Managing and Monitoring Spectrum Usage in a Wireless Network

by

Ilsun Lee

and

Yuxiang Liu

A thesis

presented to the Carnegie Mellon University

in fulfillment of the

thesis requirement for the degree of

Bachelors of Science

in

Computer Science

Pittsburgh, Pennsylvania, USA, 2006

Author's Declaration

We hereby declare that we are the sole authors of this thesis.

We authorize Carnegie Mellon University to lend this thesis to other institutions or individuals for the purpose of scholarly research.

Signature

We further authorize Carnegie Mellon University to reproduce this thesis by photocopying or by other means, in total or in part, at the request of other institutions or individuals for the purpose of scholarly research.

Signature

Abstract

Today there are a rapidly increasing number of wireless devices such as access points and wireless enabled laptops in our homes and offices. Since there are only three non-interfering channels between 2.41 GHz and 2.48 GHz, the increasing competition for these three channels of communication will cause increasing interference, thus decreasing the performance of these wireless devices. To diagnose problems caused by interference, the wireless network administrator must obtain a map of the radio propagation and interference of the site of network deployment. Our aim for this thesis is to implement a cheap, easy to deploy real time monitoring system for wireless networks. By combining local signal maps from individual monitoring stations into a global map, we can create a real time signal map of an entire wireless network. Our results indicate that there is a high correlation between RSSI variance and workload. Also, in our deployment site, most terminals are in the 802.11 b mode when the access points support 802.11 g mode. Finally, we were able to use our system to characterize the utilization at various locations in the deployment and find possible hidden and exposed terminals.

Acknowledgements

We would like to thank our thesis advisors Dr. Peter Steenkiste and Dr. Srinivasan Seshan for their support throughout the development of our thesis. This project would not be possible without their advice and encouragement. We would also like to thank Mark Stehlik, who had made our undergraduate lives so fun and interesting, and gave us this opportunity to perform research at this world class institute of learning. Finally, we would like to thank our friends who were there for us when things weren't going our way.

Table of Contents

List of Figures

List of Tables

Chapter 1
Introduction

1.1 Joint Thesis

This was a joint thesis project between Yuxiang Liu and Ilsun Lee. Yuxiang Liu was responsible for integrating and managing the database. Ilsun was responsible for the developing the passive monitoring stations. In the analysis of the data, Ilsun focused on the analysis of the general spectrum patterns in global and local view. On the other hand, Yuxiang developed a monitoring tool which provides an easy way to visualize the spectrum usage in the networks and used it to find hidden and exposed terminals. While chapters 1, 2 and 5 were written together, chapters 3 and 4 have a designated author in each section. Sections 3.2 and 3.3 were written by Yuxiang, and sections 3.1 and 3.4 were written by Ilsun. Also, in chapter 4, sections 4.2.1 to 4.2.5 were written by Ilsun and sections 4.2.6 to 4.2.9 were written by Yuxiang.

1.2 Motivation

Today there are a rapidly increasing number of wireless devices such as access points and wireless enabled laptops in our homes and offices. As the price of these wireless devices decreases and their convenience increases, their numbers will increase at an even faster rate in the near future. However, the bandwidth available to support these wireless devices is limited as these devices are only allowed to use the bandwidth between 2.41 GHz and 2.48 GHz due to government regulations [1]. Since there are only three non-interfering channels between 2.41 GHz and 2.48 GHz, the increasing competition for these three channels of communication will cause increasing interference, thus decreasing the performance of these wireless devices.

Let’s look at this problem through the eyes of a wireless network administrator. There are three stages in the life cycle of a wireless network: deployment, maintenance and expansion. In the deployment stage, the wireless network administrator must determine where to place wireless access points in order to maximize coverage and minimize interference. In the maintenance stage, the wireless network administrator must diagnose users’ complaints of poor wireless performance. Finally in the expansion stage, the wireless network administrator must determine where to place new wireless access points and how to configure them to minimize interference with the existing wireless access points.

To answer these questions the wireless network administrator must obtain a map of the radio propagation and interference at the site of network deployment. In open space, radio waves propagate according to the inverse-square law, so the strength of the radio wave at a location is inversely proportional to square of the distance from that location to the radio source [2]. However, in closed space, radio wave propagation is unpredictable as different surfaces attenuate radio waves differently [3]. Since most wireless networks are deployed in closed environments such as office buildings and houses, one cannot simply predict a signal propagation map from the floor plan of the deployment site, but rather one must measure the signal strength at each room of the floor plan.

While a site survey of the radio signal propagation can inform the wireless network administrator on where to deploy access points during the deployment stage, it does not help the wireless network administrator diagnosing network interference and poor performance in the maintenance stage. This is because the signal environment has changed with the addition of the new access points and with the continual turn over of wireless clients. Therefore, in order to diagnose real time wireless network problems with interference, a real time wireless signal map is needed.

1.3 Related Work

Providing a real time signal mapping of IEEE 802.11 networks is a relatively new research area. Adya, et al. [4] developed an architecture for detecting and diagnosing faults in the wireless network infrastructure. By placing special monitoring software on the clients and access points, they were able to detect disconnected clients, rouge clients and performance problems. The primary difference between our thesis and this paper is that our thesis focused on providing the signal map to explain and determine the causes of disconnected clients and performance problems, rather than detecting the occurrence of such problems.

1.4 Project aims

1.4.1 Implementing a cheap, easy to deploy real time monitoring system for wireless network

The first general aim was to implement a real time database for the wireless network in order to develop an accurate model of the signal environment. Easy installation and deployment of the passive monitoring station were the crucial properties to collect large and accurate data sets. It should allow user to make queries about the collected wireless network data.

1.4.2 Analyze the collected data and identify wireless network problem

The second general aim was to investigate the collected signal data in order to identify the potential wireless network problem. Problems we will try to answer are determining the location and prevalence of hidden and exposed terminal problems.

1.4.3 Provide a graphical interface to monitor the network in real time

The third general aim was to design and implement a graphical user interface to the collected data in order to present to the user a real time global view of the data. This way, the user can diagnose wireless network problems and spot patterns in the data in real time without resorting to complex offline processing.

1.4.4 Implementing an spatial application

The last aim was to develop a solution to solve the wireless problem by developing algorithm to change the configurations of the client or router. Because of the lack of the time and support from the chosen wireless drivers, this aim was not accomplished.

Chapter 2
Real time site survey

The main goal of our thesis is to provide a system that can produce a real time site survey of a wireless network deployment. At the same time, we will also try to determine the location and prevalence of hidden and exposed nodes.

2.1 Real time wireless site survey in a nutshell

A static wireless site survey is a physical survey of the wireless network deployment site to identify how radio waves are propagated in that environment. With a map of the radio wave propagation resulting from a wireless site survey, a wireless network administrator can determine the optimal location to deploy new access points to provide maximum coverage and performance at the deployment site. A static wireless site survey usually involves several people carrying hand held spectrum analyzers and measuring the radio frequency (RF) patterns at every location at that deployment site that requires wireless network access. Such process is time and labor intensive and at the same time only provide a static view of the RF propagation and interference.

A real time wireless survey monitors the wireless network deployment site in real time. This monitoring is an automated process performed by monitoring stations placed strategically within the wireless network deployment site, where each monitoring station reports the RF propagation and interference around its area in real time.

2.2 Implications

The wireless network topology changes continuously with the constant flux of wireless clients and people in the wireless network. With a real time wireless site survey, we gain an accurate and real time signal map of the wireless deployment site. Using this signal map, we can continuously monitor the strength and locations of the interference caused by access points and clients, and diagnose users' complaints of poor wireless performance. By knowing the problem areas of the wireless network, a wireless network administrator can either modify the configuration of existing access points or add new access points.

2.3 Challenges

However there are many challenges in creating a real time wireless site survey. The first challenge is how to monitor a dynamic wireless environment continuously. We need a platform that can be easily deployed and reliably gather data 24 hours a day without maintenance. The second challenge is determining the number of monitoring stations is required to provide sufficient coverage for a wireless deployment. The third challenge would be how to weave the different observations from multiple monitoring stations into one complete picture of wireless environment and mine useful information from these deep set of information.

Chapter 3
Implementation

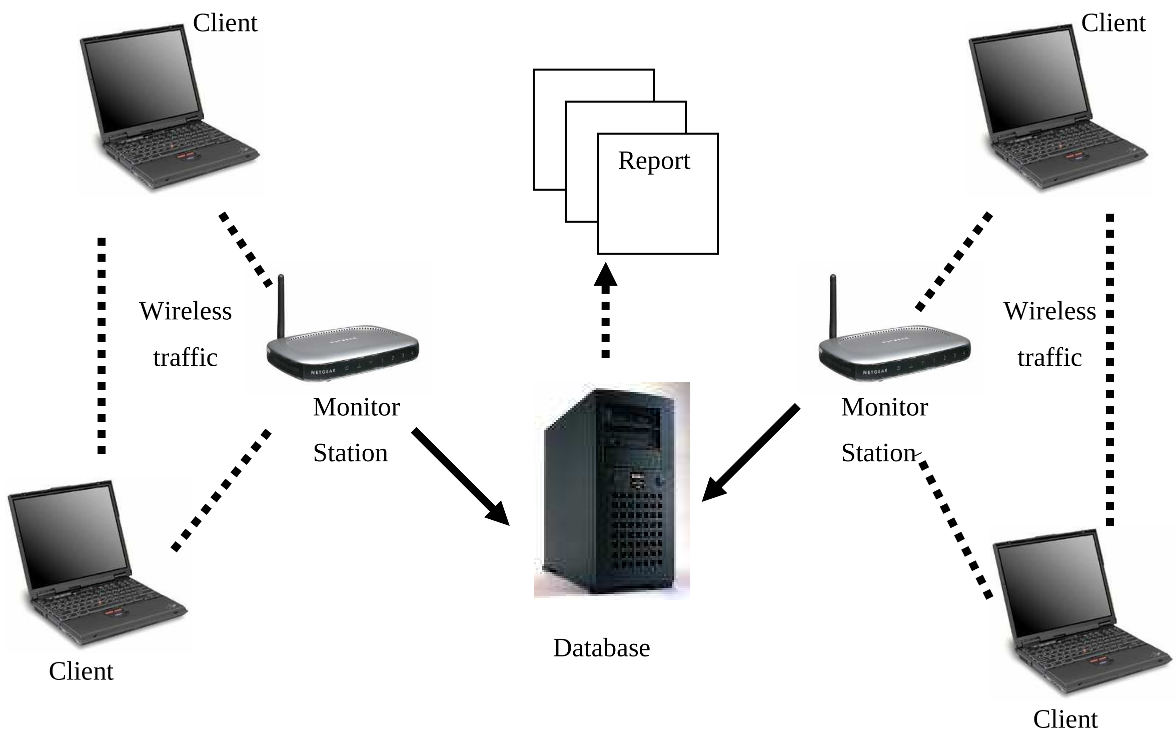

Figure 3.1 Overview of our solution to the real time signal map

Our solution to the real time signal map involved several monitoring stations strategically placed in the wireless network and set to passively monitor the wireless traffic in their area. Every minute, an aggregate of the statistics on the wireless traffic would be transmitted to the database. Finally, queries were made to the database to generate the reports required to analyze the data collected by the monitoring stations.

