An Accurate Technique for Measuring the Wireless Side of Wireless Networks

Jihwang Yeo[†], Moustafa Youssef[†], Tristan Henderson[‡], Ashok Agrawala[†]
[†] *Department of Computer Science, University of Maryland, College Park, MD 20742*
{*jyeo, moustafa, agrawala*}*@cs.umd.edu*
[‡] *Department of Computer Science, Dartmouth College, Hanover, NH 03755*
tristan@cs.dartmouth.edu

Abstract

Wireless monitoring (WM) is a passive approach for capturing wireless-side traffic with rich MAC/PHY layer information. WM can suffer, however, from low capture performance, i.e., high measurement loss, due to the unreliable wireless medium. In this paper, we experimentally show that WM can perform reliable and accurate measurements on wireless traffic, in actual, non-ideal channel conditions.

We demonstrate how to increase capture performance by merging traces from multiple monitoring devices. This merging enables WM to capture over 99% of the IP layer traffic and over 97% of the MAC/PHY frames in a controlled experiment. Our results indicate that WM enables reliable analysis of the collected traces, and should encourage the wireless research community to use this technique for a wide variety of research areas, such as traffic analysis, user mobility and handoff analysis, and MAC/PHY anomaly detection.

1 Introduction

With the growing popularity of IEEE 802.11-based wireless networks, it has become increasingly important to understand the characteristics of wireless 802.11 traffic and the wireless medium itself. A number of measurement studies have examined traffic characteristics in wireless networks [1, 5, 7, 8, 10]. These studies have measured the wired portion of the network, using wired network sniffers and SNMP polling. Wired network monitoring (WDM) can provide accurate traffic measurements as seen in that portion of the network. They may not, however, disclose characteristics of the wireless medium (the 802.11 MAC/PHY), as wired devices can only see the traffic that is successfully transmitted to the wired side of the AP. While SNMP-based approaches may be able to retrieve such detailed wireless MAC/PHY information through the use of a properly defined MIB (Management Information Base), most existing SNMP MIBs for APs (MIB-I (RFC 1066), MIB-II (RFC 1213), and 802.11 MIB (IEEE Std 802.11-1999)) provide very limited visibility into MAC-level behavior. A further drawback of SNMP-based approaches is that they require an interval between SNMP polls (typically every 1–5 minutes), and it has been shown that long poll intervals may miss wireless clients that associate with APs for less than this poll interval [7].

To overcome the shortcomings of SNMP and WDM, it is necessary to sniff the wireless medium itself. We refer to this technique as wireless monitoring (WM). Like WDM, WM involves a set of devices, commonly referred to as *sniffers*, which observe network traffic, but in WM, the sniffers are equipped with wireless cards for sniffing the wireless medium. WM has recently been adopted in both wireless networking research, e.g., [9], and commercial WLAN (Wireless Local Area Network) management product development.

There are three advantages to using WM. First, WM captures detailed wireless-side traffic statistics. Second, WM provides per-frame wireless MAC/PHY information, such as 802.11 MAC headers. Third, WM does not require any interaction with the existing network, unlike WDM, where network sniffers need to be attached directly to wired switches.

The data collected by WM can be used for many purposes. Physical layer information, such as error rates, can be used to develop accurate error models for 802.11 WLANs, and for site-planning to determine the signal strengths required to achieve a certain throughput or error rate. Link-layer data, such as the characteristics of data, control and management frames, can be used to develop 802.11 simulation models, and to identify anomalies in the operation of the 802.11 MAC protocol. The overall traces themselves can also be used for emulating 802.11 networks.

WM, however, can be complicated to conduct in practice. Unreliable and varying wireless channel conditions may lead to measurement loss. The goal of this

study is to demonstrate that WM can perform reliable and accurate measurement under such non-ideal conditions. We first demonstrate how to improve the *capture performance*, that is, the amount of the actual wireless traffic captured by a particular measurement technique, by merging multiple sniffer traces. Then, through a controlled experiment, using clients with varying signal conditions, we quantify WM's capture performance in terms of IP and MAC layer statistics.

In this paper, we address all the above problems for accurate measurement technique. However, WM has another big challenge: *scalability*, i.e. that the cost and management overhead can be significant for the deployment and management of a large number of sniffers. In this work, we limit our work to the fixed number of sniffers for relatively small coverage area (e.g., WLAN in a single floor with less than 10 APs). Based on the promising results of this work, we are currently working towards addressing the scalability problem in more general WLAN environment.

