, John Bicket, Sanjit Biswas, Glenn Judd, and Robert Morris. Link-level measurements from an 802.11b mesh network. In *SIGCOMM '04*. ACM, 2004.
- [12] Kameswari Chebrolu, Bhaskaran Raman, and Sayandeep Sen. Long-distance 802.11b links: Performance measurements and experience. In *ACM Mobicom*, 2006.
- [13] Product catalog: Cisco aironet 1500 series. <http://www.cisco.com/en/US/products/ps7156/index.html>.
- [14] Iperf. www.dast.nlanr.net/Projects/Iperf/.
- [15] Saumitra Das, Yunnan Wu, Ranveer Chandra, and Y. Charlie Hu. Context-based routing: techniques, applications and experience. In *NSDI'08: Proceedings of the 5th USENIX Symposium on Networked Systems Design and Implementation*, pages 379–392, Berkeley, CA, USA, 2008.
- [16] G. Stuber. Principles of Mobile Communication. 2000.
- [17] Cisco aironet 1500 series wireless mesh ap v5 design guide. <http://www.cisco.com/en/US/docs/wireless/technology/mesh/design/guide/MeshAP.pdf>.
- [18] Ratul Mahajan, John Zahorjan, and Brian Zill. Understanding wifi-based connectivity from moving vehicles. In *Proceedings of the 7th ACM SIGCOMM conference on Internet measurement*, 2007. ACM.
- [19] J. Camp, V. VMancuso, Gurewitz O., and E. W. Knightly. Measurement and modeling of the origins of starvation in congestion controlled mesh networks. In *IEEE INFOCOM*, 2008.
- [20] Tristan Henderson, David Kotz, and Ilya Ayzov. The changing usage of a mature campus-wide wireless network. In *MobiCom '04*, pages 187–201, New York, NY, USA, 2004. ACM.
- [21] David Kotz and Kobby Essien. Analysis of a campus-wide wireless network. In *MobiCom '02*, pages 107–118, New York, NY, USA, 2002. ACM.
- [22] Joshua Robinson, Ram Swaminathan, and Edward Knightly. Assessment of urban-scale wireless networks with a small number of measurements. In *ACM Mobicom*, 2008.
- [23] Mikhail Afanasyev, Tsuwei Chen, Geoffrey M. Voelker, and Alex C. Snoeren. Analysis of a mixed-use urban wifi network: When metropolitan becomes neapolitan. In *ACM/USENIX IMC*, 2008.