Throughput Measurements and Empirical Prediction Models for IEEE 802.11b Wireless LAN (WLAN) Installations

Benjamin E. Henty

Thesis submitted to the Faculty of the
Virginia Polytechnic Institute and State University
in partial fulfillment of the requirements for the degree of

Master of Science
in
Electrical Engineering

Dr. Theodore S. Rappaport, Chair
Dr. Ioannis M. Besieris
Dr. Scott F. Midkiff
Dr. William H. Tranter

August 9, 2001
Blacksburg, Virginia

Keywords: wireless LAN, IEEE 802.11b, throughput measurement, throughput prediction, RSSI measurement

Throughput Measurements and Empirical Prediction Models for IEEE 802.11b Wireless LAN (WLAN) Installations

Benjamin E. Henty

(ABSTRACT)

Typically a wireless LAN infrastructure is designed and installed by Networking Professionals. These individuals are extremely familiar with wired networks, but are often unfamiliar with wireless networks. Thus, Wireless LAN installations are currently handicapped by the lack of an accurate, performance prediction model that is intuitive for use by non-wireless professionals.

To provide a solution to this problem, this thesis presents a method of predicting the expected wireless LAN throughput using a site-specific model of an indoor environment. In order to develop this throughput prediction model, two wireless LAN throughput measurement products, LANFielder and SiteSpy, were created. These two products, which are patent pending, allow site-specific network performance measurements to be made. These two software packages were used to conduct an extensive measurement campaign to evaluate the performance of two IEEE 802.11b access points (APs) under ideal, multi-user, and interference scenarios. The data from this measurement campaign was then used to create empirically based throughput prediction models. The resulting models were first developed using RSSI measurements and then confirmed using predicted signal strength parameters.

This work was supported through the Bradley Fellowship, the Bradley Industrial Fellowship, and through donations from ACDC, Inc. and Wireless Valley Communications, Inc.

The author would like to acknowledge the following people for assisting in making this work possible:

Wireless Valley for the use of software and equipment. MPRG for the support, encouragement, equipment and supplies. Jiun Siew for his hard work and assistance in measurements and modeling. Max Roberts for conversations and reflections about this work. Roger Skidmore for software and additional support. Charles Lepple for his advice and assistance. Ros Neplokh for sharing his near god-like competence in configuring networks. Sam Abtek for listening. And of course, my friends and family for their continued love and support.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Wireless Data Networks

Wireless LANs (Local Area Networks) are typically installed and maintained by Information Technology (IT) staff members. These types of individuals do not generally possess detailed knowledge of wireless network performance or even a basic idea of the coverage area they can expect to achieve with a wireless LAN network. This is especially difficult because wireless networks can often experience an unexpected RF (radio frequency) environment when they are used indoors. For this reason this thesis has developed simple, accurate models of wireless LAN throughput based on indoor RF propagation. The use of these models can be used to easily predict the throughput experienced by a user at any location in the coverage area of a wireless LAN access point. This thesis also presents measurement campaign results of actual wireless LANs and the methods used to calculate the parameters of the empirical throughput prediction models.

1.2 Motivation for Throughput Prediction

There are several solutions for indoor wireless networking available today. A large number of these of products are based on the IEEE 802.11b wireless LAN standard [30]. These hardware devices are typically PC Card, PCI (Peripheral Component Interconnect), or USB (Universal Serial Bus) based based network interface cards (NICs) which can communicate with one another or through Access Points (APs) to a wired backbone network. Wireless LAN equipment based on the 802.11b standard operate between 2.4 and 2.497 GHz where they can be used in most countries worldwide. The APs are typically installed by network savvy IT staff who know little or nothing about wireless communications. Thus, the ability to simulate a wireless LAN design and predict throughput in a computer-aided design environment would be a valuable contribution to the wireless LAN industry.

IT staff members are often interested in the network performance that will be experienced by a wireless LAN user. This performance varies with the type of wireless LAN card in use, the location of the wireless LAN user, as well as a number of other factors. For instance, the 802.11b products are capable of a shared, raw data bandwidth of 11 Mbps while the 802.11 products are capable of a shared, raw data bandwidth of 2 Mbps. However, this bandwidth is not usually realized because a single NIC is not capable of utilizing the full bandwidth. Further, the available data rate fluctuates depending on the wireless channel characteristics. For this reason, it is important to be able to both accurately measure and reliably predict the actual throughput of IEEE 802.11 devices. This thesis shows that by combining measurements and prediction models, IT staff members can easily plan and deploy a wireless LAN network.

1.3 Thesis Overview

This thesis focuses on the difficulties faced in designing and deploying a wireless LAN network through developing throughput prediction models to aid in designing a wireless LAN infrastructure. To create these models, the thesis first presents the commercially available measurement products SiteSpy and LANFielder, which were created as part of the work for this thesis. These products are designed to measure wireless LAN throughput and other network performance criteria and record the results in a precise site-specific manner. This ability is key to accurate measurements of wireless LAN networks. These measurement tools are presented in Chapter 3.

To verify the accurate workings of LANFielder and SiteSpy, and to collect data for the development of empirically based prediction models, an extensive measurement campaign was conducted using SiteSpy, LANFielder and some custom RSSI (Received Signal

Strength Intensity) measurement software, in conjunction with commercially available IEEE 802.11b wireless modems, in a real network setting. The details of the measurement campaign and the results of the campaign are presented in Chapters 4 and 5, respectively.

Once the measurement campaigns were complete, models were developed using the measurements for two different test-bed IEEE 802.11b wireless LAN systems. The conceptual framework, development and accuracy of these models is presented in Chapter 6. The models are further validated in Chapter 7.

The bulk of this thesis assumes some basic familiarity with the IEEE 802.11 and IEEE 802.11b standards. However, to aid the reader who is somewhat unfamiliar with the standards, a basic tutorial is presented in Appendix A. Further, a review of basic wireless networking terms and recent research is presented in Chapter 2. Together, all of the information presented in this thesis provides an important contribution to the design of IEEE 802.11b wireless LANs.

Chapter 2

Prior Research

2.1 Wireless LAN Propagation Research

2.1.1 RF Coverage Measurement Techniques

Several authors have performed research in the area of measuring and predicting the RF coverage area of indoor wireless LAN access points. Three of the more popular techniques described in the literature to measure the path loss of wireless LAN signals are the use of continuous wave (CW) transmitters with power meter receivers, Broadband Pulse Channel Sounding, and using the signal strength percentage reported by wireless LAN cards. Some of the research which uses each of these techniques is detailed in Table 2.1 and later in Section 2.1.2. The following sections detail the benefits and shortcomings of each of these measurement techniques.

Continuous Wave (CW) Transmitter and Power Meter

The continuous wave (CW) transmission technique has long been used to measure the basic path loss of a wireless channel. Many useful statistics, such as the average fade duration, level crossing rate and the probability density function of the received signal strength, can be quantified and used in the design of a wireless modem. This narrowband technique requires less measurement hardware and measurement campaigns can be carried out more quickly and easily because a minimum amount of equipment is needed.

Table 2.1: Overview of Different RF Propagation Measurement Techniques Used in Wireless LAN Coverage Research

Measurement Technique	Prediction Technique	References
CW Signal and Power Meter	Combination of direct path loss and partition losses with ray-tracing and refraction techniques	[28]
CW Signal and Power Meter	Various exponential and partition loss based empirical models	[7]
Pulse Transmission	Exponential path loss exponent and partition losses along direct path	[11]
Network Analyzer	None	[14]
Wireless LAN Card Signal Strength in dBm	None	[23]
Wireless LAN Card Signal Strength Percentage	Empirically fit path loss to a quadratic function and used a neural network type analysis to model multipath effects	[4]
Wireless LAN Card Signal Strength Percentage, converted to dBm	Empirically fit outdoor path loss measurements to Okumura model	[9]
Wireless LAN Card Signal Strength Percentage	None	[5], [35]
Wireless LAN Card Signal Strength Percentage, converted to dBm	None	[17]

The main difficulty with this technique is that it is a narrowband measurement technique. If the bandwidth of a wireless modem is wider than the coherence bandwidth of the wireless channel then the instantaneous continuous wave path loss measurements of the signal strength will not be valid, even if local area averaged. This is especially a problem for IEEE 802.11 wireless LAN cards which have an RF bandwidth of 22 MHz [30] but operate in RF environments with typical coherence bandwidths of between 5 and 46 MHz. [11].

Broadband Pulse Transmission

Broadband Pulse Transmission overcomes many of the shortcomings of CW transmission at the expense of increased measurement cost and time. The technique uses a broadband antenna to transmit an impulse which is short in duration. By measuring the pulses and arrival times at receive locations, the exact channel impulse response can be measured. The bandwidth of the measured channel is determined by the duration of the original transmitted impulse. By measuring a bandwidth larger than the bandwidth of the wireless modem, the RF performance of a wireless modem can be exactly measured even if the channel's coherence bandwidth is smaller than the bandwidth of the wireless modem. This technique also allows the quantification of the same statistics as the CW measurement technique as well as the mean excess delay, delay spread, coherence time and coherence bandwidth of the measured channel. These statistics are extremely useful for designing wireless modems. The difficulty is that full broadband pulse systems require a large amount of equipment and time to measure.

It should also be noted that other more advanced channel sounding techniques are available with similar results, benefits and limitations, including sliding correlator, network analyzers [14] and chirp channel sounders which provide broadband channel impulse response measurements. The use of this equipment has so far not been broadly reported in the literature.

Wireless LAN Card Reported Signal Strength

The third major wireless LAN channel measurement technique used in the literature and in practice is to use the reported signal strength percentage reported by software utilities which are provided with wireless LAN cards. In this technique a laptop or other computer is used to measure the received signal strength from another wireless LAN card or a wireless LAN access point. This technique is far simpler than the previously mentioned techniques and allows the comparison of different cards from different wireless LAN vendors. In addition, measurements can be taken quickly and easily, and the results are more easily applied to actual installations and coverage prediction. However, until recently, the signal strength reported by wireless LAN cards has just been a percentage value between 0 and 100. These numbers do not have any direct, known value or meaningful relationship to path loss. They allow relative comparisons of different locations, but are not practical for actual path loss measurement purposes. These "magic numbers" are hard to apply or compare to prediction models. As a result, research using this method is typically only applicable to a single type of wireless LAN card or must be calibrated to convert the values into actual units of power. To overcome these shortcomings, IEEE 802.11b vendors, such as WaveLAN and 3Com, have recently included software tools that report the actual dBm value of the received signal strength.

2.1.2 Wireless LAN Related RF Propagation Research

This section presents some of the progress that has been made in the research on wireless LAN propagation measurement, prediction and modeling. While each of these papers present important contributions to the field, all are lacking a simple, accurate model of wireless LAN propagation based on a diverse array of measurements.

Indoor Propagation Measurements

In [11] the authors have measured the channel impulse response characteristics at 2.4 GHz using a generated impulse in two different buildings. The authors used the measured impulse responses to find the parameters of the Motley-Keenan/Seidel-Rappaport indoor propagation model ([25] and [18]). The authors used a path loss reference distance of 10λ or 1.25 meters and calculated a path loss exponent of n=3.086 which resulted in a standard deviation of 5.84 dB. In addition, the authors measured and reported the mean excess delay, delay spread and coherence bandwidth for each of the measurement cases. Thus, the authors have provided a reliable prediction model for wireless LAN received signal strength based on measured data. However, the measured data was measured using pulsed transmission and laboratory equipment. It remains to be seen if the model extrapolates to RSSI measurements made directly with wireless LAN cards that use Direct Sequence Spread Spectrum (DSSS) techniques to spread the transmitted power over a broad bandwidth. In addition, the environments measured in [11] are limited in extent. Additional validation of the authors' model is warranted.