3.1 Monitor Station

3.1.1 Router

Netgear WGT634U [5] Access Points were used to implement the passive monitoring station. Netgear WGT634U has a 200MHz MIPS32-like CPU (Broadcom BCM947XX), 32MB RAM, 8MB flash, an Atheros AR5213-based 802.11b/g card, two Ethernet interfaces, a USB 2.0 host port and run Linux 2.4 kernel. Netgear WGT634U is cheap, small and easy to install which enabled the easy deployment of the monitoring stations.

In addition, Netgear WGT634U provided versatile tools which enabled easy installation of software. Previously, Roofnet [6] project from M.I.T have used identical router as the test beds for a wireless ad-hoc network. Roofnet project provided a cross compiler, SSH daemon and other necessary tools with directions which helped enormously during the installation of the software to router.

3.1.2 Process Overview

Figure 3.2 provides the basic overview of the monitoring station's process flow. There are four stages in the monitoring station's process flow:

1. Device Driver: Responsible for the receiving incoming packet and collect hardware generated information such as RSSI.
2. Click Modular Router: Provide easy interface to add application
3. Data Aggregator: Collect and validate the collected information
4. Client Program: Send out the collected data to the central database

First, the incoming packet with data and header was received by the device driver. Then, hardware information such as RSSI value was added to the packet header by the device driver. Next, this packet was passed to the Click Software router. In the Click, only the necessary data information was extracted from the packet headers and unnecessary information such as data portion of the packet was discarded. Extracted data information was then passed to the data aggregator. The data aggregator then sorts the passed in data by its source ID and packet's transmission rate. Lastly, the client would send the collected data in every minute to the central database and clears the data aggregator for the next use. Detailed information about each stage's implementation will be described in following sections.

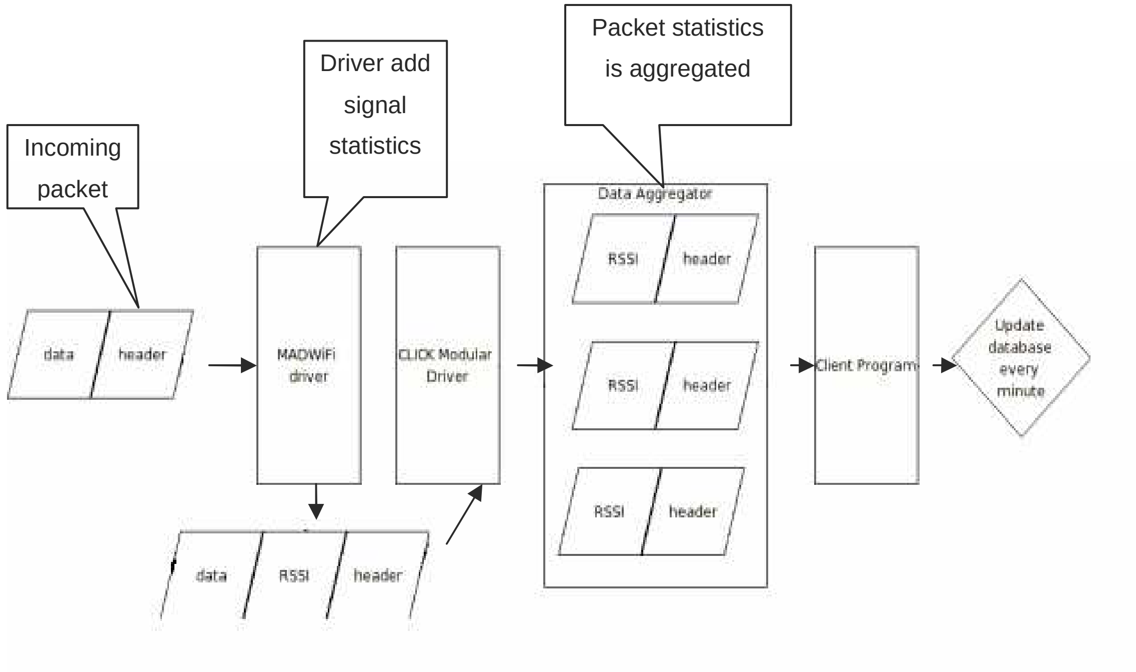

Figure 3.2 Basic overview of the monitoring station's process flow

3.1.3 Device Drivers

The main functionality for the device driver was to obtain the hardware specific data such as the Received Signal Strength Indicator (RSSI). Madwifi-old driver was used as the wireless driver in our monitoring stations. Madwifi [7] is an open source project to support a Linux kernel driver for Wireless LAN chipsets from the Atheros [8]. There were multiple versions of the Madwifi drivers in presence such as Madwifi.stripped, Madwifi-old and Madwifi-ng.

Madwifi.stripped was a fork of the Madwifi-old, which was modified to support special functionality with the Click Modular Router. However, all the Madwifi.stripped functionalities were incorporated into the Madwifi-old driver and Madwifi-old provided advanced functionality and stability compared to the Madwifi.stripped version. Madwifi-ng was the latest driver for the Madwifi. Unfortunately, Madwifi-ng was very unstable in current 2.4 Kernel environment of the router. As a result, we decided to use the Madwifi-old version because it was stable and supported all the necessary functionality for Click Modular Software we used.

3.1.4 Click Modular Router

Click [9] is modular software router written in C++. Click routers are flexible, configurable and easy to understand. A Click router consists of an interconnected collection of modules called *elements*; elements control every aspect of the router's behavior, from communicating with devices to packet modification to queuing, dropping policies and packet scheduling. One can easily construct a software router by writing configuration script, which glues elements together. We chose to use the Click Modular Router because of its easy extensibility.

To implement our monitoring software, we have modified the *PrintWifi* element. *PrintWifi* extracts the Received Signal Strength Indicator (RSSI), transmission rate, sender's MAC address and packet size from the incoming packet, then the collected information is sent to the data aggregator for the management of the data. For the configuration file, we have developed a very simple software router with *PrintWifi* element. One can easily develop a new element and add it to the configuration file to have extra functionality.

3.1.5 Data Aggregator

Data Aggregator is responsible for the management of the collected data information. Using the basic single linked list data structure, we have sorted the received packet information based on the sender's MAC address and transmission rate. Upon the receipt of each sender's MAC address and transmission rate, we recorded the mean of the RSSI, total number of the packet sent, total size of packet in bytes.

3.1.6 Client Program

Client program was responsible for sending the aggregated data to the central database. We have decided to use the existing wireless networking connection rather than the Ethernet connection for the easy installation and deployment. Frequent update to the database can cause significant interference with the existing wireless network. As a result, we have decided to send the collected data to the central base on one-minute intervals. Our simple client program was developed with the Berkeley Socket API using C.

3.1.7 Installation to Router

Because of the unique booting sequence and the different hardware specification of the router, we need to perform several steps to install our developed software into the router.

1. Install SSH Daemon: Because router does not have an output device such as monitor, we have installed a remote shell (sshd) to the router. Roofnet project have already provided an image of roofnet with sshd which enabled easy installation of the sshd. Using the web interface which was provided from the original firmware, we updated router's flash image with the image of roofnet. With the help of the sshd, we were able to login to the router via ssh.
2. Cross Compile Software: All the developed programs need to be cross-compiled in order to run on the router because of the different hardware specification. We have used hndtools-mipsel gcc 3.2.3 toolchain with the original firmware source code from Netgear to compile our software. This toolchain is identical to toolchain that was used in the roofnet project.
3. Make Permanent Change to Flash: The WGT634U's file system operates out of a ramdisk, so any changes made to the filesystem will not persist across reboots. To make a permanent change, we need to mount the flash and update the flash with our developed software.
    a. *mount /dev/mtdblock3 /mnt* will mount the flash memory.
    b. In */mnt* there will be a lot .lrp files which are .tar.gz archives. These files will be extracted to */* on ramdisk at boot, before any of the init script are run. */mnt/lrfkg.cfg* contains the list of .lrp files which will be extracted during the boot up stage
    c. We have created a new roofnet.lrp file to contain all of our necessary software and edited the init script to run our required software during boot up. Since the *roofnet.lrp* was already listed in */mnt/lrfkg.cfg* we did not modify the */mnt/lrfkg.cfg* file.
    d. Using the scp protocol, *roofnet.lrp* was transported to mounted flash directory to update the flash memory.

3.2 Database

3.2.1 Requirements

There were many possible solutions to the problem of real time storage and retrieval of the radio signal propagation data. The requirements for our solution were:

- Be able to efficiently store signal data from the monitor stations in real time and process data queries from users simultaneously.

- The database components must be flexible enough so that each component can be interchanged easily should new techniques or technologies be discovered.
- Simple to implement.
- The database must be free so we must use an open source libraries or develop our own libraries.

3.2.2 Message format

We required a flexible message to facilitate the communication between our database and the monitoring stations. While we could have used a custom messages format to save on bandwidth, we decided to use the XML based SOAP [10] message format for the messages between the database and the monitoring stations. This was because, XML based SOAP messages are human readable and easy to debug. In addition, an open format based on XML allows clients created by other people to easily interoperate with either our database or our monitoring stations' software. By making our message format to be as flexible as possible, we were able to maintain modularity in both our database and monitoring stations.

3.2.3 Front end server

The front end server was the gateway to the database. Its function was to translate the data inputs from the monitoring stations into SQL statements for insertion into the database. As such, it must be reliable enough to handle simultaneous connections from multiple monitoring stations and be efficient enough to translate their data inputs into SQL statements in real time. We could have modified the server built into most databases to handle the translation, but due to our unfamiliarity with the database server design, we decided not to approach this route. The other option would be to modify existing servers or create our own server to handle the translation of the monitor station inputs.

We decided modify a stable and open source web server, BOA [11], because we require a reliable system that can handle a relatively large number of transactions in a stable manner. By building from a stable base, we be assured that during debugging, most of the errors will not be in the low level TCP connection handling but rather in the high levels of XML parsing and database interaction.

BOA was a relatively simple and efficient web server. It was written in C for speed and memory efficiency, so all our modifications to the server must also be in C. This presents a new set of challenges since now we need to worry about memory allocation and management, and pointer manipulation. However, we believe it was worth the extra debugging effort as the final system was extremely stable as it had not crashed after months of usage and more importantly consumed very little system resources.