The rest of the paper is organized as follows. In Section 2, we discuss previous measurement studies of 802.11 WLANs. Section 3 describes the setup and implementation of our WM system. In Section 4, we describe a controlled experiment to demonstrate the accuracy of the WM technique in capturing IP and MAC/PHY layer statistics in an actual environment. Finally, we conclude the paper in Section 5 and highlight our ongoing work.

2 Related Work

There have been several measurement studies of 802.11 WLANs. These studies typically use three types of measurement techniques: WDM (wired monitoring), SNMP and syslogs (AP system logs). WDM has been used for identifying the typical traffic mix in university WLANs [7, 8, 10], or public WLANs [1]. SNMP provides information on both traffic volume and the number of active (associated) users, and has thus been used for both traffic studies [1, 8, 10] and user mobility studies [2]. Syslog records detail steps of association, and have been used effectively for studying user activity patterns [5, 8].

WM techniques have been used by [4, 6] to measure packet loss and bit error rates in a non-802.11 wireless network. They used a controlled environment to measure traffic between two wireless stations. Our work differs in that it examines a production 802.11 network for MAC traffic characterization and diagnosis.

3 WM Technique

In this section we describe our methodology, in which we use multiple sniffer devices and merge multiple datasets to improve the capture performance of the WM technique.

3.1 WM Setup

To capture wireless frames, we used three network sniffers, each comprising a PC running Linux with the 2.4.19 kernel. Each sniffer had a Prism2 chipset-based wireless network interface card; two sniffers had Demarctech DT-RWZ0-200mW-WC cards, and the third had a Linksys WPC11v3 card. To measure traffic, we used the *Ethereal* protocol analyzer (version 0.9.6) with the *libpcap* library (version 0.7). Each card was placed into 'monitor mode', which allowed the card to capture 802.11 frame information on a target channel.

The sniffers captured the first 256 bytes of each observed 802.11 frame, recording the complete view of the frame, i.e., PHY/MAC/LLC/IP/Above-IP information. PHY information, such as MAC Time and SNR (signal-to-noise ratio), can be captured using Prism2 monitor header, which is not a part of the IEEE 802.11 frame header, but is generated by the firmware of the receiving card.

3.2 Implementation of WM system

In this section, we briefly describe the WM framework, based on the techniques introduced in [11]. In that work, we demonstrated two serious drawbacks of using a single sniffer: each sniffer experiences severe loss in captured frames, and each sniffer only observes its *local view*, that is, the frames observed by one sniffer, which may differ from the AP's *global view*. Our framework aims to improve the capture performance by using multiple sniffers, placed according to SNR measurements.

3.2.1 Merging multiple sniffers

Multiple sniffers can reduce measurement loss in two ways. First, a single sniffer may not be able to observe all of the frames sent to and from a particular AP, due to radio reception and range. By using multiple sniffers, we can aggregate each sniffer's local view to create a closer approximation of the AP's global view. Second, even if a sniffer had identical radio hardware and positioning to that of an AP, it may be useful to observe the frames that the AP itself was unable to receive.

To accurately merge data from multiple sniffers, we need to be able to distinguish unique 802.11 frames for removing duplicates. We also need to prevent reordering upon merging. Reordering may occur when different

sniffers observe disjoint sets of frames. For instance, if there are four frames f_{1-4} transmitted on a WLAN, and sniffer A sees f_1 and f_3, but sniffer B sees f_2 and f_4. Although each sniffer has observed their respective frames in relative order, it is impossible to use this relative order to merge the four frames. To prevent such duplication and reordering, we need to synchronize multiple sniffers' timestamps.

Our WM framework uses 802.11 Beacon frames, which are generated by the AP, as the frame of reference for all the sniffers. Beacon frames contain their own 64-bit absolute timestamps as measured by the AP, and we can therefore uniquely identify such common beacon frames in different sniffer traces. On the timestamps of such common frames, we took one of the sniffers as a reference point and used linear regression to fit the other sniffers' timestamps to the reference sniffer.