Pre-IEEE 802.11 Wireless LAN Measurements

The authors of [17] have published a survey of different measurements that demonstrate the performance of a wireless LAN test-bed. The modems used in the test-bed were an early precursor to IEEE 802.11 wireless LAN modems but operated at 900 MHz. It is this early paper many measurements that have been repeated for wireless LAN cards operating in the 2.4 GHz band for IEEE 802.11 networks are presented. For instance, the authors measured the throughput versus the number of simultaneous users, and the throughput versus signal strength. The authors have also performed a number of measurements which have not yet been performed for 802.11 hardware. For this reason, it is this early paper which has drawn a road map for measurements that can be made in future research.

2.2 Wireless LAN Network Performance Research

Several researchers in the wireless networking area have found that attempting to use traditional wired network protocols on wireless networks has resulted in unexpected difficulties. The use of wireless packet data has opened large areas of new research areas. This section presents some of the recent work in wireless LANs and some of the terminology used.

2.2.1 Network Performance Statistics

The literature on wireless LAN network performance present measurements using various different statistics. The different network statistics used are often ill-defined. For this reason, this thesis defines the different statistics as they will be used throughout this document. Further, this thesis briefly discusses the usefulness and usage of these statistics.

Delay: Latency and Round Trip Time

One of the key statistics used in evaluating the performance of a network connection is the delay experienced by data which travels from one host to another. The term *latency* is used to describe this concept. However, care must be used in the use of the term because it is not always clear whether latency refers to the time to travel from one host to another, or the time required to transmit a packet and receive an acknowledgment, or some other delay. Unless otherwise stated, this document will always refer to latency as the time required to travel from one host to another. A related term to latency is the *round trip time* of a network connection. The *RTT* of a connection is the time required for a data packet to travel from one host to another and to return back to the original host. Latency and round trip times are typically measured in milliseconds for IP-based networks.

Throughput

Throughput is a measurement of the average rate that data (in bits) can be sent between a one user and another and is typically reported in kilobits per second or megabits per second, where kilo- is 10^3 and mega- is 10^6 (as opposed to 2^{10} or 2^{20} which are often used to define data sizes as in kilobytes or megabytes). As with latency, care must be taken in the definition of this term. The throughput of the same network connection can vary greatly depending on the protocol used for transmission (e.g., UDP, TCP, etc.), the type of data traffic being sent (e.g., HTTP, FTP, VoIP or other traffic) as well as the quality and *data bandwidth* of a network connection. This is quite different from latency which generally does not vary for different protocols or traffic types. Throughput is measured at the highest protocol level possible to reflect as accurately as possible the performance that will be experienced a user. Throughput is, thus, computed using the amount of data in the payload area of the highest protocol layer (e.g., the UDP payload size) of the transmitted packets. Overhead due to protocol headers and checksums are not included in the calculations of throughput as it is defined in this thesis.

Data Bandwidth

The *data bandwidth* or channel data rate is the maximum available, raw rate at which data can be transmitted over a network connection. The data bandwidth of a connection is similar to the throughput of a connection except that the data bandwidth is the theoretical maximum rate at which data can be transmitted if all of the overhead and checksums of the protocols used is included and the multiple access protocol is completely efficient. Like throughput, data bandwidth is measured in units of bits per second or bps. However, the data bandwidth of a of a network connection is always larger than the measured throughput of the connection. For example, later in this thesis, wireless LAN connections with 11 Mbps data bandwidths have been measured to have throughputs of 2 Mbps.

Error Rates

The *error rate* of a data connection, until recently, has not been a common metric of the network performance of a connection. Packets can be lost by routers over long trips across backbone networks due to collisions, but over short, LAN connections, typical wired or fiber optic transmission mediums have raw bit error rates on the order of 10^{-6} to as low as 10^{-14}. With the addition of error checking in many packet transmission protocols, bit errors in wired and fiber-based data transmissions are insignificant and, therefore, not often measured for local area networks. However, with the growing popularity of Voice over IP (VoIP), real-time video streaming and the use of wireless in data networks has meant that data packets are often dropped or excessively delayed with respect to the real-time requirements or shortcomings in wireless media. Thus, the error rate of a network connection has grown in importance. This has resulted in increased interest in bit error rates and packet error rates. The *Bit Error Rate* (BER) is the percentage of bits that are received in error or not received of those that are sent. The *Packet Error Rate* (PER) is the percentage of packets that are dropped or received incorrectly of those that are sent. Care must be exercised in the use the term BER because it can be used to imply the percentage of errors in only the payload data bits (e.g., excluding the header, footer and checksum bits) or to the raw data bits (e.g., including all overhead and non payload data bits). In this thesis the term BER refers to the error rate in just the payload data bits.

Delay Variation or Jitter

Delay variation or jitter is an important metric for quantifying data network performance, especially for VoIP and video streaming applications in which the protocol relies on regular arrival rates of data packets. As a result, the delay of a packet sent from one host to another is extremely important, but so is the *Delay Variation* or *Jitter*. The delay variation or jitter of a packet is defined to be the average variation in the arrival time of a packet and is reported in milliseconds or other appropriate time scale [8].

2.2.2 Network Protocols

The overwhelming majority of network measurements are carried out using either Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) packets. This is not surprising since most of the traffic on the Internet consists of mainly these two protocols. *UDP*, or *User Datagram Protocol* is a simple protocol. UDP does not guarantee whether or not packets will arrive, or if they do arrive they are not guaranteed to arrive in the order they were sent. In either case the sender will not receive an acknowledgment of the transmission. These characteristics makes UDP ideal for video or voice streaming applications in which retransmissions and acknowledgments are a waste of bandwidth and the transmissions can make use of the low overhead associated with UDP. *TCP*, or *Transmission Control Protocol* guarantees that packets arrive and arrive in the order transmitted. TCP also attempts to avoid network congestion by sometimes delaying the transmission of packets. TCP however has a higher overhead than UDP. This overhead is not significantly larger than for UDP though, so it is widely used in non-streaming applications.

2.2.3 Network Performance Measurement Techniques

There are many techniques and software tools available to measure the network performance of any data network which supports Internet Protocol (IP) transmissions, such as IEEE 802.11 wireless LANs. However, a relatively small number of methods have been applied to the actual measurement of wireless networks. This section outlines some of the popular techniques for wireless LAN measurement and the prior research that has gone into measuring the performance of wireless LAN connections. Additional information about how each of these techniques has been used in the literature is shown in Table 2.2 and Section 2.2.4.

Table 2.2: Network Performance Measurement Techniques Used in Wireless LAN Research

Measurement Technique or Software	Prediction Technique	Reference
A variant of ttcp, a UNIX, command line based program	None	[33]
Test file transfer using FTP	None	[27]
Test file transfer using FTP	None	[5]
Chariot 3.1, Test Packet Based Software	Test Packet Based Software	[10]
Harris WLAN Evaluation Software, Test Packet Based Software	Predicted minimum adjacent channel interference power ratio for different channel settings	[13]

Due to the diverse array of techniques available for the measurement of network performance statistics, i.e., those parameters defined in Section 2.2.1), an overview of the popular techniques is presented here. The following network performance measurement techniques and software products are intended to quantify the network performance of the network connection between two hosts. In general, the measurement solutions will send some sort of test data from one host to another using a certain protocol and a certain test pattern. Measurements are repeatedly made of the test data and then averaged for better accuracy. Some of the measurement techniques and software programs used to achieve this process are outlined in this section.

UNIX Command Line Based Software

There are several command line utilities that run under UNIX operating systems. These command line utilities are simple and intuitive to use (for individuals familiar with IP networks), and are often used to quickly test and diagnose problems with networks on a daily basis. Some examples of these programs include *ping*, and *TTCP*, but there are a great deal of others. Ping is a simple program that allows the user to send packets to a specific Internet host which acknowledges the packet. Ping allows the user to control the size of the data in a packet and to send repeated packets. However, the accuracy of Ping is low (on the order of 1 millisecond for the latency) and it is only capable of measuring the average latency for a packet to travel from one host to another, and the packet error rate. Ping has the advantage of not requiring a second software program to answer the test packets sent by the originating host since this is done by the operating system software. Alternatively, TTCP sends packets from one host to another and provides users with significantly more control of the way packets are sent between two hosts. TTCP allows the user to transmit data using the TCP or the UDP protocol and measures the latency and throughput of a network connection.

File Transfers Using FTP

A simple and inexpensive method for measuring the throughput of a connection is available using FTP software. In this technique, test files of a know size are transferred from one host to another host. The time required to transfer the file is then used to calculate the throughput. This technique, however, does not give any information about the latency of the connection. In addition, it is limited in that most of the data is sent in a single direction, and has a steady stream of maximally-sized packets sent continuously. Thus, the measured throughput using this technique tends to be optimistic because many network applications have vastly different network traffic characteristics, including packet size variations and fewer uni-directional transfers of data.

Test Packet Based Measurement Software and Chariot 3.1

Chariot 3.1 is a popular network measurement software tool. The from NetIQ Corporation, is representative of a class of network measurement software packages that precisely control the characteristics of test packets sent over a connection to measure a range of network performance statistics. Chariot 3.1 is capable of emulating a wide range of traffic types, such as Web traffic or Voice Over IP traffic, and a wide range of network protocols such as TCP, UDP and IPX [19]. For any of these protocols or traffic types, Chariot can measure the throughput, latency, jitter, and packet error rates [19]. The advantage of a software product such as Chariot is accurate control of the test traffic, combined with the ability to measure a wide range of network performance statistics. The difficulty with these types of products is that they can often be too costly for research and, thus, are not necessarily practical to use as a standard for comparing network performance.

2.2.4 Network Performance Measurement and Prediction Research

This section presents some of the research on wireless LAN network performance measurement, prediction and modeling. None of these references present a simple means for predicting the throughput or other network performance statistic. However, the diverse array of research does show how measurements of the network performance of wireless data networks have been performed. Additionally, several papers indicate the plausibility of predicting throughput based on a small number of parameters such as the path loss at a receiver location.

BER and Throughput Correlation to Delay Spread

Maeda, Takaya and Kuwabara published a measurement of wireless LAN performance and the validity of a ray-tracing technique to predict the performance of a wireless LAN [15]. The measurements were tracked in a small, highly-controlled radio frequency (RF) environment and indicated that the wireless LAN throughput and BER were correlated to the delay spread of the wireless channel. The researchers have not, however, presented any way to actually predict the bit error rate or throughput from the predicted delay spread profile output by a ray-tracing technique.

Early Wireless LAN Network Performance Measurements

Xylomenos and Polyzos explored the performance of UDP and TCP packets sent over several fixed IEEE 802.11 wireless LAN network connections in [32] and [33]. The research focused on throughput limitations caused by software implementation issues and operating system shortcomings. The researchers used their own modified version of the command line utilities ttcp, tcpdump and nstat under Linux to perform UDP and TCP throughput tests. All measurements were taken between three fixed locations and focused on varying the wireless LAN card types (PCMCIA or ISA) and the end-user computer hardware (i.e., Pentium 150 with 48 MB of RAM versus a Pentium 200 MMX with 64 MB of RAM). The researchers make recommendations for changes in the implementation of network protocols and Linux operating system enhancements. The measurements did not consider the effects of different physical locations, signal strength, or the effect of variations in the wireless communications channel on the network throughput.

IEEE 802.11 Throughput Measurements in a Hallway

Duchamp and Reynolds presented packet throughput measurement results for varying distances for IEEE 802.11 wireless LANs [6]. These measurements were performed in a single hallway. Thus, these measurements, too, suffer from failing to measure a representative environment. The researchers do not present a model to predict their results nor do they attempt to validate any sort of computer prediction technique. Their work does not consider multiple users and was more focused on estimating the potential range of the wireless LAN in a nearly free space environment.

Idealized Wireless LAN Performance Measurements

Bing presented measured results of the performance of an IEEE 802.11 Wireless LAN. In [2], Bing presents delay and throughput measurements as well as theoretically based throughput and delay estimations for various wireless LAN configurations. The results are given as optimal results and were measured on a lab bench rather than in an actual building environment. Therefore, the results presented are an upper bound on best possible results and do not extend into a site-specific wireless LAN performance prediction technique.