Since the monitor stations' messages were in SOAP / XML format, we added a XML parser to the BOA server in order to translate the monitor stations' data into SQL statements. Due the dominance of JAVA in the SOAP / XML web services domain, it was rather difficult to find a stable and open source XML parser written in C. We decided to use the open source XML parser, EXPAT [12] to parse the monitor stations' messages. While EXPAT does not have as much features as most JAVA based XML parsers, it is sufficient to handle the simple XML messages from the monitor stations.

Finally, the parsed XML messages are passed to our own module where the messages are converted to SQL statements and passed onto the database via the database's API. We wrote this module ourselves in order to maintain the maximum flexibility in our database design.

3.2.4 Database structure

Since the database itself did not directly handle monitoring stations' data input, there was not as much pressure to keep the database as efficient as possible, but rather we emphasized on the flexibility to accommodate the various database designs we wanted to try. There are many open source and full featured databases to choose from and we choose POSTGRESQL [13] as it is the database that we were most familiar with.

There were three tables in our database: nodes, data_dump and location. The nodes table stored each source and destination MAC address of the wireless traffic detected by the monitoring stations. The data_dump table stored the individual 1 minute aggregates of the wireless traffic data from the monitoring stations. The location table recorded the locations of the identified nodes such as access points and monitoring stations.

The main goal behind this database design was to facilitate the time and spatial nature of most queries that we needed to generate a real time signal map. Most of the time, we just want to know what are the traffic conditions at all locations. Since wireless transmissions at different times do not interfere with each other, we do not need to correlate the data across measuring times. However, different transmitters can interfere with each other's transmissions, so we need to how these transmissions are propagated. By indexing the traffic data by their location where they are overheard, we constructed a web of interference through comparing the transmission sources overheard at each monitoring station.

3.3 Graphical Network Visualization

The main goal of the graphical network visualization was to present to the wireless network administrator a global view of the combined signal environment data gathered from individual monitoring stations. By unifying the signal environment data from multiple monitoring stations, we can present a relatively complete picture of the signal environment at the deployment site in real time. This way, a wireless network administrator may spot trends or problems with the wireless network at one glance, rather than relying on offline processing of the collected data.

We chose to use the JAVA graphics 2D animation tool kit to visualize the signal environment in our network due to the extensive tutorials on animation with JAVA available online [14]. While JAVA animation does require more CPU resources to run, its simplicity and flexibility more than make up this flaw with a relatively fast development time. Since JAVA animation applets is also platform independent, our network visualization should run equally well on Linux and Windows based systems.

3.3.1 Localization

While the monitoring stations can tell us many interesting information about the interference in the wireless environment at their locations, they cannot determine the source of the interference. There are many algorithms that can be used to determine the location of a particular transmitter [15] [16]. Since localization was not the main focus of our thesis, we simply chose an easy to implement and relatively accurate localization algorithm called pattern matching in order to localize the source of the interference.

The basic localization algorithm contains two parts: training and localization. In the training phase, we used a laptop to transmit to an access point from various locations in the deployment site. At the same time, we set the monitoring stations to record our laptop's RSSI values. By correlating the set of RSSI values recorded by each monitoring station and the laptop's transmission location, we were able to create a map of RSSI patterns at each location. In the localization phase, the transmitter's RSSI pattern recorded by the monitoring stations was then compared against this map and the location of the RSSI pattern with the lowest difference was returned as the possible location of the transmitter.

3.3.2 Network visualization

Figure 3.3 Screenshot of the Network Visualization Applet

The network visualization's goal was to give the wireless network administrator a real time view on the difference in propagation of wireless signals in a deployment. We can estimate the signal propagation and loss of each packet by looking at the loss of the received signal strength of the packets at each monitoring station. So for example, if a transmission originating at the 100 corridor arrives at monitoring station 1 at 30 dbm and at monitoring station 2 at 20 dbm, we can estimate the loss between monitoring station 1 and 2 at 10 dbm for transmissions originating at 100 corridor.

First the visualization presented to the user a floor plan of the deployment. The colored dots with the label AP represented the access points on that floor and the dots with label MS represent the monitoring stations. Please note that each monitoring station was a different color. The boxes in the visualization represented the transmitter sources in that floor. You can think of the boxes and colored dots as the transmission sources of various packets that are overheard at various monitoring stations. Next to these boxes and dots were numbers of various colors. The number represented the received signal strength of the packets overheard at the various monitoring stations and the color of the number correlate to the specific monitoring station that overheard the packet. So if a box in the 100 corridor had three numbers associated with it, it means a transmission originating at that box's location were overheard at three monitoring stations.

3.3.3 Utilization visualization

Figure 3.4 Screenshot of the Utilization visualization applet

The utilization visualization's goal was to give the wireless network administrator an idea on network utilization at the locations around each monitoring station. Network utilization was the percentage of the time that the monitored channel was busy carrying packets. Since at any single location, a single channel can only support one transmission at a time, areas of high network utilization means clients will need to contend harder for an open transmission slot and transmissions will have a higher rate of collisions at those areas.

Channel utilization at each minute was calculated by summing the times required to transmit all the packets each monitor station overhears in that minute. The time require to transmit a packet is the size of the packet divided by the transmission rate that packet is transmitted at.

3.3.4 Hidden terminal and exposed terminal visualization

Figure 3.5 Screenshot of the Hidden terminal visualization applet

Two of the problems that a wireless network administrator must solve are finding the locations and prevalence of hidden terminals and exposed terminals. Hidden terminals are nodes that are out of transmission range of other nodes. Suppose two nodes that are hidden relative to each other are communicating with the same access point, their transmissions could potentially arrive at the access point at the same time and interfere with each other's signals. Depending on the strength of the signal that arrives at the access point, either one or both nodes' transmission will not be received correctly by the access point. Exposed terminals are nodes that are within transmission range of each other but the access points that they are communicating with are out of transmission range with the access point that the other node is communicating with. So theoretically, both nodes should be able to simultaneously transmit interference free to their associated access points, but because both nodes are within transmission range of each other, they are unable to transmit simultaneously to their respective access points.

Because we did not install the monitoring software on the access points and clients, we were unable to determine exactly which nodes were hidden or exposed. However, we could find potential hidden and exposed terminals by assuming the channel conditions experienced by monitoring stations close to access points to be channel conditions at the access points, and by assuming the channel conditions experienced by monitoring stations at other locations to be channel conditions experienced by clients at those locations. So to determine potential hidden terminals, we set one monitoring station to the "access point" and two other monitoring stations to be the "clients". Then a terminal was determined to be a potential hidden terminal when its transmissions were heard by only one of the "clients" and by the "access point". So transmissions from these potential hidden nodes could potentially interfere with each other at the access point. Potential exposed terminals were determined by finding the terminals that are heard by the "client" but associated to the "access point". These nodes were exposed relative to the "access point" because while they are transmitting, other clients associated with the access point would believe the channel was busy but when in fact the transmissions by these potential exposed nodes would never arrive at the access point.

The visualization layouts for hidden terminal and exposed terminal were similar to the previous two visualizations. However the changes were that there are four additional buttons that allows one to select the "access point" monitoring station, and two "client" monitoring stations, and starting the hidden or exposed terminal detection.

3.4 Analysis Tools

3.4.1 Problems during analysis

We have analyzed various views of the wireless signal environment. While it was easy to analyze and graph the global views of the environment, it was problematic to analyze the local, per source view of the signal environment due to the large number of different sources. As a result, we have developed a MATLAB program which will plot the each source's individual signal environment automatically. In addition, the MATLAB program provided an easy way to plot different parameter with very small modification of the program.

3.4.2 Usage of Analysis Tools

The analysis tool consisted of two parts. The first part was the simple parser program which received an input of the global result of the data and output per source based data. The second part used per source based data and generated the graph for each source.

1. Parser: For the simple parser, we have developed a simple java program. The program required the input of the global data with the following format (time, monitoring station ID, source ID, parameters separated by the comma) then the program will output to the file with the filename of monitoring station ID_source ID. For example the monitoring station ID with

1308 and source ID with 253253 will produce file 1308_253253 which contains all the signal information that was collected in monitoring station in 1308 from source of 253253.

2. Graph Utility: The graph utility was developed using the MATLAB program. MATLAB program will prompt to enter the directory which contains all the output from the Parser program and the two parameters that will be plotted. After successful input the program will generate a graph each source.

Chapter 4

Experimental Results

4.1 Experiment setup

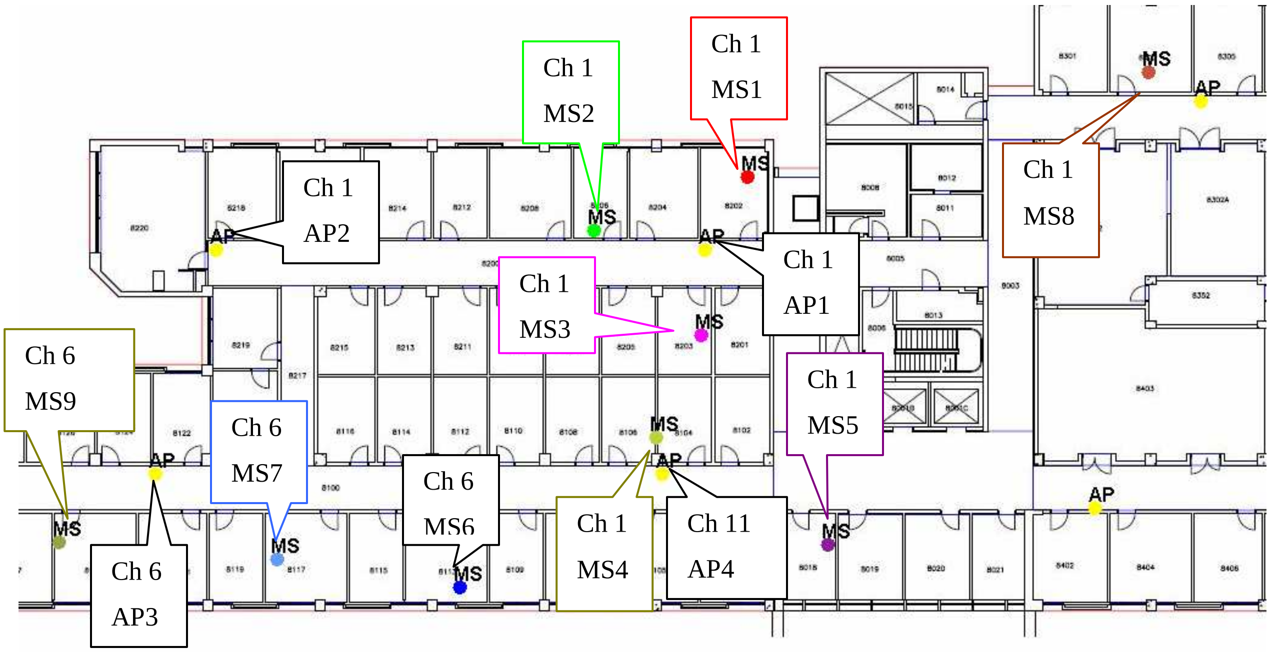

Figure 4.1 Deployment of the monitoring stations (MS) and locations of access points (AP) on Wean 8th floor

Monitoring Stations	Room Numbers
MS1	8202
MS2	8206
MS3	8203
MS4	8104
MS5	8018
MS6	8113
MS7	8117
MS8	8303
MS9	8123

Table 4.1 Location of the monitor stations

We decided to deploy our monitoring system in the 8th floor of Wean Hall at Carnegie Mellon University. The exact placements of the nine monitoring stations are shown on Figure 4.1. This figure also contains the IDs (MS1 through MS 9) of these nine monitoring stations. When we mention monitoring stations MS 1 through MS9 for the rest of this thesis, please refer back to this figure.