To prevent duplication and reordering, the time synchronization error (the difference between two timestamps of different sniffers for the same frame) needs to be less than *half* the minimum gap (G_{min}) between two valid IEEE 802.11 frames. In the IEEE 802.11b protocol, the minimum gap, G_{min}, can be calculated as the 192 μs (microsecond) preamble delay plus the 10 μs SIFS (Short Inter-Frame Space) and the 10 μs minimum transmission time for a MAC frame (for the case of an Acknowledgement frame) to be a total of 212 μs. Therefore, the time synchronization error needs to be less than 106 μs. Applying linear regression for each Beacon interval ($\approx$ 100ms) on 24 hours of traces from our test setup, we measured synchronization errors on the Beacon frames from another AP. We observed a maximum error of 30 μs, which is well below the 106 μs requirement. Our setup was thus suitable for measurement using multiple sniffers.

3.2.2 Sniffer placement

We used SNR measurements to place our multiple sniffers. One sniffer was placed adjacent to the AP, to be responsible for capturing the From-AP traffic and the traffic of clients near the AP. The other sniffers were placed as close as possible to the wireless clients. Assuming that clients are uniformly distributed over the coverage area, this meant placing the sniffers so that they cover as much of the AP's coverage area as possible. Generally, if we have n sniffers to place, we split the AP coverage area into n equal areas and place the sniffers in the center of mass of these areas.

To determine the AP coverage area, we first used the SNR (obtained from Prism2 header) seen in Beacon frames from the target AP to draw the *contour lines* (as shown in Figure 1). The AP coverage area was then determined by choosing a particular SNR contour, e.g., the 15-dB contour line.

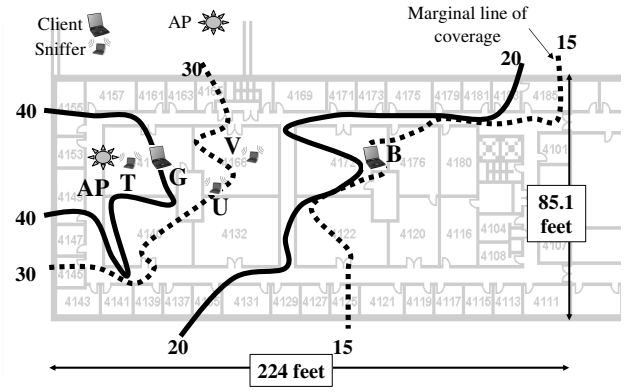

Figure 1: SNR Contour Map for controlled experiment: SNR Contour lines for 40, 30, 20 and 15 dB were obtained from SNR measurements. We placed the wireless clients at the locations with different signal conditions based on SNR measurement. Sniffers were placed at locations T, U and V.

We can refine this strategy by noting that, in an environment where multiple APs are installed, the coverage area of an AP may be reduced to the *Association Area* of the AP. The Association Area of an AP is the area at which a client will favor this AP for association compared with other APs in the area. This behavior may be device-specific and may also vary depending on whether a client has roamed to an area or has just powered on their radio. For the purposes of sniffer placement, we assume that clients will associate with the AP with the highest SNR [9].

4 Analysis of WM Technique

In this section, we discuss the accuracy of the WM technique in capturing IP and MAC layer statistics.

For IP layer statistics, we conducted a controlled experiment, using four different measurement techniques: Application-level measurement (APP-Level), SNMP, WDM, and WM. We then analyzed the differences between WM and these other techniques. For the MAC layer statistics, we analyzed the MAC sequence numbers from the same experiment to obtain the number of missing frames in the WM results.

4.1 Controlled Experiment

4.1.1 Application level measurement

To estimate the exact measurement loss, we needed to use reliable application-generated sequence numbers. We conducted a two-way UDP packet exchange experiments using an end-to-end traffic measurement tool, *NetDyn* [3]. Our NetDyn setup consisted of three different

processes, *Source*, *Echo* and *Sink*. *Source* inserted a sequence number in the payload, sent the packet to *Echo*, which also added a sequence number before forwarding it to *Sink*. The *Source* and *Sink* processes ran on a wireless station, while the *Echo* process ran on a server connected to the wired LAN. Using the sequence numbers generated by the *Source* and *Echo* processes, we were able to determine which packets were lost in the path from the *Source* to the *Echo* and vice versa.

In our experiment, *Source* sent 20,000 1500-byte packets at 100 packets/second. We ensured that no fragmentation occurred on either side of the AP. Therefore, for each NetDyn frame on the wireless side, there was a corresponding frame on the wired side. We used the number of the NetDyn packets (that were not lost in the paths), as the baseline for comparison with other monitoring techniques.