Causes of Throughput Variation in IEEE 802.11 Networks

Demir, Komar, and Ersoy compared the effects of different system configuration factors on IEEE 802.11, 2-Mbps DSSS wireless LAN performance as measured by throughput [5]. The authors measured the throughput of an FTP-based file transfer and the signal strength percentage reported by the wireless LAN card hardware for 1, 2 and 3 simultaneous users. The authors used a 2^k factorial design and sign table, which is a technique to determine what factors have the most significant impact on the throughput experienced by a wireless LAN user. The authors considered the SNR, the number of simultaneous users, and the file size used in the data transfer. The authors concluded that the number of simultaneous users has the greatest effect on the throughput experienced by a single user, that the SNR level had a moderate effect on throughput, and that the file size had no effect on the throughput. The difficulty with these results is that the authors have not used their data to present a model for predicting throughput; the SNR measurements are for an arbitrary, unknown vendor.

Wireless LAN Performance Issues

Prasad, et. al., review many of the important factors in wireless LAN deployment, including a basic review of the IEEE 802.11 standard, standard indoor propagation models and interference and coexistence concerns with wireless LANs [22]. The authors also present some measurement results, although without any information about how the results were produced, in which an IEEE 802.11b wireless LAN user's throughput is compared to the received signal strength. The results show that the user achieves a consistent throughput of about 4.8 Mbps until the received signal strength reaches about -85 dBm. From this point, the throughput in Mbps falls in an almost linear fashion relative to the dBm value of the received signal strength until it reaches zero at approximately -97 dBm. These results seem reasonable for a single user in a highly controlled environment.

Multiple-User Measurements of a Wireless LAN

Chariot was used by Kamerman and Aben to measure the throughput for on a wireless LAN using TCP packet transfers for each of the four data bandwidths defined in the IEEE 802.11b standard [10]. The authors performed the measurements for 1, 3 and 5 simultaneous users and compared it to an analysis of the data bandwidth which is used by various overhead, collisions and unused bandwidth. The authors claim a maximum actual throughput of of 0.82, 1.52, 3.41 and 5.17 Mbps for the data bandwidths of 1, 2, 5.5 and 11 Mbps, respectively. It should be noted that these results are for ideal RF scenarios in which the data bandwidth has been forced to a desired setting.

Impact of Using Overlapping DSSS Channels in Close Proximity

An exhaustive means of evaluating the interference between overlapping DSSS channels of IEEE 802.11 access points has been performed by Leskaroski and Mikael [13]. In this work, the authors used a simple network measurement tool called LANEval which was available from Harris/Intersil, the maker of several popular wireless LAN chipsets. The measurement tool sends test packets from a client to a server and quantifies the throughput on the link between the two hosts. The authors have used this tool to measure the performance of a wireless LAN link operating on channel 1 in the presence of another interfering LAN operating on a varying channel between 1 and 6. The authors have used this data to calculate a minimum path loss between clients that allows the clients to operate without significant impact on the throughput of one another. This is an extremely useful and practical result.

2.3 Summary of Prior Research

After extensive review of the current research into wireless LANs, it can be concluded that there is a need for simple, accurate throughput prediction models. While some basic information is available in the literature, no attempts have been made to predict wireless LAN throughput based on site-specific information. Further, the ability to predict the throughput of a wireless LAN in different locations can be critical to designing an efficient wireless network. For this reason, this thesis focuses on first developing a means to measure realistic, non-optimistic, site-specific throughput, testing this measurement technique with an extensive measurement campaign, and creating throughput prediction models to allow wireless LAN design using throughput as a design parameter.

Chapter 3

Software Development for Throughput and RSSI Measurement

Several major software packages were used to perform the measurements that form this work. However, one of these packages was developed and released as a commercial product as part of the work for this thesis. The commercial wireless LAN measurement products LANFielder and SiteSpy were developed, marketed and sold at Wireless Valley Communications, Inc. The process of simultaneously developing a measurement product for research and commercial applications was a key contribution of this thesis. This chapter presents the key concepts and functioning of the LANFielder and SiteSpy products as well as some of the other custom software that was used in measurements conducted as part of this research. These measurements are presented further in Chapter 4.

3.1 Motivation for a Wireless LAN Measurement Tool

3.1.1 Difficulties with Data Networks

Data networks have, since their inception, presented design and performance measurement issues. The key problem is that in order to design a well functioning distributed data network such as an IEEE 802.3, Ethernet network or an IEEE 802.11 or 802.11b wireless network, accurate information about the performance of the design is needed. To

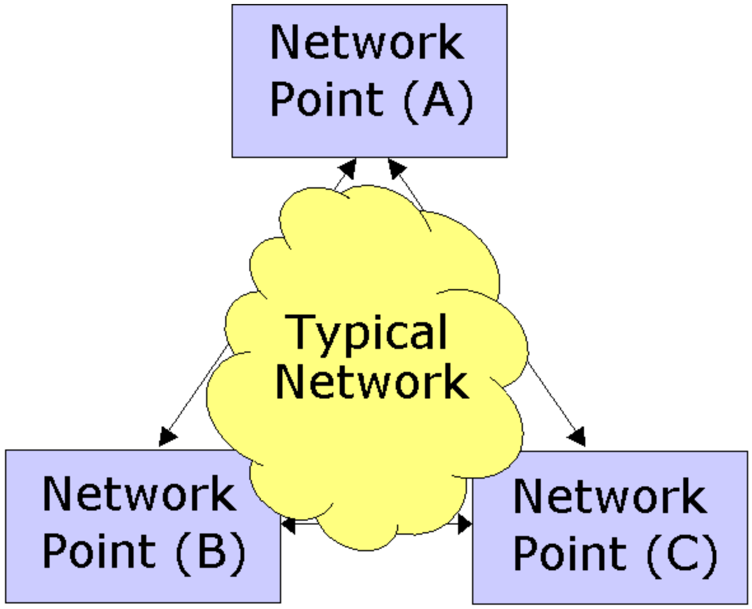

Figure 3.1: An example data network.

accurately measure the performance in a data network between different points or locations, data network measurement tools typically require two measurement agents. The reason for this is that by controlling the data flow between point A and point B, as in Figure 3.1, accurate information about the performance of the link between A and B can be known. If a single agent at A is used, it is impossible to know if the performance between point A and B or between point A and C or a combination of both is being measured due to the nature of distributed packet networks. If instead agents at A and B send test data back and forth over the single link between A and B, the performance of that link can be accurately determined.

3.1.2 The Client/Server Architecture

Data networks are often used as a way for many users to access a shared resource at the same time. For example, members of a company may access email from a single computer or people all over the world may access the same web page using the world wide web. The nature of this shared access has led to the development of the client/server architecture. This architecture, in its most basic form, allows many different clients (i.e. users at different network locations) to access a single shared resource on a server. The server allows clients to connect and access information on that server. Faced with the difficulties of monitoring and quantifying the performance of data networks, the Client/Server architecture have often been used to accurately measure specific network connections. This architecture both mimics the way in which data networks are most often used and also allows the performance of specific network links to be identified. Referring to Figure 3.1, by installing a measurement server at network point A, and clients on points B and C, it is possible to reliably and accurately measure the performance of the links between point A and B as well as the network link between points A and C.

3.1.3 Wireless Data Network Challenges

Several techniques are available for the measurement of network performance as discussed in Section 2.2.3. The goal of this thesis is to predict a realistic throughput that reflects the actual throughput that a user can expect to see during typical operation of a wireless LAN. For this reason, the optimistic, uni-directional throughput measured by a FTP measurement technique, while popular, was not selected. Although many different protocols can be easily implemented and certainly could be added at any time, this thesis focused on a test packet based solution using UDP, which could measure the throughput, packet error rate and latency of a wireless LAN connection, averaged over both directions, with was selected.

It was found that the network measurement products which perform packet based network measurement are incapable of handling the site-specific nature of wireless LAN based data networks. To clarify, wireless LANs can have a different throughput at different locations. The throughput of the wireless LAN experience will vary as the user moves closer or further away from an access point. For this reason, a measurement tool is needed which can record the network performance of a wireless LAN at specific locations. LANFielder and SiteSpy were designed as add-on modules to SitePlanner, an indoor, site-specific wireless communications system design tool. Using LANFielder and SiteSpy, measurements of various network performance statistics in a site-specific manner could be taken. In addition, the software tools utilize a realistic measurement technique to ensure accurate measurement of wireless LAN networks.

3.2 LANFielder and SiteSpy Details

LANFielder is a plug-in software module to the Wireless Valley Communications, Inc. design and measurement tool, SitePlanner. SitePlanner 2000 was used with the LAN-Fielder 1.1 plug-in module. SiteSpy is a stand-alone version of LANFielder that can operate without the use of SitePlanner's precise site-specific modeling. Both products simply provide a different interface to the same network measurement core. LANFielder utilizes the full three-dimensional model provided by SitePlanner to record measurements site-specifically. SiteSpy utilizes a text based method of associating a location with a throughput, and is more useful as a stand-alone server or traffic generation tool. Patents were filed for this product at Wireless Valley Communications. As will be seen in Chapter 4, this is exactly how LANFielder and SiteSpy were used for the measurements performed for this thesis.

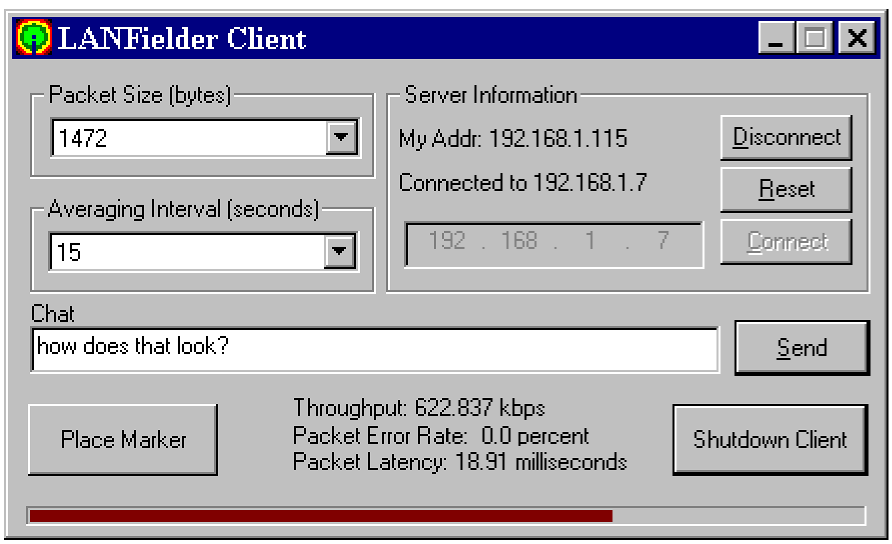

Figure 3.2: A LANFielder Client in operation

3.2.1 Client/Server Architecture

LANFielder and SiteSpy use the Client/Server architecture to measure the performance of a wireless LAN link. When LANFielder or SiteSpy are started, the user is given the option of using the program as a client or a server. Measurements can be made using a LANFielder client and server, a LANFielder client and a SiteSpy server or any other variation thereof. However, 1 client and 1 server are needed to perform measurements.