Wean Hall is 8 stories high and was constructed with mainly steel and concrete. The walls between offices were also constructed of concrete. All of the outer offices contained one large window, while the inner offices were surrounded by concrete walls on all four sides. We have placed two monitoring stations in the inner offices and 7 in the outer offices with windows.

The monitoring stations that we used are described in section 3.1. The channels that each monitoring station monitors are shown on Figure 4.1. Each monitoring station was placed between waist and eye level to simulate the location of potential clients in the area and were free of obvious obstructions that could interfere with its reception. After initial setup, the monitoring stations were allowed to run non stop, except for occasional maintenance or configuration change, and a regular reboot at 4 AM everyday. There were also 6 access points in Wean 8th floor. The channels they transmitted at are also shown on Figure 4.1.

The monitoring stations are configured to transmit its gathered data to a server in Wean hall using the wireless network itself. While this adds additional interference and load to the wireless network, they more than made up this deficiency with the easiness of deployment and configuration as we could put them anywhere with a power outlet.

4.2 Analysis

4.2.1 Usage distribution of transmission rates

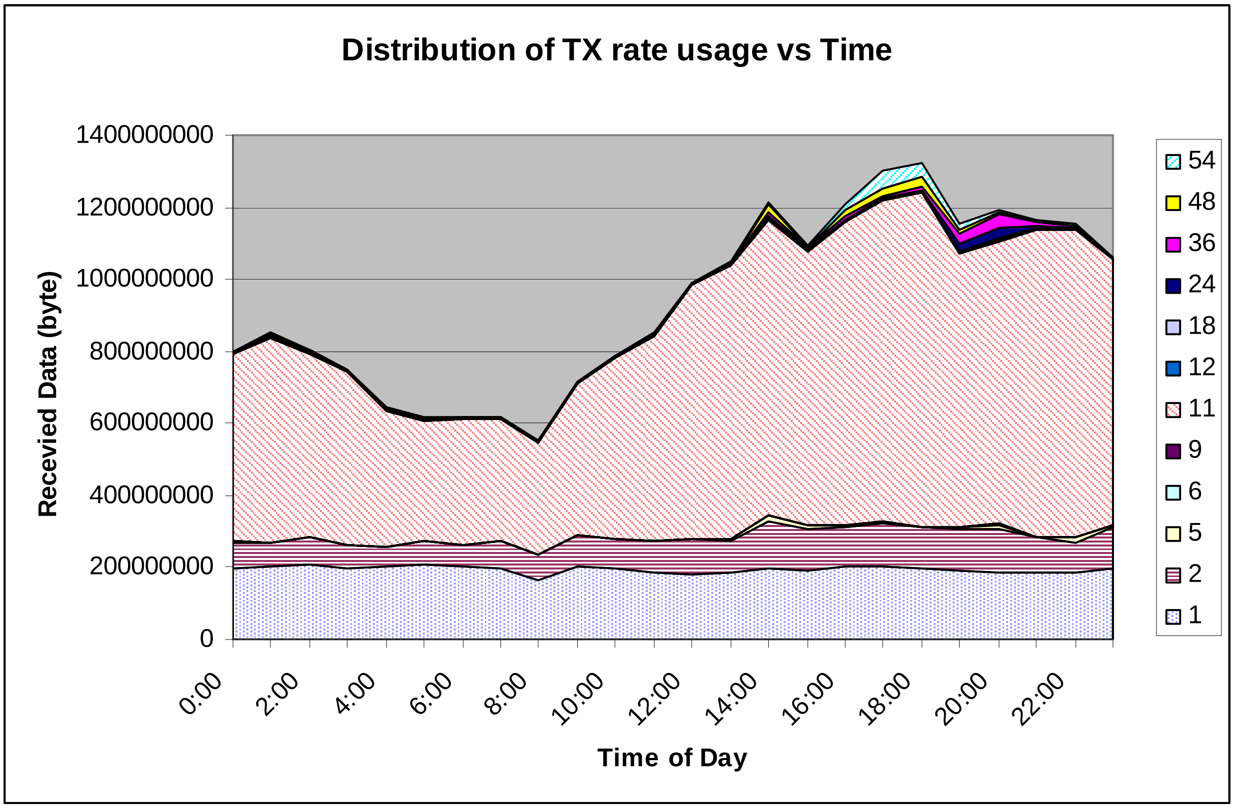

Figure 4.2 Distribution of TX rate vs. time of day from April, 22, 2006

Since all monitoring stations are configured to listen in G mode, they should be able to hear packets with transmission rates of up to 54 Mb/s. Figure 4.2 shows the distribution of the transmission rate on an hourly basis. The most popular transmission rates were 11 Mb/s, 2 Mb/s and 1 Mb/s. Figure 4.2 also shows us that while 54 Mb/s rate was available to users, most of the traffic was still concentrated at legacy 11 Mb/s and 2 Mb/s rates. This could be due to most wireless cards were still using B mode.

In Figure 4.2, we see stable workloads in transmission rate 1 Mb/s and 2 Mb/s. Most of the control packets such as beacons, probe request, RTS/CTS packets were broadcasted in lowest transmission rate. As a result, we saw very steady workload in transmission rate of 1Mb/s and 2 Mb/s.

4.2.2 Window Office VS Inner Offices

	MS1	MS2	Difference between MS1 and MS2	MS6	MS7	Difference between MS6 and MS7
Average Number of Sources	114.45	41.37	73.08	65.58	28.12	37.46

Table 4.2 Average number of terminals detected at different monitoring stations

	MS1	MS2	Difference between MS1 and MS2	MS6	MS7	Difference between MS6 and MS7
Percentage of 1 Mb/s and 2Mb/s packets	39.43	24.87	14.56	50.18	33.38	16.8

Table 4.3 Percentage of low bit rate packets at different monitoring stations

The monitoring station had collected very different data depending on how close the monitoring station is to the window. The MS1 was located within 30cm from the window while the MS2 was located about 2m from the window. The relative distance from the window to the locations of MS1 and MS2 was relatively close, so most of the packet that was received by the MS1 should have reached the MS2 as well. MS1 and MS2 should have detected a similar number of terminals. However, the collected data showed a different view. MS1 detected much more number of sources than MS2. We have analyzed the mean of the average number of the detected source for each of the MS1 and MS2 for the five days. The result is displayed in Table 4.2. MS1 have detected about 73 more terminals than MS2. It seemed that the MS1 received and decoded the packet that was coming though the window while MS2 didn't. It was highly likely that the packets that came through the window could only be received and decoded by the stations which were very close to the window.

The packets which were coming though the windows seemed to have very low transmission rate. The low transmission rate has a higher resistance to interference by having more redundancy in encoding pattern. The packet that came though the window must go through two barriers, one when exiting from other floor or building and other one when entering the different floor or building. Thus these packets needed to be highly resistance to the interference in order for the monitoring station to decode it property. To validate our assumption, we have analyzed average percentage of packets which are sent at 1 Mb/s and 2 Mb/s for five day period. The result was displayed in Table 4.3. The MS1 had about 40% and MS2 had about 25% with difference of 14.56% which was significant. Also, we have analyzed the MS6 and MS7 which are in similar configuration as MS1 and MS2. The analysis from MS6 and MS7 also showed the difference of about 37 in number of detected source and difference of about 16% in percentage of 1 Mb/s and 2 Mb/s packets.

4.2.3 Workload pattern

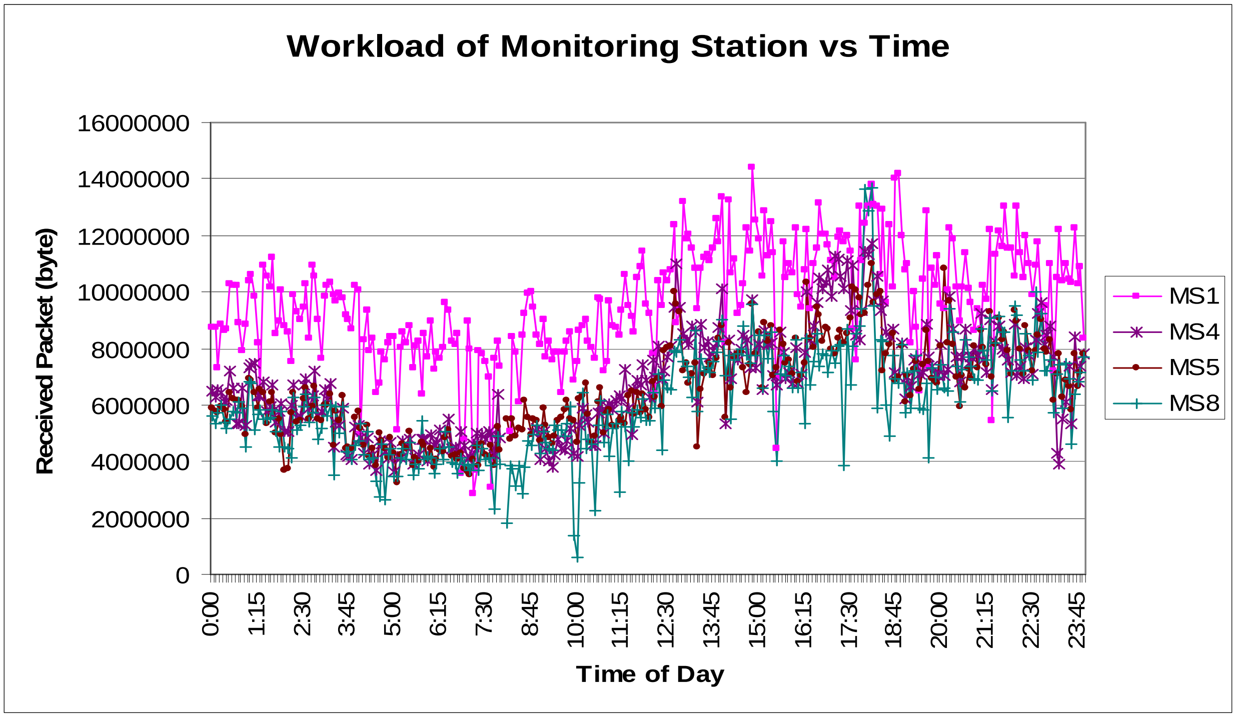

Figure 4.3 Workload data from April, 22, 2006

The Figure 4.3 displays workload of several monitoring station throughout the day. We carefully selected monitoring stations to cover the most of the area. Even though the monitoring stations are located in different area, they seemed to have similar pattern.