4.1.2 SNMP and wired monitoring

A wired sniffer running *Ethereal* was installed on the same LAN as the AP and the NetDyn *Echo* machine through a *Century Tap*, a full-duplex 10/100 Ethernet splitter. The same sniffer machine also ran a SNMP client that was configured to poll the AP for SNMP statistics every 60 seconds.

4.1.3 Experimental setup

As shown in Figure 1, we used two wireless clients at two different locations corresponding to two different signal conditions. The "Good" client G laid in an area of good AP coverage, in terms of SNR (the 40 dB-line), while the "Bad" client B laid in an area of bad AP coverage (the 15 dB-line). We also had three wireless sniffers (T, U and V) capturing the wireless traffic between *Source*, *Sink*, and the AP. Sniffer T was placed adjacent to the AP, while U and V were placed.

4.2 Application Layer Capture Accuracy

Table 1 compares the three traffic measurement techniques, using the NetDyn results as the baseline. To-AP traffic represents traffic from the clients to the AP, while From-AP traffic represents traffic from the AP to clients. Note that for SNMP statistics, we based our analysis on MIB-I counters as in [1] as well as the MIB-II counters (RFC 1213, RFC 2665). MIB-II provides many variables for calculating inbound/outbound error statistics more accurately than MIB-I [11].

From Table 1, we can make the following observations:

- WM has comparable performance to the other techniques for the *common* information that can be captured by other techniques, such as traffic on the wired side of the AP.

Table 1: Comparison between APP-measurement with NetDyn, WM, WDM, and SNMP, in terms of capture percentage

	NetDyn	WM	WDM	MIB-I	MIB-II
From	*To-AP Wireless Traffic*				
G	100	98.6	100	N/A	N/A
B	100	100.1	100	N/A	N/A
Total	**100**	**99.3**	**100**	**100.2**	**100.2**
To	*From-AP Wireless Traffic*				
G	100	99.4	103.4	N/A	N/A
B	100	102.6	103.5	N/A	N/A
Total	**100**	**100.9**	**103.5**	**102.0**	**99.9**

- The MIB-I and MIB-II SNMP statistics cannot reveal per-client information. The 802.11 MIBs and AP-specific MIBs may provide per-client information; we do not consider these here.

- Wired monitoring can provide accurate To-AP information about the wireless medium through the proportion of successfully-transmitted frames, as the probability of the loss on the wired medium is much lower than the probability of loss on the wireless medium. If, however, frames are fragmented on the wireless medium, we cannot obtain correct statistics on the wireless frames from the wired side.

- For the From-AP traffic, although wired monitoring can provide per-client information for the wired segment, it *overestimates* the actual traffic compared to WM. This is due to the noisy characteristics of the wireless channel, which lead to the loss of many packets on the wireless side that wired monitoring cannot capture.

- It is interesting to notice that even the SNMP statistics may differ from the true view of the wireless client. For example, in Table 1 the MIB-II total number of successfully transmitted packets is less than the number of packets received by the NetDyn *Sink*. This can be explained by noting that there may be packets that were successfully received by the *Sink* after three retransmissions, and the corresponding MAC-level ACK was sent back. This ACK, however, was not received by the AP, and so the AP did not count it as a successful transmission.

- Due to client B's bad location, the number of successful transmissions at client B was smaller than at client G. For example, client B successfully

Table 2: Capture percentage by sequence number analysis: Merging data from sniffers T, U and V significantly increases the number of captured frames. The adjusted column ignores frames that are sent by a device on channels other than the target channel (e.g., probe request frames).

From	*T*	*U*	*V*	*T+U+V*	***Adjusted***
AP	97.85	96.88	96.14	98.53	98.53
G	97.48	93.65	92.92	97.72	97.72
B	61.09	88.51	88.67	88.96	**93.31**
Total	88.92	93.91	93.40	95.74	**96.97**

exchanged 18,490 (92.5% out of 20,000) NetDyn packets, while client G successfully transmitted 19,905 packets (99.5% out of 20,000). Using the number of successfully-received NetDyn frames as the 100% benchmark, we observe that for client B, the sniffers captured more than 100% of the NetDyn frames, e.g., 102.6% for traffic from AP to B in Table 1. In other words, for clients with bad signal strength conditions, WM can capture more frames than are successfully transmitted to/from the clients at application layer.