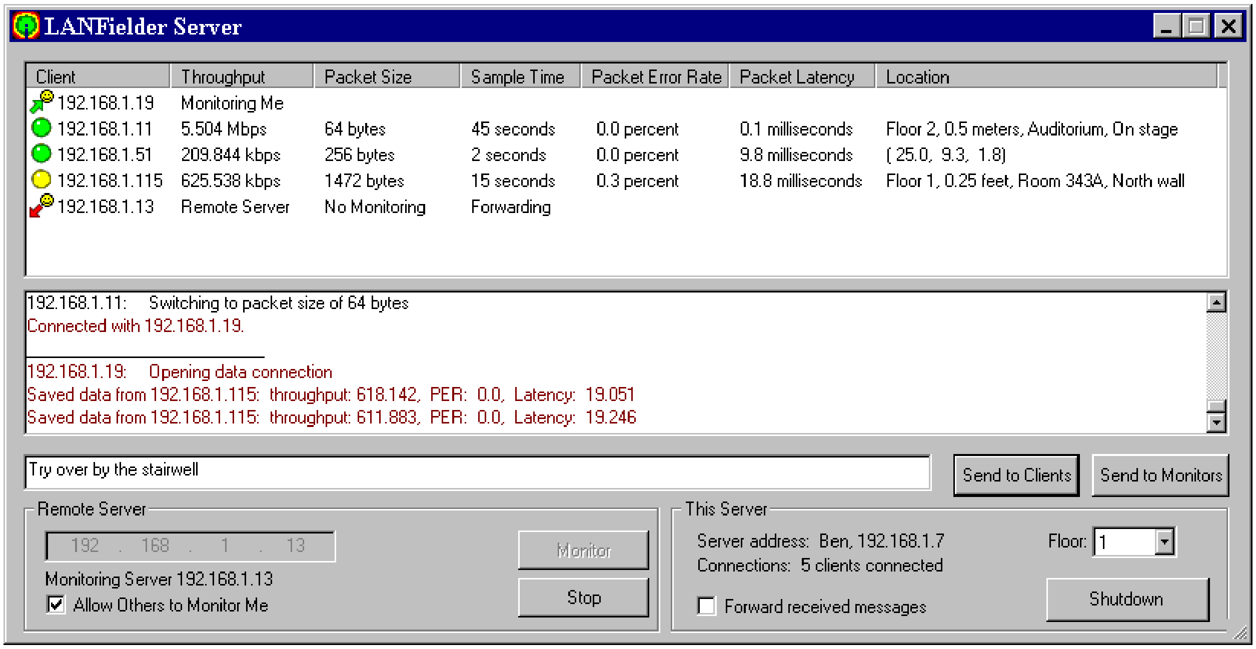

Figure 3.3: A LANFielder Server with several clients connected.

3.2.2 Site-Specific Data Recording in LANFielder and SiteSpy

SiteSpy does not support true three dimensional measurement recording. Rather, SiteSpy is typically used as a stand-alone server, and as a traffic generator to provide a source of interference. However, a SiteSpy client can save measured data by associating 4 textual tags with a measurement point. These 4 tags are intended to give an approximate location label to each recorded measurement point. The four tags give a floor number, receiver height, room number and general location to each measurement location. For convenience, these location tags can be loaded from a pre-defined file containing pre-set measurement locations or can be loaded on the fly as measurements are performed. To save the data, a user simply clicks the Save Data button and the current settings and current measured statistics are saved to a logfile along with a time stamp and the four text-based location tags.

LANFielder works in much the same way as SiteSpy, except instead of assigning textual tags to each measurement point, a precise X,Y,Z coordinate is used. This coordinate is derived from a mouse-click at a location in a three dimensional model of the environment and a predefined receiver height or can automatically stored using a positioning system such as GPS. The current settings, measured statistics, time stamp and X,Y,Z coordinate are saved to a logfile and also directly within the model of the environment.

3.2.3 Measured Network Performance Statistics

LANFielder and SiteSpy can both be used to measure the throughput, packet error rate and packet latency at each location. The definition of these terms are given in Section 2.2.1. To measure these statistics, each client is typically configured to transmit a test packet with the maximum amount of data possible. The format of a test packet is discussed further in Section 3.2.4.

3.2.4 Test Packet Format

The format of a test packet at the Ethernet layer is shown in Figure 3.4. This test packet is encapsulated by the MAC and then the PHY layer as is shown in Figures A.8 and A.7 respectively.

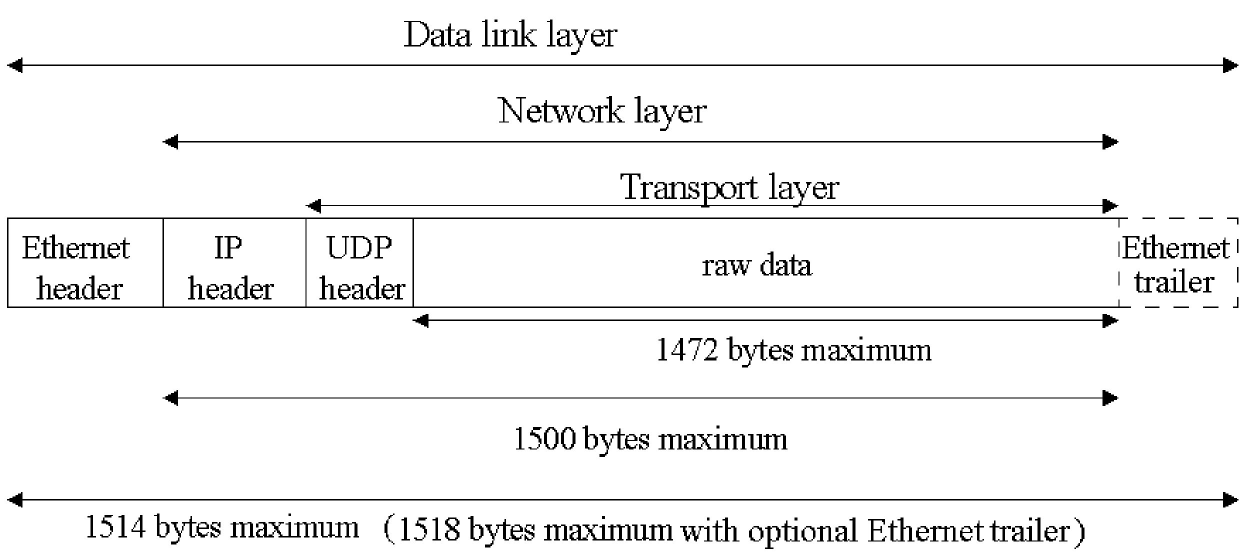

Figure 3.4: Format of LANFielder Test Packet

A single test packet is sent using UDP and consists of random data that is sent from the client to the server. The server echos an exact copy of the packet back to the client. This process is repeated for a timed measurement period. This period is called the averaging interval because after it is completed, the throughput, packet latency and packet error rates are calculated based on the average over all of the test packets transmitted.

3.2.5 Additional Configurable Parameters of SiteSpy and LANFielder

In addition to the packet size, SiteSpy and LANFielder allow the user to set the averaging interval and timeout period of a measurement period. The averaging interval is the period over which test packets are sent and the network performance statistics are calculated. Longer averaging intervals provide greater repeatability, but can result in time costly measurement campaigns. Since only one test packet is ever in transit for a single client/server pair, a timeout is needed to identify when a packet has been lost. The timeout interval set by the user is the maximum time after a packet is sent that the client will wait for a response from the server before sending a new packet.

3.2.6 Network Performance Calculations

The equations used to calculate the network performance statistics are shown in Equations 3.1, 3.2 and 3.3. Several important points need to be made about these calculations. First, note that throughput is based solely on the packets which are sent and received successfully. As shown in Equation 3.3, the latency is defined as the round trip time, from client to server and back, divided by two. The time required for this trip is measured by the client. Additionally, note that there is a direct relationship between throughput, packet error rate and latency. This relationship is given in Equation 3.4. This relationship brings up an important additional parameter. There is a timeout parameter that LANFielder uses to determine if a packet has been dropped or lost. During the averaging interval, if the client does not get a packet returned from the server for the duration of the timeout period, starting from when a test packet is sent, the packet is considered lost, and a new packet is transmitted. When a packet is dropped in this manner it is not counted in the $payload\_bits\_sent\_successfully$ parameter, or in the $payload\_bits\_received\_successfully$ parameter, but it is counted in the $packets\_not\_received$ and $total\_packets\_sen$ parameters.

$$Throughput = \frac{payload\_bits\_sent\_successfully + payload\_bits\_received\_successfully}{averaging\_interval\_duration} \tag{3.1}$$

$$PacketErrorRate = \frac{packets\_not\_received}{total\_packets\_sent} \tag{3.2}$$

$$PacketLatency = \frac{1}{2}(mean\_round\_trip\_time) = \frac{averaging\_interval\_duration}{2 * total\_packets\_sent} \tag{3.3}$$

$$Throughput = \frac{(bits\_in\_packet\_payload) * (1 - Packet\_Error\_Rate)}{Packet\_Latency} \tag{3.4}$$

3.2.7 Rationale Behind Network Performance Measurement Technique

The measurement technique described above has its advantages and disadvantages for evaluating the latency, error rate and throughput of a connection. First of all, compared to FTP measurement methods and certain FTP emulating packet based measurement techniques, this measurement technique will often measure a smaller throughput but a more consistent packet latency. That is, the measurement technique used by LANFielder is a pessimistic measurement of throughput. There are, however, several good reasons for making measurements in this manner.

1. The FTP measurement technique, by nature, provides an optimistic measurement because it is continuously sending a large quantity of data in one direction. Actual network connections may not be so directional. Thus, the ability to submit large frames of data for transmission (i.e. multiple packets) at once is not typical of many types of network transmissions. LANFielder test traffic is split evenly between the up-link and down-link direction. In addition, only 1 packet is transmitted at a time. Thus, a more realistic, albeit pessimistic, throughput is measured.

2. LANFielder avoids the problem of the congestion avoidance algorithm found in TCP used on wireless LAN connections [34]. TCP makes the assumption that packet losses are caused by network congestion. Thus, when a packet is lost, TCP will de-

crease the rate at which packets are sent. This can be a real problem in wireless LAN connections where packet losses can be caused by a signal fade or a bursty interfere in the wireless channel. Since fading and interference is a often much more transient than channel congestion, the TCP congestion avoidance algorithm can often be tricked into decreased transmission even when congestion is not present. This phenomenon can be measured using TCP packets as is done using the FTP and other techniques. The difficulty with this technique, though, is that different TCP algorithms (e.g. TCP Tahoe, TCP Reno, etc) use different congestion avoidance methods (e.g. Fast Restart etc.) [21]. For this reason, LANFielder uses UDP packets. This means that no congestion avoidance will take place. However, the effect of different congestion avoidance methods can be evaluated analytically using statistical techniques and the measured packet error rate and packet latency statistics found using UDP transmissions.

Due to the differences in the measurement technique from other published results (such as [33], [27], [5] [10] [13]), care should be used in comparing these measurements to other published results. Since throughput measurements are so sensitive to the measurement technique, it is often difficult to compare these measurements carefully. A possible solution to this is that measurements of throughput, made using different techniques, can first be normalized to the maximum measured throughput. This has the potential to minimize the differences (i.e. normalize) caused by the particular throughput measurement technique.

3.3 Custom RSSI Measurement System

3.3.1 Motivation of RSSI Measurement Technique

The received power from DSSS signals such as IEEE 802.11b modems, is spread across a wide bandwidth and thus can be well below the noise floor. For this reason, measuring the RSSI (received signal strength intensity) using a spectrum analyzer can be quite difficult. In addition, a spectrum analyzer is unable to identify times when the wireless LAN is not transmitting. Thus, any measurement of power would need to be performed when the LAN is carrying as much traffic as possible, but would still slightly underestimate the actual power received. For the same reason, a power meter would have difficulty measuring the RSSI from a DSSS signal. Alternatively, measuring continuous wave signals neglects the effect of frequency selective aspects of the wireless channel. A full wideband channel sounder or network analyzer could have been used to measure the impulse response and thus the RSSI at various locations. However, a more practical and applicable method of RSSI measurement was desired.

3.3.2 Justification for RSSI Measurement Technique

The goal of this thesis is to allow installers of wireless LANs to easily predict the network performance of wireless LANs. If network performance models are to be based on measurements, then it is imperative that it be simple to measure the needed values. For this reason, this thesis has measured the RSSI of wireless LAN signals using a calibrated wireless LAN card which directly reports the RSSI of a wireless LAN signal in dBm. This value is the power which is spread across the 22 MHz bandwidth which is received by the wireless LAN card after despreading. Such data is sometimes provided by the software included with a wireless LAN card.