The workload was very low from midnight to 8:00 am when the most of the offices are empty and idle. Within this period, most of the mobile wireless devices have exited the wireless network. The most of workload comes from control packets which are broadcasted by the stationary wireless devices, such as access points. From 8:00am to 12:00pm, we saw an increase in workload. As people came into the building with mobile wireless devices and started to use those devices, extra workload in networks was created. During daytime, from 12:00 pm to 9:00 pm, the workload reached its high points. This period was the most active period during day with most of people in their offices. Then the workload starts to drop from 9:00 pm as people started leaving their offices with their mobile wireless devices.

The MS1 seemed to have a very large workload compare to other existing monitoring stations. The reason is that MS1 was very close to the windows and sees much more wireless devices as mentioned in section 4.2.2.

4.2.4 Utilization of Monitoring Station

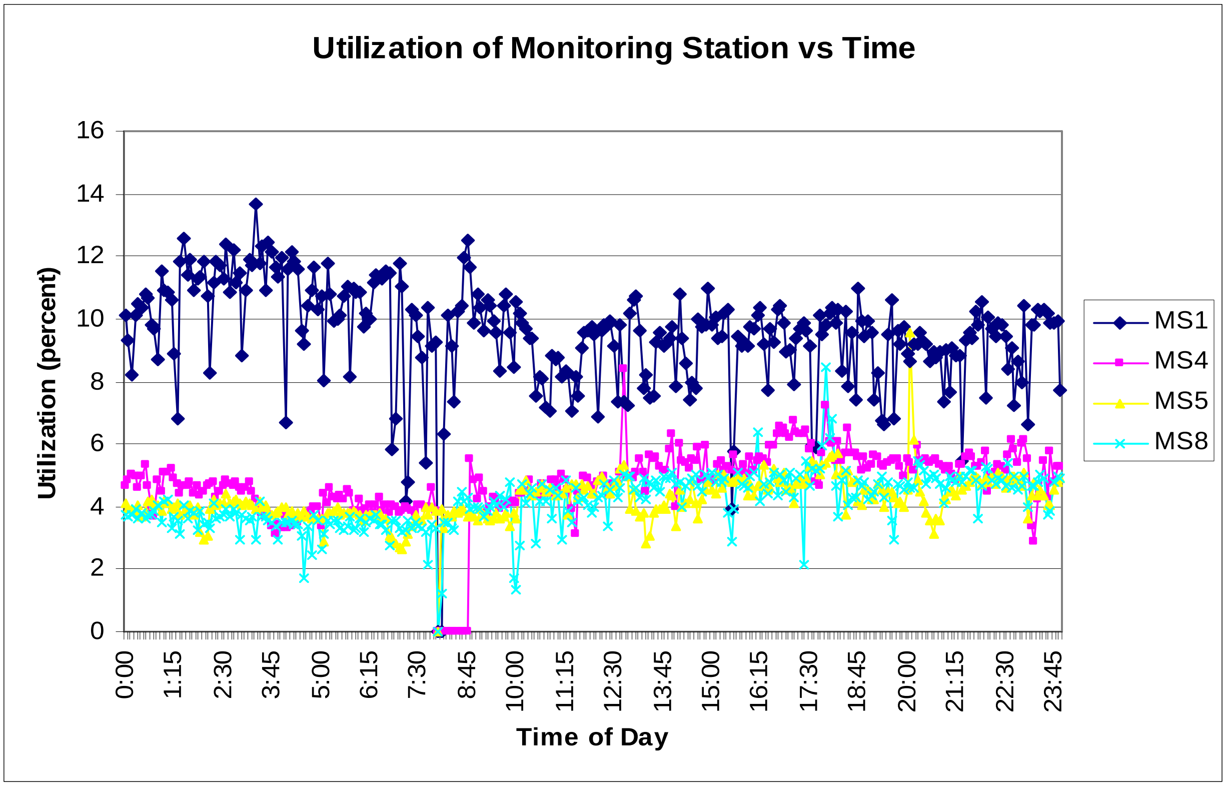

Figure 4.4 Channel utilization data is from April 22, 2006

The utilization of a monitoring system is calculated by following equation:

$$Utilization = \sum_{x} \frac{\text{All Packet Received at Rate } x}{\text{Transmission Rate } x \times \text{Duration of Observation Time}} \quad \text{where } x \text{ is all possible transmission rates}$$

Equation 4.1

Utilization indicates percentage of time that the wireless medium is active around the monitoring station environment. The two important factors in calculating the utilization is the amount of the packet that was received and the transmission rate at which the packets were sent. During the high workload period, the utilization seemed to increase slightly. However, we did not see a clear distinction during high workload and low workload period. The reason why the utilization difference was very small can be found from the transmission rates. Utilization had an inverse relationship with transmission rates. From Figure 4.2, we can clearly see that most of the workload increases came from 11 Mb/s packets. Thus, even though the workload increased significantly, the utilization increase remained small due to the relatively high transmission rate of 11 Mb/s.

MS1, which was located very close to window, received large amount of low transmission rate packets compare to other inner located monitoring stations as we described in section 4.2.2. As a result, the MS1 showed high utilization average compare to other monitoring stations. Surprisingly, MS1 showed a higher utilization in low workload period unlike other monitoring stations. This was possible because of the increase in low transmission rate packets during low workload period. During low workload period, the MS1 was able to receive and decode more low transmission rate packets which came though the window due to less interference in environment. As a result, MS1 had higher utilization in low workload period.

4.2.5 Signal environment at individual monitoring stations

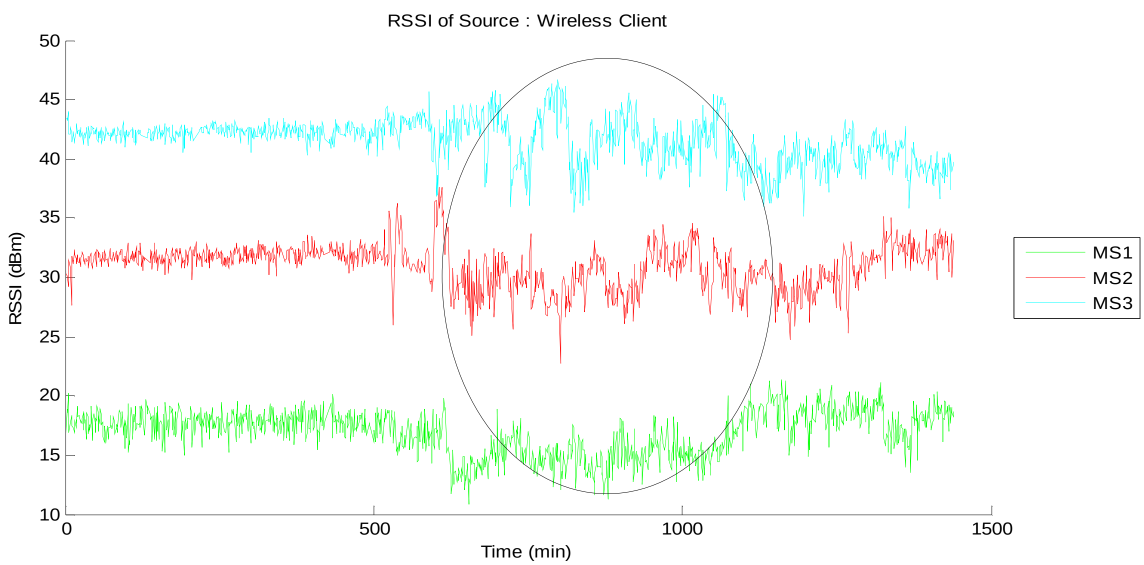

Figure 4.5 Displayed data was collected on March 29, 2006

While previous section focused on global pattern of the wireless environment, it is also very important to analyze how an individual wireless device behaves in the wireless environment. To investigate the individual activity, we have plotted the source based RSSI using an automated MATLAB program.

It was very difficult to characterize specific patterns for the individual wireless device because of large number of the short lived devices. Because we did not have enough data on short lived individuals to see patterns, we focused on individuals which had a long history of activities. One interesting pattern we found was the fluctuation of RSSI during high workload period in some of the clients.

Figure 4.5 displays one of the wireless client's RSSI with fluctuation in high workload period. The circled area in Figure 4.5 corresponds to high utilization and high workload period in Figure 4.3 and Figure 4.4. Congestion resulted from high workload can cause RSSI fluctuation in wireless network by creating more interference. To validate if there was any relationship between RSSI fluctuations with workload, we have analyzed the variance of one specific source in different time period. We have used the source that was used to graph Figure 4.5. The data result is displayed in Table 4.4.

	MS1			MS2			MS3		
	0am-10am	10am-6:30pm	6:30pm-0am	0am-10am	10am-6:30pm	6:30pm-0 am	0am-10am	10am-6:30pm	6:30pm-0am
Day 1	0.49	2.44	1.87	0.76	3.01	4.19	0.57	2.87	1.23
Day 2	0.99	2.26	1.89	0.80	4.47	3.56	0.45	4.88	2.07
Day 3	1.14	4.83	1.55	1.12	2.32	1.72	2.37	3.59	2.43
Day 4	1.39	3.40	1.71	1.27	5.39	0.88	1.07	4.51	2.00
Day 5	3.34	2.45	1.60	1.16	2.12	0.58	3.43	2.31	1.86
Day 6	0.82	2.12	0.89	0.29	0.86	0.45	0.57	1.52	0.85
Avg	1.36	2.92	1.58	0.90	3.03	1.90	1.41	3.28	1.74

Table 4.4 Variance in RSSI during different time periods of one specific source

In all monitoring station, the RSSI variance reaches its highest value during 10am to 6:30pm period except Day 5. Also, the average variance of RSSI in this period was almost double of the variance in other periods. These results imply that this specific client experienced the RSSI fluctuation continuously on high workload period. However, if high workload actually caused this fluctuation, other wireless devices present should have experienced the fluctuation in RSSI as well. To validate our assumption, we have calculated the variance of RSSI of all sources for 6 days in three different time durations as well. The results are shown in Table 4.5.

	Midnight–10 am Low workload	10 am – 6:30 pm High workload	6:30 pm – Midnight Middle workload
Day 1	1.09639292	2.965807111	2.701974414
Day 2	0.50137857	0.536210981	0.459565255
Day 3	1.307616281	10.2084206	2.398461648
Day 4	4.04990127	6.081151032	0.475076917
Day 5	0.487673156	0.284420617	0.246989772
Day 6	0.491255362	7.14644341	0.41531664
Average	1.322369593	4.537075625	1.116230774

Table 4.5 Variance in RSSI during different time periods of multiple sources

Except Day 5, we saw the most variance in RSSI between 10 am to 6:30pm, which indicated that RSSI fluctuated greatly during this period for majority of sources. Also, the variance between Midnight and 10am and between 6:30pm and Midnight were low, mostly under 2.5. It indicated that most of clients during high workload period experienced greater fluctuation in RSSI, which supports our assumption.