Finally, we note that the WM technique statistics are within **1%** of the actual application layer statistics.

4.3 MAC Layer Capture Accuracy

In this section, we examine WM's accuracy in measuring the MAC layer traffic. To count the measurement loss with only MAC frames, we exploited the IEEE 802.11 MAC sequence number. A wireless device increments the MAC sequence number whenever it sends a new Data/Management frame. We focus our analysis in this paper on the Data and Management frame types, as Control frames do not contain a sequence number. If a device retransmits a frame, then it uses the same sequence number as the original frame. Since the maximum MAC sequence number is 4095, a wireless device reuses the same sequence number every 4096 unique frames. We denote the difference in sequence numbers of two consecutive captured frames as the *gap size*. For example, if consecutive frames have sequence numbers, 4094 and 1, then there is a gap of size three ($mod_{4096}(1-4094)$) between the frames, and there are two missing frames between them.

Table 2 analyzes the sequence numbers for the same controlled experiments as represented in Table 1. As MAC sequence numbers are generated per device, we examined the MAC sequence numbers of the AP, client G and client B separately. Table 2 shows that as data

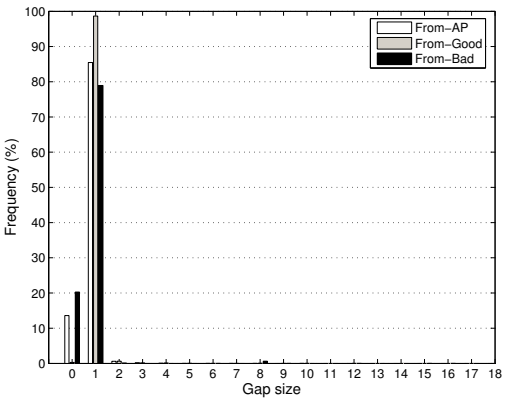

Figure 2: Distribution of sequence number gaps in From-AP, From-Good, and From-Bad MAC traffic.

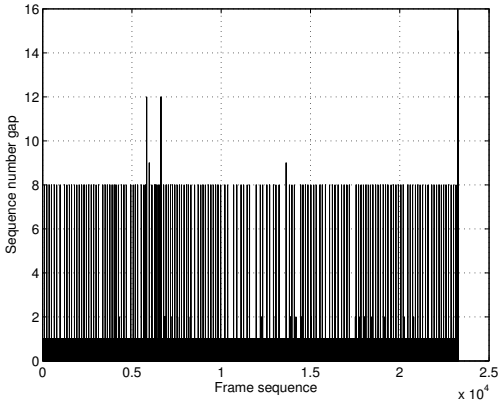

Figure 3: Sequence number gap in From-Bad MAC traffic.

from sniffers T, U, and V are merged, the capture performance for AP, G and B increases to 98.5%, 97.7%, and 89% respectively (with a further increase to 93% in the discussion below). The numbers in the table reflect the gap between different frames captured by the sniffers.

Figure 2 shows histograms of the gap size between two consecutive frames for the case of using the three sniffers. Note that a gap of zero means that a frame has been retransmitted, while a gap of 1 means that there are no missing frames between these two frames. To calculate the number of missing frames, we count the gaps of greater than 1. Let $freq_i$ denote the number of occurrences of gap size i in Figure 2. Then, we can calculate the number of missing frames (N_{miss}) by $\sum_{i=2}^{i=\infty}(i-1) \times freq_i$. The column labeled 'T+U+V' in Table 2 can be obtained by $\frac{N_{cap} \times 100}{N_{cap}+N_{miss}}$, where N_{cap} denotes the number of distinct captured frames.

Figure 3 takes a closer look at the Bad client case. The x-axis represents the received frame sequence number and the y-axis represents the gap before this frame. There was a periodic gap of 8, i.e., a measurement loss of

7 frames. By looking into the traces we found that these periodic behavior was due to the Bad client performing periodic active scanning searching for better APs. Since this process involves sending probe request frames on *different* channels [13], the sequence numbers were not captured by our sniffers, which sniffed the traffic on only one channel. If these missed probe request frames are added to the loss statistics in Table 2, the capture accuracy for the Bad client increases to 93.31% and the overall capture performance increases to 96.97% (as shown in the *Adjusted* column in Table 2). This means that the WM statistics differ by at most **3.03%** from the actual statistics.