3.3.3 RSSI Measurement Software

The RSSI Measurement created for this thesis has a simple design. The software basically polls a wireless LAN card driver at a regular interval. The driver software, which is provided with every wireless LAN card, reports the signal and noise levels averaged since the last time the information was requested from the card. These averages are generally dB averages due to the logarithmic amplifiers in the modem receiver. The signal and noise powers should, however, be averaged linearly, and not in dB values. Thus, software was written to poll the wireless LAN card with a high frequency (5 times per second) and to linearly average the data over a period of 5 seconds. Thus, each measured signal and noise power measurement is a linear average of 25 measurements.

This is an important design decision. In order to more accurately measure the received signal strength, linear averaging of RSSI samples should be used. However, since the card drivers only perform dB averaging, the wireless LAN card driver should be polled as often as possible and the polled values can then be linearly averaged. However, there is a limit to the rate at which the polling can take place, due to the nature of the operating system and the wireless LAN card drivers. In addition, requesting the received signal strength from the wireless LAN card driver will impede data transfers. For this reason, the RSSI was measured separately from the throughput.

3.3.4 RSSI Measurement Hardware

A wireless LAN, IEEE 802.11b wireless LAN card was used in to perform the RSSI measurements. The RSSI is measured by the wireless LAN card by measuring the strength of beacons transmitted by the wireless LAN access point. The noise is measured by measuring the received signal strength when no data transmission is detected. The signal strengths are averaged (in dB) repeatedly by the wireless LAN card until requested by the driver that interfaces the card to the operating system.

3.4 Summary of Measurement Software

In this Chapter, the details of software that was developed for Wireless LAN measurements are given. LANFielder and SiteSpy were two tools developed at Wireless Valley to allow the site-specific measurement of wireless LAN systems. In addition, RSSI measurement software was developed which interfaces to a CISCO Aironet API and allows the measurement of received signal strength by actual wireless LAN cards. The developed software allows the precise, location based measurement of throughput, packet error rates, packet latency and received signal strength and noise power for test-bed wireless LANs. The capabilities of this software was used in a measurement campaign detailed in the next chapter.

Chapter 4

Wireless LAN Measurement Campaign

4.1 Measurement Campaign Overview

In order to design a wireless LAN, it is often necessary to perform a measurement campaign to ensure the wireless LAN provides the needed coverage and capacity (i.e. throughput) for the intended users. With this need in mind, a full measurement campaign was performed in which multiple access points were setup in three different typical indoor environments. Measurements of the RF performance and network performance of the wireless LANs were measured for these test installations.

A primary goal of the measurement campaign was to compare the performance of two IEEE 802.11b wireless LAN systems from different vendors under various different operating situations and to collect performance data to be used to develop empirical models of wireless LAN performance. The full details of the measurement campaign and some of the results are presented in this chapter. The measurement results were used to create empirical models of wireless LAN network performance, which is presented in Chapter 6.

4.1.1 Measurement Test-Bed

Two different IEEE 802.11b wireless LAN systems were compared using network performance measurements. The two systems used were made by WaveLAN (Lucent/ORiNOCO) and 3Com (AirConnect). Some of the specifics of these two systems are presented in Table 4.1. These two systems were configured as an infrastructure network in which a single access point was placed in a certain location. For some of the measurements, a BayNetworks access point and a laptop with a BayNetworks DSSS IEEE 802.11, 2 Mbps wireless LAN card was setup in a fixed location. For these measurements the BayNetworks equipment acted as interference sources. Measurements were then made in particular locations using laptop computers with wireless LAN cards corresponding to the vendor of the access point. The laptops are referred to as "clients" throughout this thesis. These two test-bed systems were setup at different times but all access points and clients were placed in the same exact locations so that the performance of the two systems could be compared in the identical environments.

Table 4.1: 2.4 GHz Wireless LAN Equipment Used in Network Performance Measurements

	Lucent (ORiNOCO) WaveLAN System	3COM AirConnect System	Bay Networks
System Type	IEEE 802.11b DSSS	IEEE 802.11b DSSS	IEEE 802.11 DSSS
Maximum Throughput	11 Mbps	11 Mbps	2 Mbps
Access Point Type	AP-500 Access Point	11 Mbps Wireless LAN Access Point	BayStack 660 Access Point
Access Point Transmit Power	15 dBm	18 dBm	15 dBm
Access Point Antenna Gain	2 dBi	2 dBi	2 dBi
Client Card Type	ORiNOCO Gold PC Card (PCMCIA)	11 Mbps Wireless LAN PC Card (PCMCIA)	BayStack 660 PC Card (PCMCIA)
Client Card Transmit Power	15 dBm	15 dBm	15 dBm
Client Card Antenna Gain	2 dBi	2 dBi	2 dBi

A diagram of all of the equipment used in the test-bed setup is shown in Figure 4.1. This figure will be used throughout this chapter to explain how measurements were performed. Note that all of the equipment was not used for all measurements. The laptop clients used to make all of the actual measurements are shown in the figure and are named Laptop1 and Laptop2. All other equipment is present as part of the test-bed to allow measurements to be taken in different usage scenarios in which 1 or 2 users and with or without interference setups were considered.

LANFielder and SiteSpy, the software developed for this thesis as described in Chapter 3, were used to measure the network performance and RSSI measurement of a wireless LAN user at different specific locations. Using the RSSI measurements, the path loss from

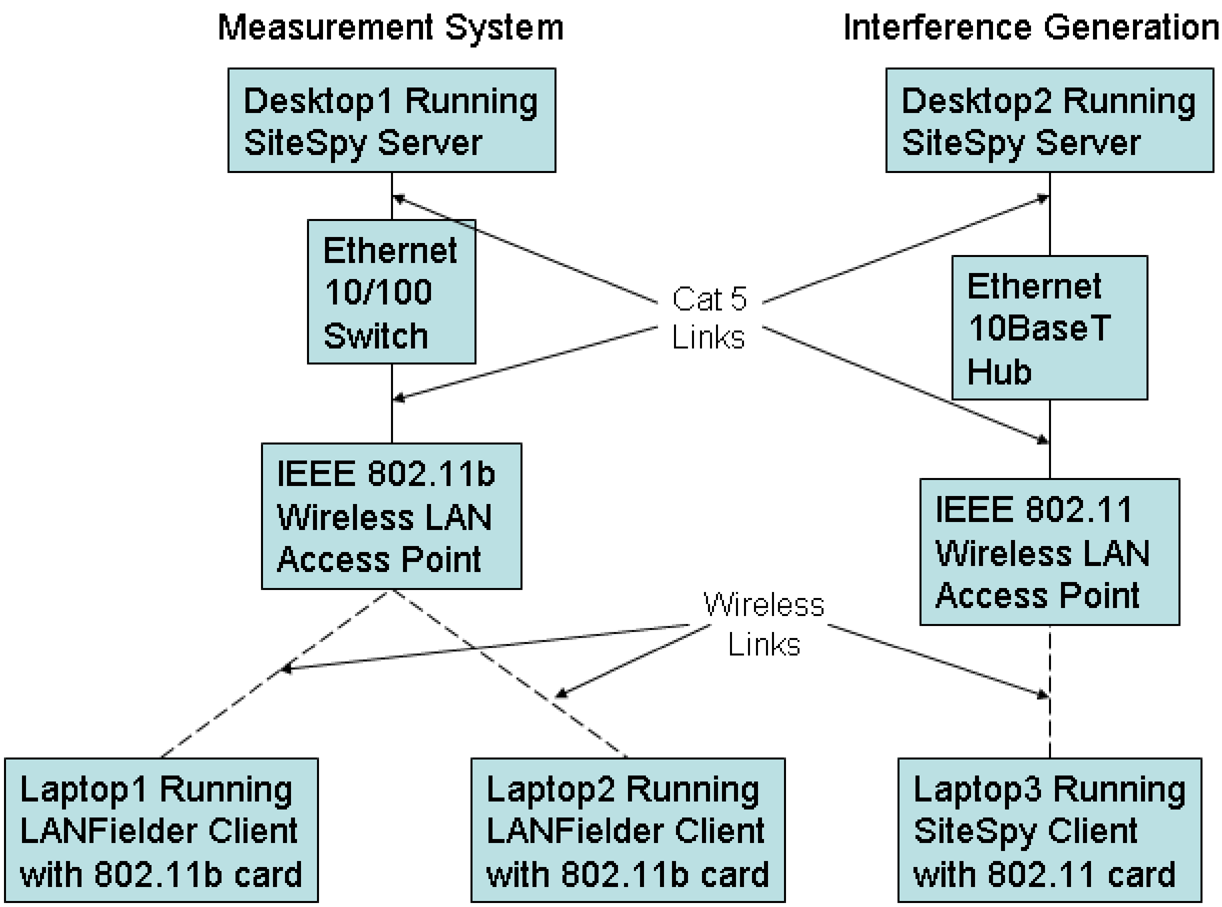

Figure 4.1: Logical layout of measurements and the software used.

the access point to the measurement user or client was known. The software was run on the client laptops. The software used on each individual computer is shown in Figure 4.1. Since LANFielder and SiteSpy use a client/server technique for measurements, SiteSpy was also run on desktop computers which were used as servers. These servers computers are also shown in Figure 4.1. It was necessary to connect the server computers to Ethernet hubs or switches as is shown in the figure. Ethernet hubs and switches were needed because the access points do not support being plugged directly into computers.

For some of the measurement scenarios, an interfering IEEE 802.11 wireless LAN system was used to evaluate how the test-bed IEEE 802.11b system performed in the presence of interference. The interference system was a BayStack 660 manufactured by Bay Networks (since acquired by Nortel Networks). The equipment is no longer available for sale. However, the relevant specifications for this equipment is listed in Table 4.1. In addition, this equipment functioned well as an interfering source and was setup using SiteSpy to generate data traffic that would interfere with the test-bed system. The interfering IEEE 802.11 equipment is shown in Figure 4.1. Throughout this thesis, the inferring equipment is referred to as the *interfering access point* and *interfering client*. The test-bed system is referred to as the *desired access point* and *desired client*. This terminology is used to distinguish the equipment.

4.1.2 Additional Equipment

Several computers, wireless LAN cards, access points and Ethernet hubs were used to perform network performance measurements. The configuration of the equipment is shown in Figure 4.1. The settings and specifics of the computer equipment used is shown in Table 4.2. All of the computers used were chosen to a) be as nearly identical as possible and b) reflect typical equipment for a wireless LAN-based network. With these two goals in mind, all of the computers used in the throughput measurements had the same operating system installed and had similar hardware configurations. To reflect typical usage scenarios, all mobile clients were laptops with PCMCIA wireless LAN cards in use. The fixed server computers were desktop computers with PCI 10/100 Ethernet cards installed. Additionally, the Ethernet hub (10 Mbit) and Ethernet switch (100 Mbit switched) are capable of supporting significantly more traffic than the wireless LAN connections. Thus, the measured throughput and latency effects caused by these back end network devices should be minimal.

4.1.3 RSSI Measurement Configurations

The path loss between the desired access points and each client location was measured using the custom software from Wireless Valley, described in Section 3.3. This software does not require a server computer to operate. Rather, the RSSI measurements can be made using a single laptop and an access point. The layout used for the measurements is shown in Figure 4.2. The RSSI measurements were made at all client locations with the laptop oriented in each of the four cardinal directions (i.e. North, East, South . For each location and direction, four different measurements were performed. The first two

Table 4.2: Computers and Settings Used in Measurements.

Computer Name	Operating System	Processor	RAM	Measurement Software	Network Card(s) Used
Laptop1	Windows 2000	Celeron 400 MHz	192 Megabytes	SitePlanner with LAN-Fielder Client	3Com, Wave-LAN IEEE 802.11b cards
Laptop2	Windows 2000	Pentium II 266 MHz	128 Megabytes	SitePlanner with LAN-Fielder Client	3Com, Wave-LAN IEEE 802.11b cards
Laptop3	Windows 2000	Pentium III 700 MHz	128 Megabytes	SiteSpy Client	BayNetworks IEEE 802.11 cards
Laptop4	Windows ME	Pentium III 850 MHz	256 Megabytes	RSSI Measurement Software	CISCO Aironet 340 IEEE 802.11b PCMCIA card
Desktop1	Windows 2000	Pentium Pro 200 MHz	128 Megabytes	SiteSpy Server	10/100 PCI Ethernet Card
Desktop2	Windows 2000	Pentium Pro 200 MHz	128 Megabytes	SiteSpy Server	10/100 PCI Ethernet Card

were of the received signal strength from the two different desired access points at the desired client locations. The third measurement was the received signal strength of the interfering access point at the desired client locations. Lastly, the fourth measurement was from the interfering access point, placed at the interfering client location and orientation, to the desired client locations. All four measurements were taken individually for each direction at each measurement point for a duration of 5 seconds.