4.2.6 Distribution of channel utilization

Next we tried to use the channel utilization applet to determine the location of any hot spots of activity in the wireless deployment. If there was, it may offer some clues on the locations of future access points.

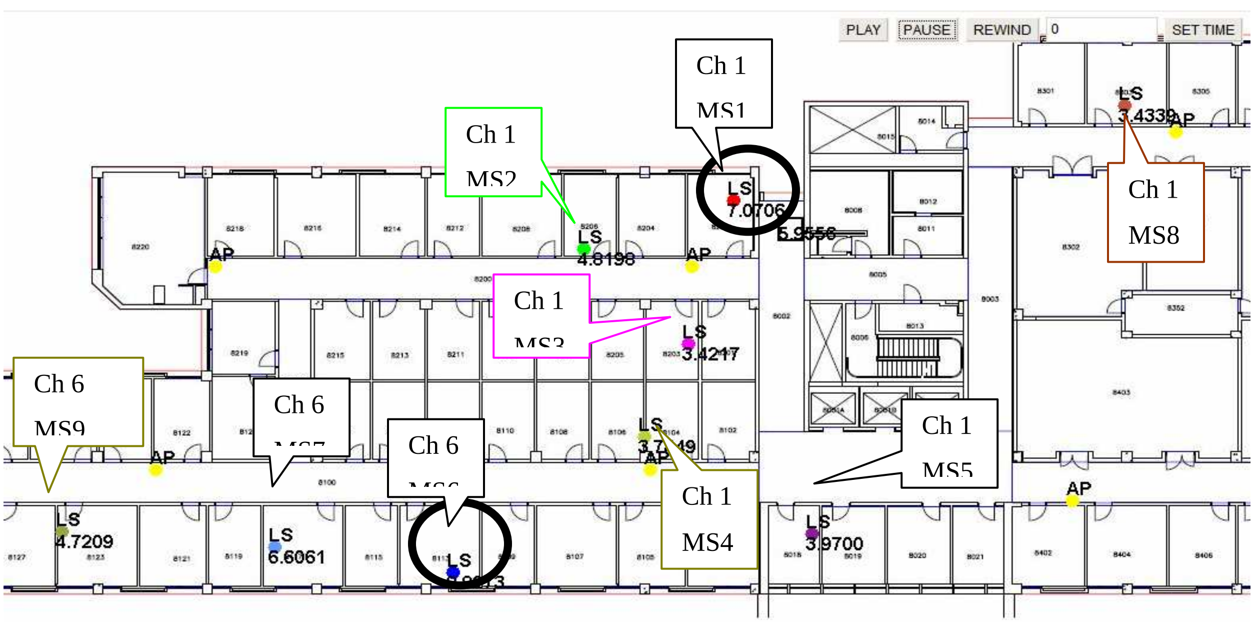

Figure 4.6 Screenshot of the channel utilization visualization applet

As Figure 4.6 shows, areas around MS1 and MS6 had the highest channel utilization of 7% and 9% respectively. As mentioned in section 4.2.2, there is a correlation between window offices and high channel usage. In addition to the increased number of terminals observed by monitoring stations in window offices, window monitoring stations also observe much higher channel utilization than non window offices. This effect can be seen by the 9% channel utilization in MS6 VS the 3.4% channel utilization in MS3. One also observed from this figure that while MS 7 and MS 9 were also in window offices, their channel utilization of 4.7% and 6.6% was much lower than the channel utilization observed by MS 6. This is because while both MS 7 and MS 9 were in window offices, these monitoring stations were placed about 2 to 3 meters away from the window, while MS 1 and MS 6 were placed within 1 meter of the window.

4.2.7 Channel mapping

One of the goals of the network visualization applet was to provide an integrated view of the signal environment from multiple monitoring stations to the network administrator. Figure 4.7 showed one view of this with how strongly was a transmission received at each monitoring stations. In the figure, both monitoring stations and access points were the sources. Next to each source was a list of the RSSI received at each different monitoring station.

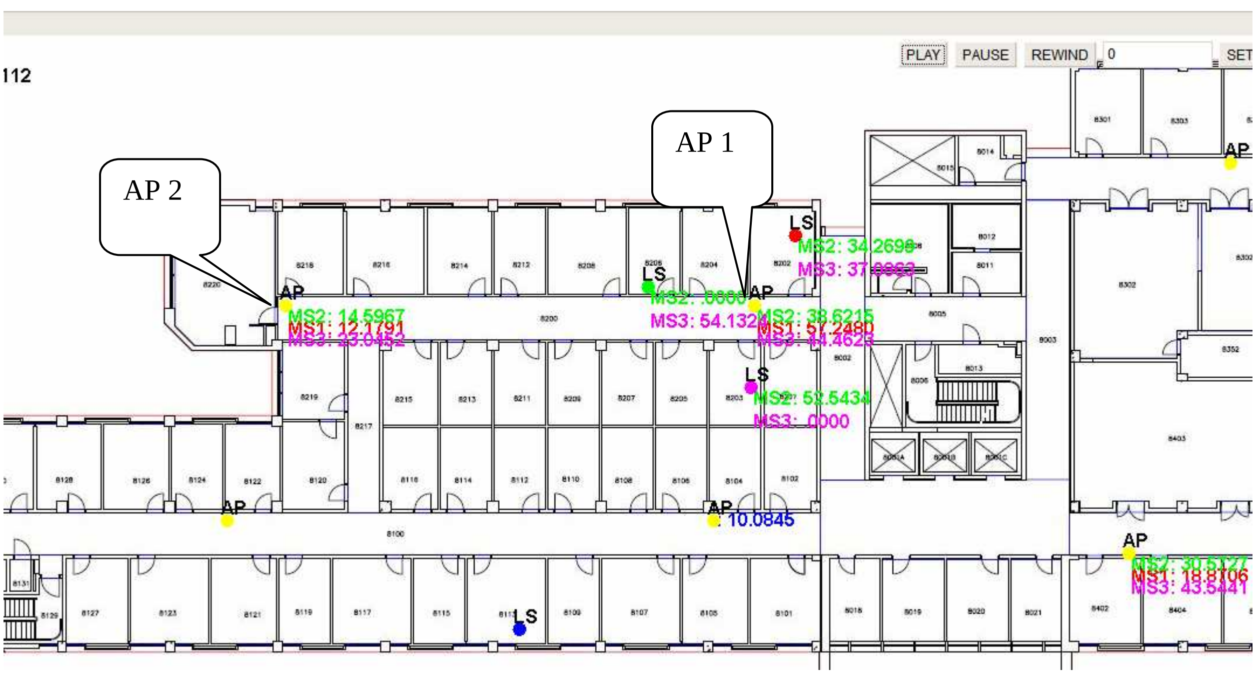

Figure 4.7 Screenshot of the network visualization applet

As you can see, the signal attenuates very differently even in areas that were relatively close to each other. For example, let's look at the access point AP 2 in Figure 4.7. Its signal was received at 23 dBm at MS3 and 14.5 dBm at MS2. While MS2 was closer to the highlighted access point than MS3, the RSSI at MS2 was greater by 9 dBm. This confirmed that there was no strong relationship between RSSI and distance.

The other result that can be gleaned from this figure is that even in relatively small areas, the RSSI varies greatly. For example, in Figure 4.7 the signal from AP 1 was received at 38 dBm at MS2, 57 dBm at MS 1 and 44 dBm at MS 3. While MS 1, MS 2 and MS 3 were all about the same distance to AP 1 and were very close to AP 1, the received signal strength differs by as much as 10 dBm. Because of the variation seen even in relatively small areas, we may need to deploy the monitoring stations more densely to capture this variation.

4.2.8 Hidden terminals

Wireless networks rely on Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA) to control access to the communication channel. During the start of every data transmission, the transmitter will first sense if there are any jam signals from another transmitter. If there is, it will wait for a random amount of time and try again. If there is no jam signals, it will transmit a jam signal and then a data frame.

Hidden terminals are transmitters that are out of carrier sense range from each other, but are trying to transmit to a single receiver, usually an access point. So while they cannot sense any jamming signal in the channel before they start transmission, their transmissions may arrive at the same time at the access point. If their signals are strong enough, one or both of their transmissions may not be decoded correctly by the access point.

Since our monitoring software does not run on the clients and access points, we are unable to detect exact collisions at the access points. However we can find potential hidden terminals by assuming clients and access points close to a monitor station will have the same carrier sense range as the monitor station. Figure 4.8 illustrates our approach to finding potential hidden terminals. The three circles represent the carrier sense range of three monitoring stations. We first substituted what MS1 sees in the network for what the access point would see because MS1 was located close to the access point. Next we assumed all terminals within the carrier sense range of MS2 to share the network view of MS2, and all terminals within the carrier sense range of MS3 to share the network view of MS 3. This means we assumed the terminals within the carrier sense range of a monitoring station share the same carrier sense range of the that monitoring station. With this in mind, potential hidden terminals are transmitters associated with the same access point, but are out of carrier sense range of each other.

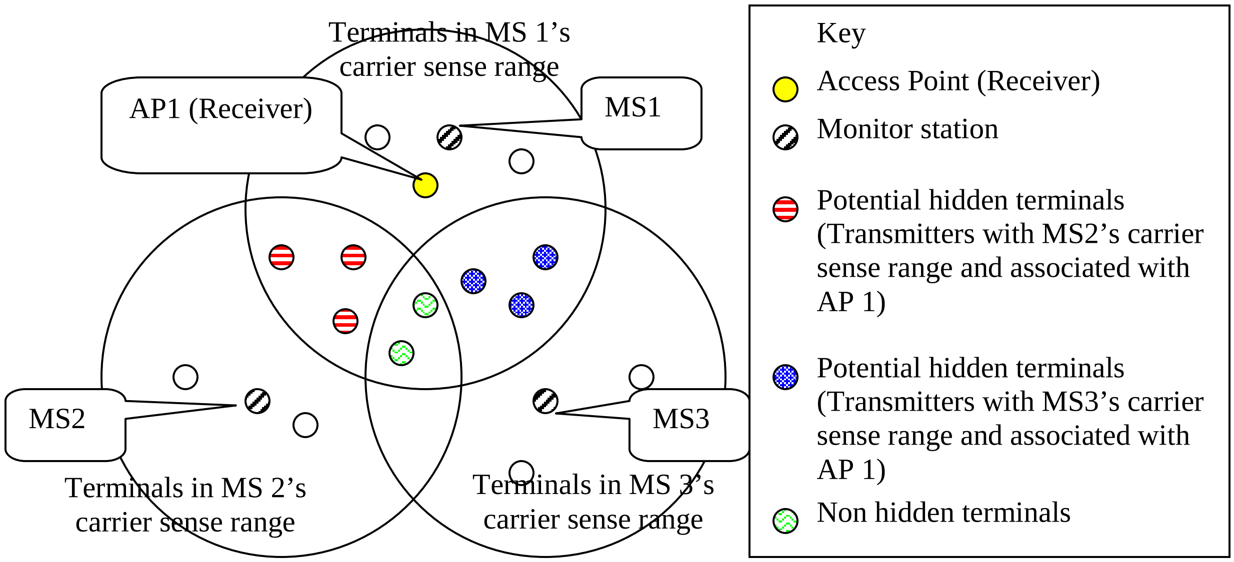

Figure 4.8 Definition of hidden terminals

In addition to being out of carrier sense range, a terminal was only considered hidden if its signal strength was high enough to disrupt communication at the access point. However since we do not know what at RSSI will the communication be disrupted, we have tabulated the average ratio of hidden terminals to non hidden terminals at various RSSI ratio thresholds at Table 4.6. RSSI ratio is the ratio of RSSI of potential hidden terminal transmissions to the RSSI of non hidden terminal transmissions.