5 Conclusion

In this paper, we demonstrated that a WM technique can perform reliable and accurate measurement on the 802.11 traffic, in non-ideal channel condition. We discussed how to merge traces from multiple sniffers to increase the wireless monitoring capture performance. Results from a controlled experiment, with clients having different signal conditions, show that the WM technique captures wireless side statistics with 1% and 3% error bounds at IP and MAC layers, allowing a reliable analysis of the collected traces.

We believe that the described controlled experiment presents a worst case scenario for the WM technique. This is because our Bad client case, where the client is located on the marginal line of AP coverage, is severe and unlikely to occur in a real environment. The MAC layer capture performance of the WM technique would be much better in a WLAN environment where APs are positioned so that clients in most locations of interest can find at least one AP with good signal conditions.

Based on the results in this paper, we are currently using the WM technique to analyze the WLAN traffic in a Computer Science department environment. Some initial results can be found in [12] on traffic characterization and in [11] on anomaly detection. Besides characterizing WLAN usage patterns, we are using the traces for multiple APs to analyze user roaming patterns, co-channel interference and interactions between different APs. In such experiments, we expect that combining wired monitoring data, such as *Inter Access Point Protocol* information (IEEE Std 802.11f), with WM analysis, would give better measurement capabilities, e.g., on the roaming behavior of the mobile users and the handoff process.

We believe that our results should encourage the wireless research community to use WM techniques in many research areas, including traffic analysis, user mobility and handoff analysis, and MAC/PHY anomaly detection.

References

[1] A. Balachandran, G.M. Voelker, P. Bahl, and V. Rangan. Characterizing User Behavior and Network Performance in a Public Wireless LAN. In *Proceedings of ACM SIGMETRICS '02, Marina Del Rey, CA*, June 2002.

[2] M. Balazinska and P. Castro. Characterizing Mobility and Network Usage in a Corporate Wireless Local-Area Network. In *Proceedings of MOBISYS '03, San Francisco, CA*, May 2003.

[3] S. Banerjee and A. Agrawala. Estimating Available Capacity of a Network Connection. In *Proceedings of IEEE ICON '01*, September 2001.

[4] B.J. Bennington and C.R. Bartel. Wireless Andrew: Experience building a high speed, campus-wide wireless data network. In *Proceedings of MOBICOM '97*, September 1997.

[5] F. Chinchilla, M. Lindsey, and M. Papadopouli. Analysis of Wireless Information Locality and Association Patterns in a Campus. In *Proceedings of INFOCOM '04, Hong Kong, China*, March 2004.

[6] D. Eckardt and P. Steenkiste. Measurement and Analysis of the Error Characteristics of an In-Building Wireless Network. In *Proceedings of SIGCOMM '96*, August 1996.

[7] Tristan Henderson, David Kotz, and Ilya Abyzov. The changing usage of a mature campus-wide wireless network. In *Proceedings of MOBICOM '04*, pages 187–201. ACM Press, September 2004.

[8] D. Kotz and K. Essien. Analysis of a Campus-wide Wireless Network. In *Proceedings of MOBICOM '02, Atlanta, GA*, September 2002.

[9] M. Shin, A. Mishra, and W. Arbaugh. Improving the Latency of 802.11 Hand-offs using Neighbor Graphs. In *Proceedings of MOBISYS '04, Boston, MA*, June 2004.

[10] D. Tang and M. Baker. Analysis of a Local-Area Wireless Network. In *Proceedings of MOBICOM '00, Boston, MA*, August 2000.

[11] J. Yeo, M. Youssef, and A. Agrawala. A Framework for Wireless LAN Monitoring and its Applications. In *Third ACM Workshop on Wireless Security (WiSe'04), Philadelphia, PA*, October 2004.

[12] J. Yeo, M. Youssef, and A. Agrawala. Characterizing the IEEE 802.11 Traffic: Wireless Side. In *CS-TR 4570, Dept. of Computer Science, University of Maryland*, March 2004.

[13] M. Youssef, L. Shahamatdar, and A. Agrawala. The IEEE 802.11 Active Probing Mechanism: Analysis and Enhancements. In *CS-TR-4613, Dept. of Computer Science, University of Maryland*, August 2004.