It was necessary to reconfigure the measurement client for each of these four measurements to talk to the necessary access point. The measurements of the signal strength from the desired access point was performed for both access point manufacturers, 3Com and WaveLAN. For these measurements the RSSI measurement client used Wireless LAN DSSS Channel 1, which has a center frequency of 2412 MHz [30]. The measurements of the interference signal strength were performed at the frequency of the interfering wireless LAN system. That is, the measurement client measured the received signal strength from the interference system while both the interfering access point and the measurement client were using DSSS Channel 3. The wireless LAN measurement card could not measure the signal strength from Channel 3 on Channel 1.

To make these RSSI measurements, the software development laptop which had the custom software, Laptop4, was used with a commercial IEEE 802.11b wireless LAN card.

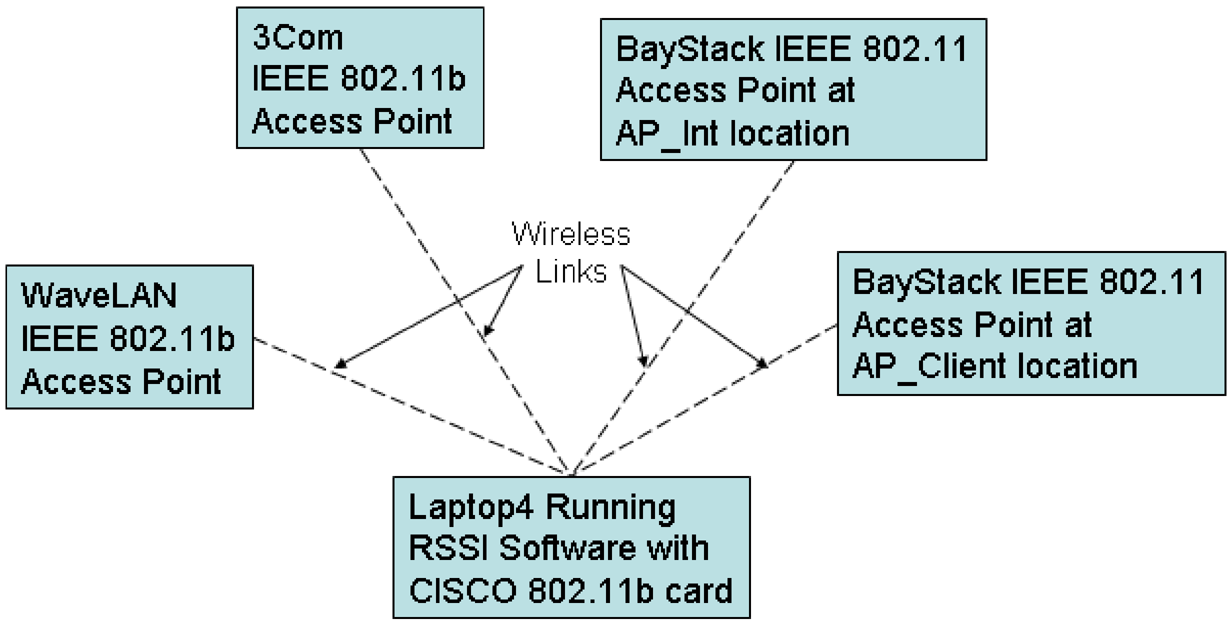

Figure 4.2: Logical layout diagram of the RSSI measurement hardware configurations

The specifications of the laptop used is given in Table 4.2. The wireless LAN card used was selected because it was capable of reporting the received signal strength from an access point more often than either of the 3Com or WaveLAN cards. In addition, the RSSI measurement card had a 2 dBi gain as do the other two client cards. It was assumed that the difference in the pickup pattern of the card was negligibly different from the patterns of the WaveLAN or 3Com cards. The measurements were averaged at each location and entered manually into SitePlanner for creating propagation models of the environment.

4.2 Organization of Measurements

The measurements were divided into 3 cases and 4 scenarios. Cases were the general environments used for measurements, Scenarios were how the equipment was configured. For each of the 12 combinations of cases and scenarios, 2 different IEEE 802.11b wireless LAN systems were measured. These two different wireless LAN systems were the Lucent/ORiNOCO WaveLAN and the 3Com AirConnect system. The configuration and specifications of these products was discussed in Section 4.1.1 and shown in Table 4.1. This section now discusses the different cases and scenarios that were used to evaluate these two wireless LAN systems.

4.2.1 Measurement Cases

The three measurement cases reflect different indoor environments that are likely to be used for wireless LAN networks. All three of these environments were selected to reflect a range of typical installation locations. The Cases are identified as I, II and III and are, respectively, an Open/Large Office Environment, a Long Hallway Environment, and a Closed/Small Office Environment. Measurements were taken on a single floor at various locations throughout these environments with the wireless LAN access point located in a central location. Specific details of each of theses cases are given in the following sections. Figures 4.3, 4.6, and 4.9 show an overhead view of the blueprint of these different cases. Different colors are used to represent the different partition types. The numbered locations represent the locations where measurements were taken.

Case I - Large/Open Office Environment

Case I was selected to reflect the performance that could be expected in a large, mostly unobstructed office environment consisting mainly of movable cubicle type partitions, drywall, and glass doors. A layout of this area is shown in Figure 4.3. This figure is a screen capture from SitePlanner. The numbered dots represent the measurement locations while the Larger circles are the omnidirectional patterns of the antennas used in modeling the access points. The interfering client location is also shown with a "*".

As can be seen from the blueprint, the main room in this area measured 9.4 meters wide, 15.7 meters long, and 2.7 meters high. The walls are made of primarily drywall and are shown in blue in Figure 4.3. There was a drop ceiling made of particle board ceiling tiles and the cubicle partitions (shown in gray in Figure 4.3) did not extend to the ceiling but were rather 1.7 meters high. The floor was carpeted in the large center room and tiled in all other areas shown. Two elevators of different sizes can be seen in the picture near measurement point 2 and are represented in orange. The furniture in this case consisted of mainly primarily wooden desks which were built-in to the cubicle partitions, computers, rolling chairs, as well as some shelving, lab benches and a refrigerator. A picture of the desired Access Point location is shown in Figure 4.4 and a picture of a typical client receive

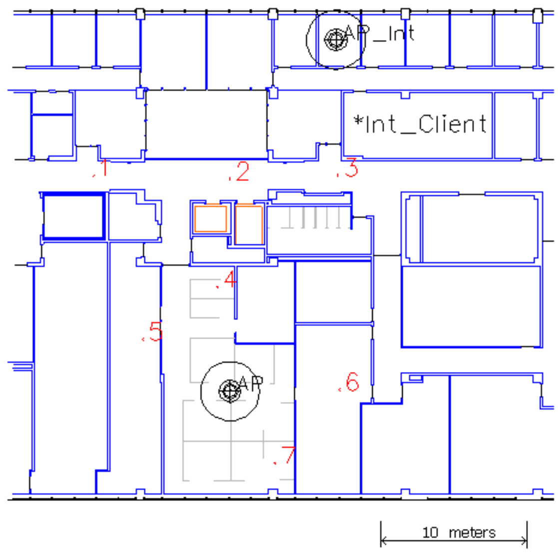

Figure 4.3: Physical Layout of Measurement Equipment in Case I

location is shown in 4.5.

Figure 4.4: A picture of the wireless LAN access point transmitter location for Case I

As can be seen in these pictures, stands and tripods were used to support the laptops and wireless LAN access points. These allowed the measurements to be taken without a user near the measurement platform.

Figure 4.5: A picture of the wireless LAN access point receiver location for Case I

Case II - Hallway Environment

Case II was selected to consider the performance of a wireless LAN in a hallway where waveguide effects could be have an effect on transmission. A blueprint layout of Case II is shown in Figure 4.6.

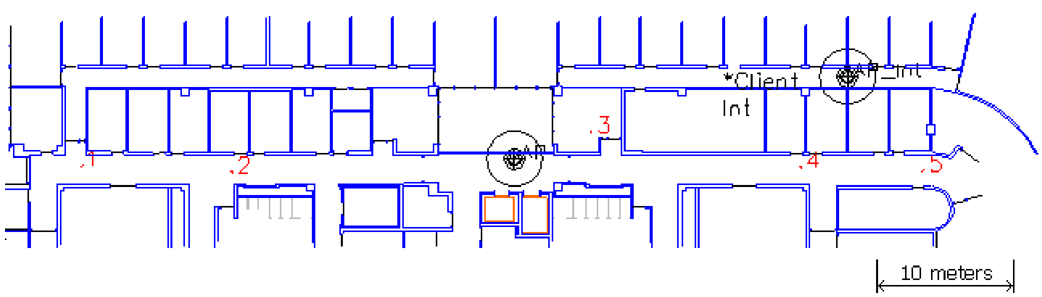

Figure 4.6: Physical Layout of Measurement Equipment in Case II

The walls of the hallway were painted cinder-block (shown in blue in 4.6). A drop ceiling at 2.7 meters high was again present. The hallway is approximately 27.5 meters long and 2.1 meters wide and has a tile floor. Note the orange squares shown in the center of Figure 4.6 near the Access Point. These squares represent the elevators in this area. Note that these are the same elevators in Case II. The access point and measurement clients were setup in the same manner as in Case I as is illustrated in Figures 4.7 and 4.8.

Figure 4.7: A picture of the wireless LAN access point transmitter location for Case II

Case III - Closed/Small Office Environment

Case III is intended to reflect a more closed office environment consisting of mainly closed offices rather than cubicles. The area is shown in blueprint form in Figure 4.9.

The rooms shown in the blueprint are almost entirely made up of offices. The smaller rooms typically measure 3.75 meters long and 3 meters wide. The larger rooms typically measure 4.6 meters long and 3 meters wide. Other larger rooms are also present in the area. The area has a carpeted floor and a particle board ceiling tile drop ceiling. Doors are made of wood and glass (shown in pink in Figure 4.9), and walls are primarily drywall (shown in blue in Figure 4.9). Office furniture consisted of wooden desks an chairs and some bookshelves and computer and photocopying equipment. The measurement clients and access points were setup in the same manner as was done in the previous 2 cases and is illustrated in Figures 4.10 and 4.11.

Figure 4.8: A picture of the wireless LAN access point receiver location for Case II

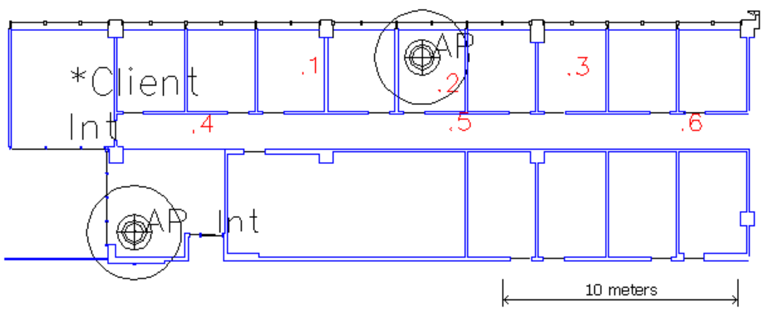

Figure 4.9: Physical Layout of Measurement Equipment in Case III

4.2.2 Measurement Scenarios

The four measurement scenarios were chosen to reflect the effects of different usage in which 1 or 2 users and 1 interfering client were combined in different ways. The four scenarios are identified as A, B, C and D. Scenarios A, B, C and D consist of 1 client, 2 clients, 1 client with interference and 2 clients with interference respectively. That is, referring to Figure 4.1, Scenario A used Desktop1, Laptop1, the Ethernet Switch and the IEEE 802.11b Access Point. Scenario B used all of the equipment of Scenario A and Laptop2. Scenario C used all of the equipment of Scenario A and added Desktop2, Laptop3, the Ethernet Hub and the IEEE 802.11 Access Point. Scenario D used all of the equipment shown. The actual locations of all of this equipment for each case is shown in Figures 4.3, 4.6, and 4.9.