	Ratio of hidden terminals to non hidden terminals			
RSSI Ratio Threshold	0.4	0.5	0.6	0.7
	0.122575	0.105565	0.089725	0.073991

Table 4.6 Ratio of hidden terminals to non-hidden terminals at MS2 and MS3

So if we assume a RSSI ratio of 0.4 will cause transmission decoding errors, then on average, 12% of the terminals were hidden. If we set the higher ratio for transmission decoding errors at 0.7, then on average 7% of the terminals were hidden. This shows that while MS 2 and MS 3 are relatively close to MS 1, at least 7% of the terminals detected by MS 1 were potentially hidden from each other. While this number may seem very high for hidden nodes, please keep in mind that these were only potential hidden terminals.

Next we decided to look at terminals that are further away from each other and from the access point so see if there are more hidden terminals at longer distances. We decided to look at MS2 and MS4 because both monitoring station were still monitoring in channel 1 and were relatively far from each other. The results were tabulated in Table 4.7 Ratio of hidden terminals to non-hidden terminals at MS2 and MS4. As you can see, on average at least 15% of terminals were potentially hidden from each other compared to the 7% potentially hidden terminals at a closer distance to the access point.

	Ratio of hidden terminals to non hidden terminals			
RSSI Ratio Threshold	0.4	0.5	0.6	0.7
	0.182587	0.170467	0.157604	0.146388

Table 4.7 Ratio of hidden terminals to non-hidden terminals at MS2 and MS4

As the distance between terminals increase, it was more likely that these two terminals will be hidden relative to each other because it was more likely they will fall within each other carrier sense range. Since we are substituting the carrier sense range of different monitoring stations as the carrier sense range for these terminals, the fact that there was an increase in the number of hidden nodes detected when we increased the distance between monitoring stations was expected.

To determine where these potential hidden terminals were located, we used the hidden terminal visualization applet. A screenshot of the applet showed two areas where potential hidden terminals are commonly located is at Figure 4.9. In the screenshot, one can see there were six dots labeled AP. These represent the access points. The dots labeled with MS are the monitoring stations. The boxes in the screenshot represented the location of the hidden terminals that were detected using the algorithm described earlier in this section. As one can see, there were two locations in the 8^{th} floor of Wean were likely to contain potential hidden terminals. These two locations are very far apart from each other, so the fact that they cannot detect the transmissions from each other was very likely. However, the terminals in both of these locations were transmitting in channel 1 and therefore were connected to one of the two access points labeled on the screenshot. If they were actually associated with the access point in the middle, then it was very possible for their transmissions to interfere with each other.

Figure 4.9 Screenshot of the hidden terminal visualization applet

4.2.9 Exposed terminals

Exposed terminals are transmitters that are within carrier sense range of each other but their recipients are not within carrier sense ranges of each other. So due to carrier sense, exposed terminals cannot simultaneous transmit to their respective receivers, when in fact they could without causing interference at the receivers. Figure 4.10 illustrates this definition, where S1 and S2 are the transmitters and R1 and R2 are the receivers. S1 and S2 are within carrier sense range of each other, but the receivers R1 and R2 are not.

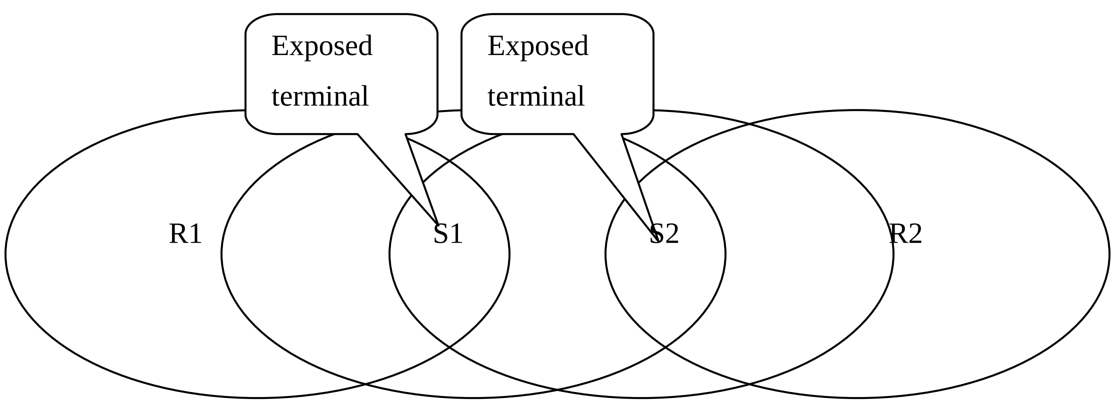

Figure 4.10 Definition of exposed terminals

We have adopted a simplified definition of the exposed terminal. So instead of using two sources and two receivers to define exposed terminals, we just look at two sources and one receiver, S1, S2 and R1. If S1 and S2 are within the carrier sense range of each other but only S1 transmission can reach R1, then S2 is an exposed terminal since S2 must be associated with another receiver. If we assume that terminals within the carrier sense range of a monitoring station will share the network view of that monitoring station, and we set MS 1 to act as the access point, then terminals outside of the carrier sense range of MS 1 but within the carrier sense range of MS 2 will be exposed relative to the terminals that is within the carrier sense range of both MS 1 and MS 2. This is illustrated with Figure 4.11. So exposed terminals are transmitters that are not associated with AP1 but are within the carrier sense range of transmitters that are associated with AP1

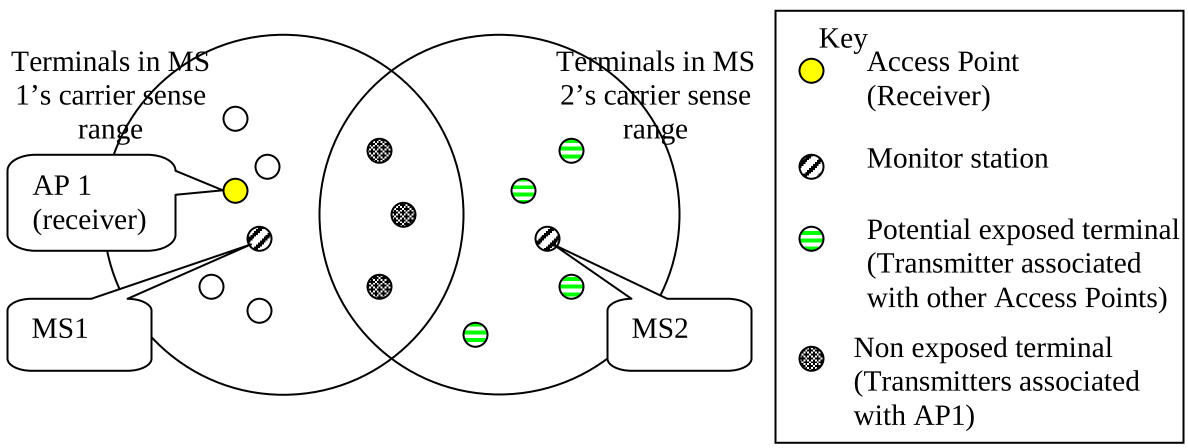

Figure 4.11 Definition of potential exposed terminals

We have graphed the ratio of the utilization of exposed terminals to non exposed terminals in Figure 4.12 through Figure 4.14. Utilization was calculated using Equation 4.1 in section 4.2.4. In Figure 4.12, we graphed the utilization ratio of the exposed nodes to non-exposed nodes at a monitoring station MS2. Since both MS1 and MS2 were on the window side and were relatively close to each other, their network views were very similar and therefore, MS2 had very little exposed terminals relative to MS1. In Figure 4.13, MS3 was further away from MS1 and also was located in the inner offices. So more terminals in MS3 may be far enough from the access point in MS1 such that they would not be affected by transmission from the non-hidden nodes to the access point in MS1. Therefore, the number of exposed terminals in MS3 increased greatly when compared to MS2.

MS4 was also located in an inner office but even further away from MS 1. This results in a little increase in the number of exposed terminals when compared to MS3. In Figure 4.13 and Figure 4.14, there were times where the utilization ratio reached close to 1. This was because at those times there was no activity in the non-exposed terminals. So in effect, most of the terminals at those times were "exposed". But because at those times there were very little non-exposed terminals to be affected by the exposed terminals, we could safely discard the utilization ratios at those times.