Figure 4.10: A picture of the wireless LAN access point transmitter location for Case III

Scenario A – One User

In this scenario, a single laptop was setup to transmit test packets over a wireless link to an access point, over a wired link, through an Ethernet switch and to a desktop computer. The single laptop was then moved to various different measurement locations while the access point remained fixed for each case. Between 5 and 7 points were measured for each case. These points are identified by numbers in Figures 4.3, 4.6 and 4.9. For each measurement point point, four measurements were made in each of the four ordinal directions (North, East, South, West). That is the Ethernet card's antenna was positioned to face in each of the directions North, East, South and West in that order. This process was performed for both access points and the wireless LAN cards of the same manufacturer using Laptop1, the Ethernet Switch and Desktop1. At each location the throughput was measured using LANFielder as a client and SiteSpy as a server as shown in Figure 4.1.

Figure 4.11: A picture of the wireless LAN access point receiver location for Case III

Scenario B - Two Users

The two client scenario was performed much like Scenario A, except two laptop computers were moved to different locations. Both laptops transmitted test packets over the same access point to the same desktop computer. Again, the access point remained in a fixed location for each case. However, for the 2 client scenarios, all different combinations of the different measurement points were considered. That is, each possible pair of measurement points was measured. Measurements were made with clients at points 1 and 2, 1 and 3..., 2 and 3, 2 and 4,..., etc. However, measurements were not made for all possible permutations of locations. That is, interchanged measurement locations were not additionally made at locations 2 and 1, or 5 and 3, for example.

For each combination of measurement locations, the same direction was measured simultaneously. That is, the measurement clients were placed in two measurement locations with both wireless LAN cards facing in the same direction (e.g. North). Both clients then were used to measure their respective throughputs at each location simultaneously. The clients were then turned to the next direction and another measurement was taken. Thus, four measurements were taken at each pair of measurement points. Different combinations of directions were not measured.

For these measurements, Laptop1, Laptop2, Desktop1 and the Ethernet switch were used, as shown in Figure 4.1. The laptops were running LANFielder while the desktop was running SiteSpy. The chat facility of LANFielder was used to ensure measurements were taken in synchronization.

Scenario C – One User with Interference

Scenario C was performed in a similar manner to Scenario A except that another IEEE 802.11 system was operating in a static setup while measurements were performed. That is, a new laptop, Laptop3, was setup with an IEEE 802.11 (not an IEEE 802.11b) card to transmit to another computer, Desktop2. Both computers had SiteSpy installed and simply streamed the maximum amount of data possible to act as interference. The interfering client and access point remained in a fixed location while the same measurements that were made in Scenario A were repeated. Referring to Figure 4.1, Laptop3 and Desktop2 exchanged test packets to create interference. Laptop3 sent test packets using SiteSpy over a BayNetworks, IEEE 802.11 card to a BayNetworks Access Point, from there to a 10 Mbps Ethernet hub to a Desktop computer. While this occurred, Laptop1 was moved to various different locations and transmitted test packets to Desktop1 over the original IEEE 802.11b wireless LAN connections. Measurements were taken in the exact same locations and manner as was done for Scenario A.

The BayNetworks access point was configured using a different, incompatible Service Set Identifier (SSID) and a different, but overlapping, channel as the IEEE 802.11b wireless LAN access points. This ensured that the interfering network cards never connected to the desired access points or vice versa. For these measurements the desired access point was set to channel 1, as it was for all cases, while the interfering access point was set to channel 3. The center frequencies of these channels are 2412 and 2422 MHz, respectively. The 22 MHz RF bandwidth of the channels [30] means the two channels overlapped by 12 MHz. These channels were chosen because it has been found that at this overlap some effect on throughput occurs, but not usually a catastrophic interference [13].

Scenario D - Two Users with Interference

Scenario D was carried out in the same manner as scenario C, except that two clients were used, as was done in scenario B. That is, the two client measurements made in scenario B were repeated with the interfering wireless LAN system in use at the same static location as was used in scenario B. As a result, all of the equipment shown in Figure 4.1 was utilized for this scenario. Measurements were made using LANFielder on Laptop1 and Laptop2 which streamed data to Desktop1. These measurements were made with the interference system setup which consists of Laptop3, Desktop2 and the interfering Access Point.

4.3 Measurement Equipment

4.3.1 Physical Locations

For all cases the desired access points and the interfering access points were placed on metal tripods with wooden platforms. The tripods were extended so the base of the access point was 2.5 meters above the floor. The access points were oriented as they would be installed in a typical installation. The measurement client laptops were placed on wooden platforms on top of PVC and plastic stands. These stands were rotated in place to measure laptop throughput with North, East, South and West orientations. Measurements were recorded using LANFielder by clicking at the location in the model of the environment. Additionally, the software automatically saved all data to the Desktop1 computer as a backup.

4.3.2 Equipment Configuration

For all measurements the desired access points were configured to use channel 1. For the interference scenario measurements (scenarios C and D), the interfering access points were configured to use channel 3. The rational for this selection is given in Section 4.2.2. The access points were configured to only use the Distributed Coordinated Function (DCF) for all transmissions. That is, no polling or medium reservation using RTS/CTS packets was allowed. This was done by setting packet fragmentation thresholds at the maximum levels. All power-save features of the wireless LAN cards were fully disabled and no WEP encryption was used.

4.3.3 Software Configuration

The different software packages used in the measurement campaign are shown in Figures 4.1 and 4.2, and detailed in Table 4.2. All throughput measurements were made in LANFielder using a 15 second averaging interval. That is, when a measurement was performed, a full 15 second period was used to measure the network performance during which the measurement researcher(s) were motionless and distant from the client laptop. The value of the timeout duration used by LANFielder was set to 400 milliseconds. These values were selected through experimentation with various settings. It was found that these values provided a consistent, repeatable measurement at each location without causing overly excessive measurement times. The packet size was 1472 bytes (i.e. the maximum allowable payload). Identical settings were used for all SiteSpy clients. RSSI measurements were made using the custom software which recorded the signal strength reported by the wireless LAN card 5 times per second for a total of 5 seconds. This resulted in 25 samples of the received signal strength which were linearly averaged to yield a single RSSI measurement.

4.4 Measurement Precautions and Verification

A large number of precautions were taken to ensure the viability and repeatability of the measurements. Some of these precautions are listed here:

- All measurements were taken at night or during periods when little or no human presence would affect the measurements.
- The measurement software had a timer that allowed the researchers performing the measurements to step back several feet from the laptops, well out of the nearfield of the antennas, and remain motionless while the measurements were performed.
- Also, the platforms used to support the wireless LANs were made of wood to minimize the nearfield effects on radiation from the wireless LAN antennas.
- Care was also taken to ensure that measurements were not made while the elevator was on the same floor nor was it in motion when measurements were being taken.

- Measurements were never made when any of the microwave ovens, photocopiers or other potential 2.4 GHz noise sources on the floor were in operation.
- The interfering access points were configured to use a different SSID than the desired access points. This prevented desired wireless LAN cards from connecting to the interfering access point and vice versa.
- Reciprocity confirmation measurements were made in several random locations to ensure that the path loss from the transmitter to the receiver was the same as the path loss from the receiver to the transmitter. All measurements were found to be within 3 dB or better of the expected values. This is well within the tolerances of the CISCO wireless LAN card measurement capabilities.
- Several random measurements were made at non-measurement point locations and compared to predicted values from SitePlanner. All measurements were found to be within 5 dB of the predicted value.
- RSSI and throughput measurements were made separately, but in the same locations. This allowed a reliably accurate measurement of the RSSI and prevented the polling of the wireless LAN card from lowering the measured throughput.
- All wired Ethernet connections were significantly faster than the wireless links that they were attached to. Thus, for the IEEE 802.11b networks, all wired Ethernet links operated at 100 Mbps which is significantly faster than the 11 Mbps wireless LAN link. The IEEE 802.11 links were used with 10 Mbps Ethernet links. This careful pairing of wired and wireless Ethernet data bandwidths was selected to ensure that the wired links were not the constraining factor of the measurement links.

4.5 Measurement Conclusions

This chapter has outlined the measurement campaign that was designed to measure a wide variety of IEEE 802.11b usage situations. Two IEEE 802.11b systems were measured in three different environments, and with 1 or 2 users in operation with or without the presence of an interfering DSSS IEEE 802.11 system operating. Measurements were made in 18 different points and in each of four directions at every point. This wide range of measurement scenarios meant that 288 RSSI measurements and over 1600 throughput measurements were taken.

Chapter 5

Wireless LAN Measurement Results

5.1 Format Used to Present Throughput Data

This Chapter presents the measurement results that were obtained from the extensive measurement campaign described in Chapter 4. To present this data in an orderly format, MATLAB plots are used. The MATLAB plots presented throughout this chapter are arranged with the measured signal to noise ratio (SNR) for the desired link from the client to the access p on the x axis and the measured throughput on the y axis. These plots are created by pairing the throughput and SNR measured at each point. By combining the SNR, throughput pairs for all three cases, a wide range of throughputs and SNR levels considered.

5.1.1 Meaning of SNR

Note that the same SNR measurements were used for all scenarios. That is, the SNR between the access point and the client location was measured with 1 client operating. This value was seen to change slightly when two clients were operating, but not significantly. The SNR will most certainly change when interference is present, but since it is difficult to predict this value the SNR was not measured in these circumstances, instead another parameter, the SIR, was measured. The SIR value is intended to be a measurement of the net interference power for a measurement location. Additional plots of the throughput versus the SIR parameter as well as the computation of SIR itself are discussed further in Section 5.7.2. All plots presented in this section use the same SNR values and are labeled ”theoretical” for scenarios B, C and D because the SNR used was actually only measured in Scenario A and extrapolated for use in Scenarios B, C and D.

5.1.2 Overview of the Data Measured

Each scenario has two sets of figures, one for each wireless LAN system (WaveLAN and 3Com). Each set has three figures. These three figures present the measured data for from all 3 Cases, but just a single Scenario in unaveraged, and in two different averaged formats. The first of the three figures that make up a set is a plot in which each data point represents a different data point and cardinal direction. The second figure is a plot of the average of all four cardinal directions (i.e. North, East, South, West) forming a single, spatially averaged data point. In the second figure there is a single data point for each measurement point, rather than four data points for each measurement point as is the case for the first figure. The third figure is a smoothed version of the first figure. Each data point in the third figure represents the SNR of a single data point and a windowed average of the data points throughput and the throughput of the 4 data points closest in signal strength (i.e. RSSI) to the main data point and the data point itself. Stated another way, each of the points in the third plot of the set represents a (5 point) window average over data points with similar signal strengths. The actual number of 5 data points used in the window average to calculate the throughput was selected by trial and error to provide a degree of smoothing to the data while minimizing the smearing caused by selecting too large a number of data points.

5.2 Scenario A Data - One Client

5.2.1 WaveLAN, Scenario A, All Data

Figure 5.1 shows the single user, WaveLAN data. The data shown is for each cardinal direction. The general trend of a maximum throughput until the SNR reaches a critical level is clearly evident. Below this critical level the throughput becomes less stable and trends downward.