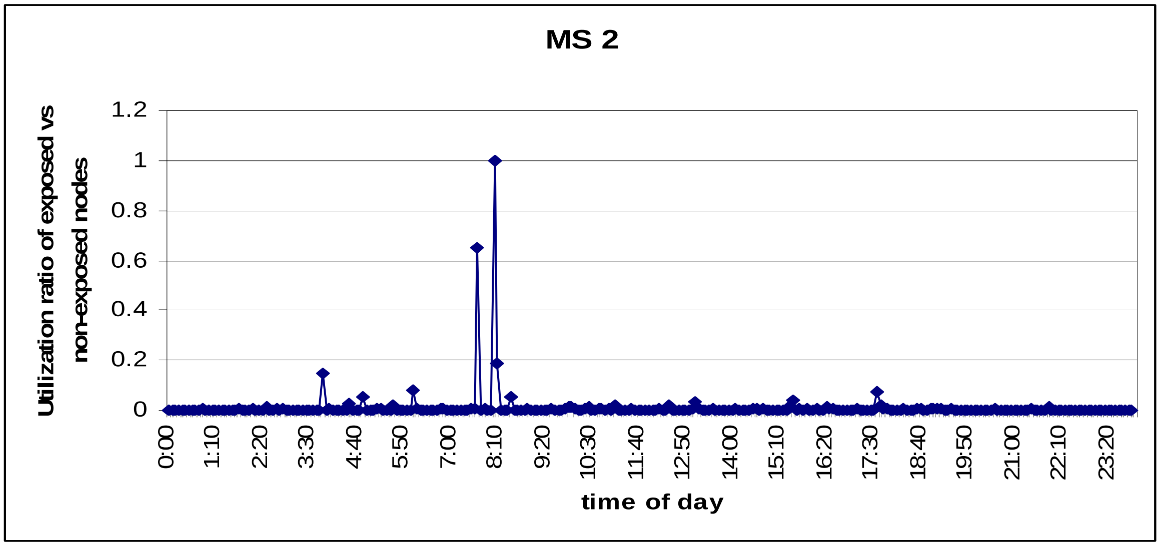

Figure 4.12 Utilization ratio of exposed terminals vs. all terminals at MS2

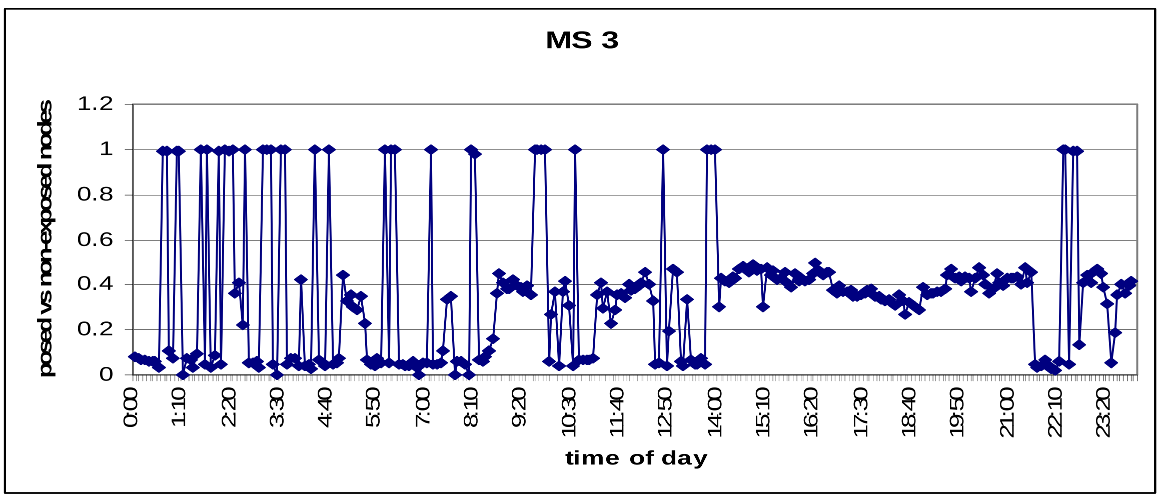

Figure 4.13 Utilization ratio of exposed terminals vs. all terminals at MS3

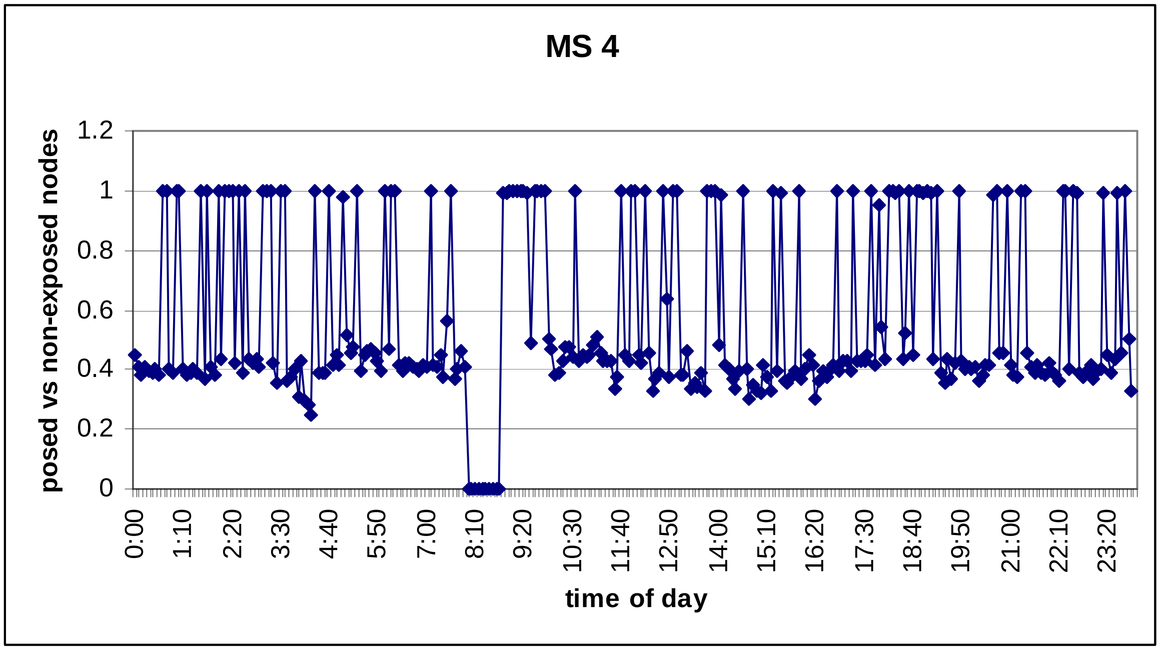

Figure 4.14 Utilization ratio of exposed terminals vs. all terminals at MS4

When compared to the number of hidden terminals, the number of exposed terminals was greater. For example, if we look at the area around MS3, about 8% to 12% of the terminals were potential hidden terminals, while 34% of the terminals were potential exposed terminal. This effect was even more dramatic at the area around MS4 which was further away from the access point. The data showed 14% to 18% of the terminals were potential hidden terminals while about 44% of the terminals were potential exposed terminals. This was expected because in an office environment that was relatively small in area but contains many access points, terminals that were connected to different access points have a very high chance of being in the same carrier sense range of each other. In addition to the small area of the deployment, there were many offices with windows where signals from terminals outside of the wireless network could enter, and could make the terminals inside the deployment believe the channel was busy when in fact it was not.

Chapter 5
Conclusion

The primary motivation of this thesis was to develop a real time signal mapping system and use it to diagnose potential problems in the wireless network. With the exception of the implementing a spatial application to take advantage of our system, we have achieved all aims outlined in section 1.4.

5.1 Achievements

We have developed and deployed a system that provides a real time wireless site survey for a wireless network deployment. Our system have solved the challenges to deploying a real time wireless site survey described in the Section 2.3 with a deployment of cheap monitoring stations strategically placed in the wireless network environment. These monitoring stations were made from commodity wireless routers and controlled open source software.

We have also created a set of visualizations that integrates the local signal maps gathered by individual monitoring stations into a comprehensive signal map of the wireless network. By unifying the signal environment data from multiple monitoring stations, we can present a relatively complete picture of the signal environment at the deployment site in real time. This way, a wireless network administrator may spot trends or problems with the wireless network at one glance, rather than relying on offline processing of the collected data.

With our system we have also characterized the wireless network here at CMU. First, we found that while most access points support operating in G mode, most packets were sent out in B mode data rates. Second, we discovered that a significant portion of the wireless traffic was made up of control packets such as beacons, probe request, RTS/CTS packets were broadcasted in lowest transmission rate. Third, we found that there were significantly more interference and traffic at areas near windows. Fourth, we found while the workload varies with the daily patterns of people, the channel utilization does not. Finally, we have also made an interesting correlation between the variance of RSSI and the channel workload. We discovered greatly during periods of high activity, between 10 am to 6:30pm, and remained stable during periods of low activity.

To evaluate the usefulness of our system, we performed a more in depth analysis of the signal data to determine the prevalence of hidden and exposed terminals in the wireless network.

5.2 Hindsight and Future Direction

In hindsight, we should have not concentrated so much of our time on developing a dynamic power control spatial application. While we hoped that dynamic transmission power control would added to the Madwifi drivers, we should have not depended on other people to implement this functionality. We also should have deployed our system on a site that is more active or on different floors to see the differences in the signal map at different environments.

In additional new insights, our thesis also brought new questions. One question would be why was there such a drastic difference in traffic pattern in two different window offices on the same side of the building. Could this really be due to the difference in 1 meter in the placement of the monitoring station from the window? We would like to answer this question by more densely deploy the monitoring stations at this trouble spots. Another question would be the reason behind the apparent high correlation between the variance of RSSI and network workload. We would like to perform some controlled experiments to determine whether network workload actually affects RSSI.

In addition to answering these interesting questions, another future step we would like to take would be developing spatial applications to take advantage of our real time signal monitoring in the following areas:

- **Configuration**: Selecting the least congested access point rather than the access point with the strongest signal by querying about the network congestion in a specific area.
- **Interference reduction**: Reduce interference by querying for the exact power needed to communicate with an access point or a client.

Bibliography

[1] Code of Federal Regulations, Title 47, Volume 1. Revised as of October 1, 2004.

[2] Boithais, Lucien, *Radio Wave Propagation*. McGraw-Hill Book Company, New York. 1987.

[3] Atul Adya, Paramvir Bahl, Ranveer Chandra, Lili Qiu. Architecture and Techniques for Diagnosing Faults in IEEE 802.11 Infrastructure Networks. In *Proceedings of the MobiCom Conference*, Philadelphia, PA, September 2004.

[4] Konstantina Papagiannaki, Mark Yarvis, W. Steven Conner. Experimental Characterization of Home Wireless Networks and Design Implications. *Proceedings of the Twenty-Fifth Annual IEEE Conference on Computer Communications (INFOCOM 2006)*, Barcelona, Spain, April 2006.

[5] http://www.netgear.com/products/details/WGT634U.php

[6] Roofnet [MIT Roofnet], http://pdos.lcs.mit.edu/roofnet

[7] Madwifi, http://sourceforge.net/ projects/madwifi

[8] Atheros Communications, http://www.atheros.com

[9] Eddie Kohler, Robert Morris, Benjie Chen, John Jannotti, and M. Frans Kaashoek. The Click modular router. ACM Transactions on Computer Systems,18(4), November 2000.

[10] http://www.w3.org/TR/soap12-part1

[11] http://www.boa.org/

[12] http://expat.sourceforge.net/

[13] http://www.postgresql.org/

[14] JAVA Graphics2D API, http://java.sun.com/j2se/1.4.2/docs/api/java/awt/Graphics2D.html

[15] Bahl, P., AND Padmanabhan, V. N. RADAR: An In-Building RF-based User Location and Tracking System. In *IEEE Infocom 2000* (March 2000), vol. 2, pp. 775–784.

[16] Bahl, P., Padmanabhan, V. N., AND Balachandran, A. Enhancements to the RADAR User Location and Tracking System. Tech. Rep. MSR-TR-00-12, Microsoft Research, February 2000.