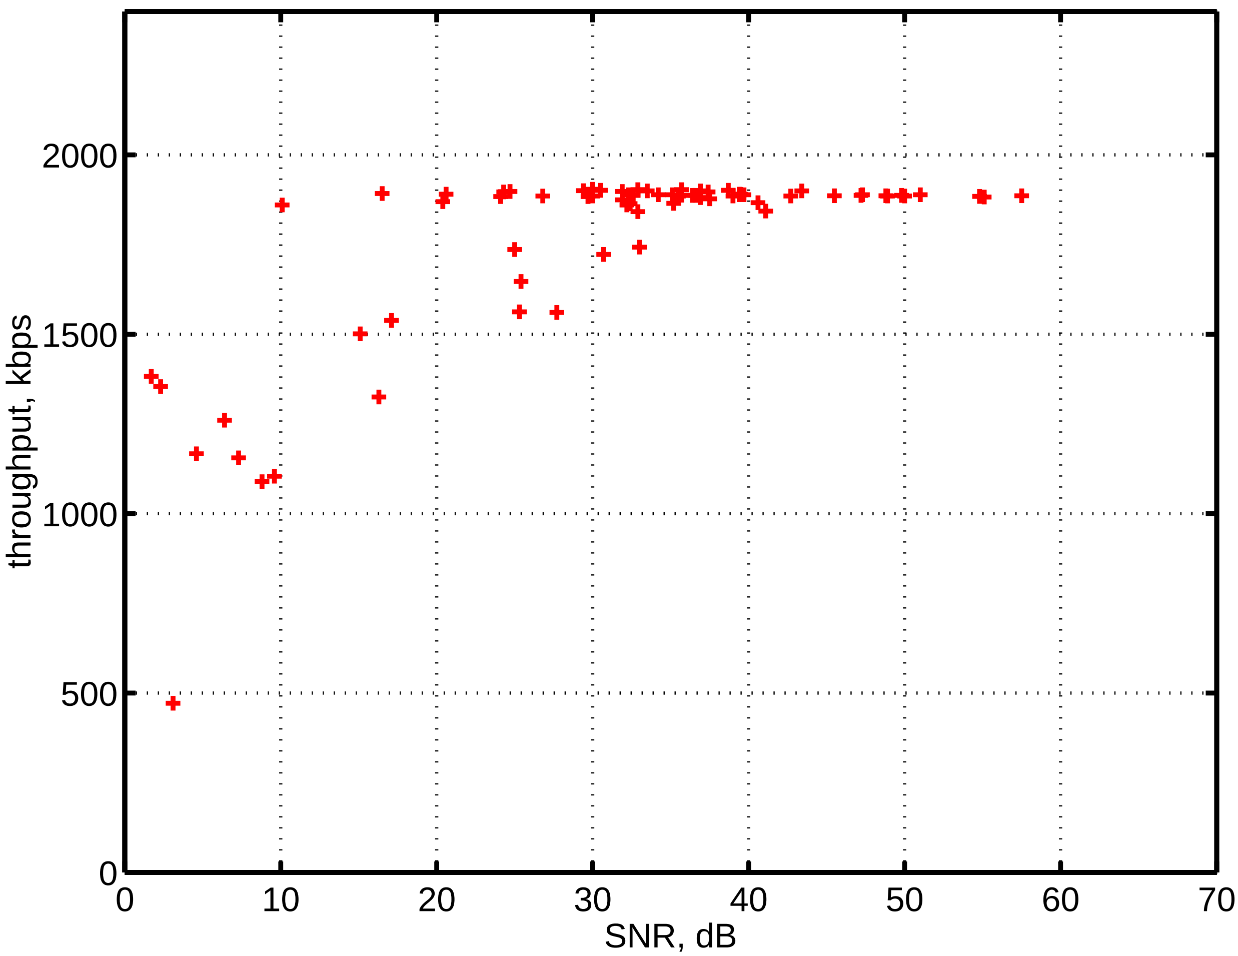

Figure 5.1: All measurements for Scenario A, WaveLAN System

5.2.2 WaveLAN, Scenario A, Spatially Averaged Data

In Figure 5.2 the positional averaging has smoothed out the throughput below the critical SNR level. This averaging has resulted in exposing the general trend of a dropping off throughput as the SNR decreases below the critical level.

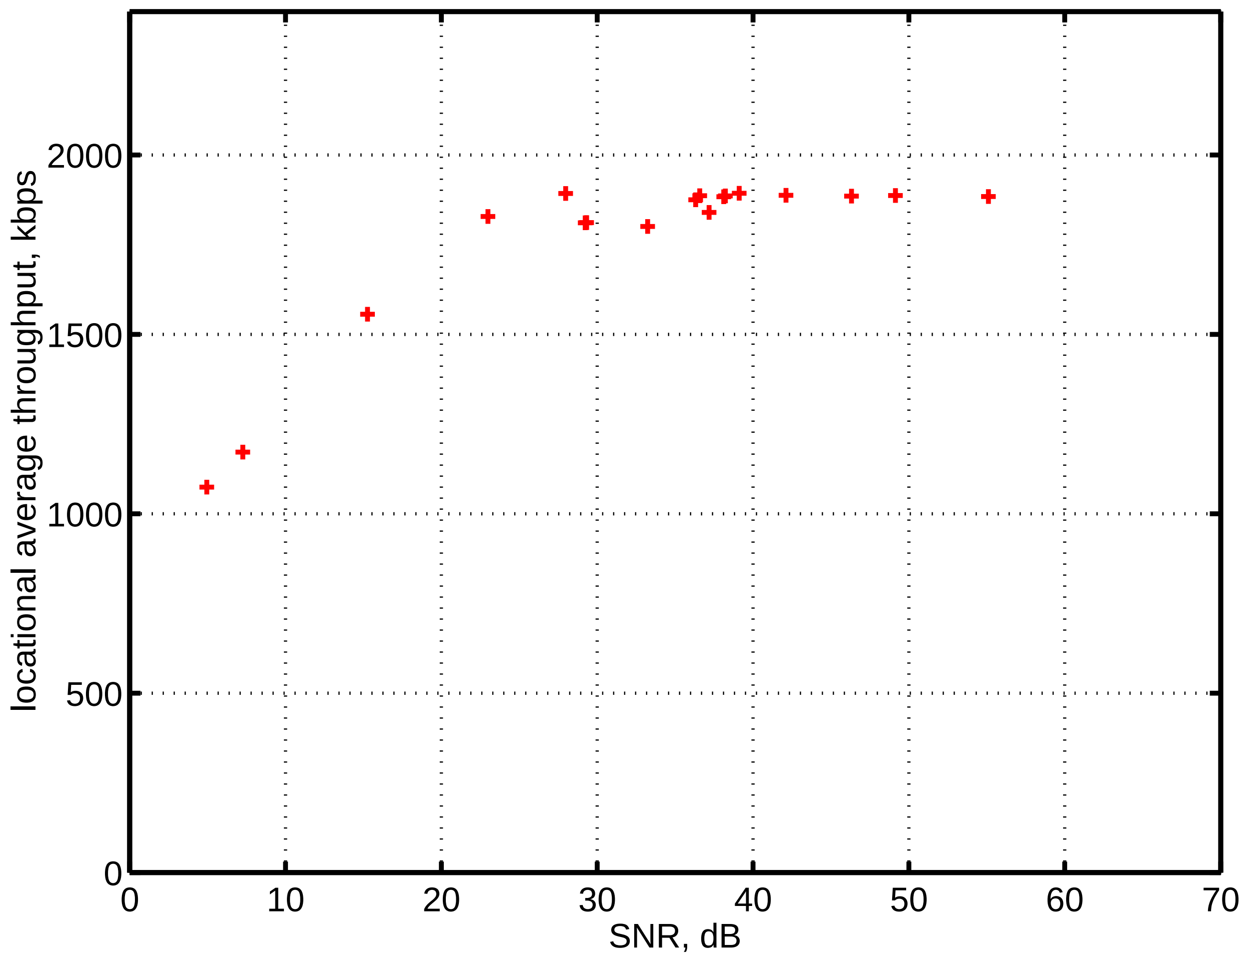

Figure 5.2: All measurements averaged for all directions at each measurement location for Scenario A, WaveLAN System

5.2.3 WaveLAN, Scenario A, Window Averaged Data

Figure 5.3, like Figure 5.2, also shows the general linear trend of the performance below the critical SNR level. The window averaging, however, has resulted in a clearer presentation of this trend.

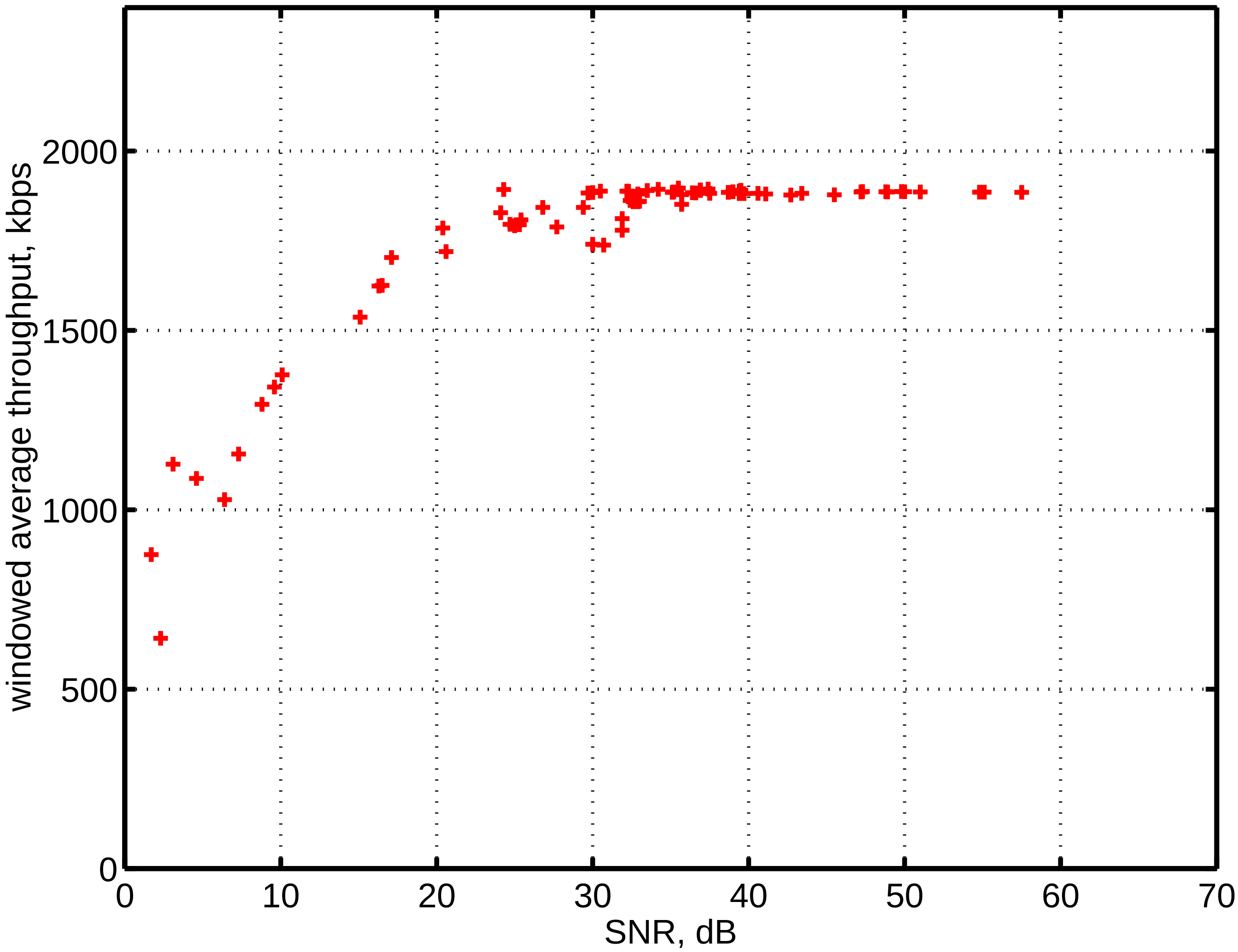

Figure 5.3: All measurements averaged over similar signal strength (i.e. RSSI) locations for Scenario A, WaveLAN System

5.2.4 3Com, Scenario A, All Data

Figure 5.4 shows the single user, 3Com data. The data shown is for each cardinal direction. The general trend of a maximum throughput until the SNR reaches a critical level is evident. Below this critical level the throughput becomes less stable and trends downward.

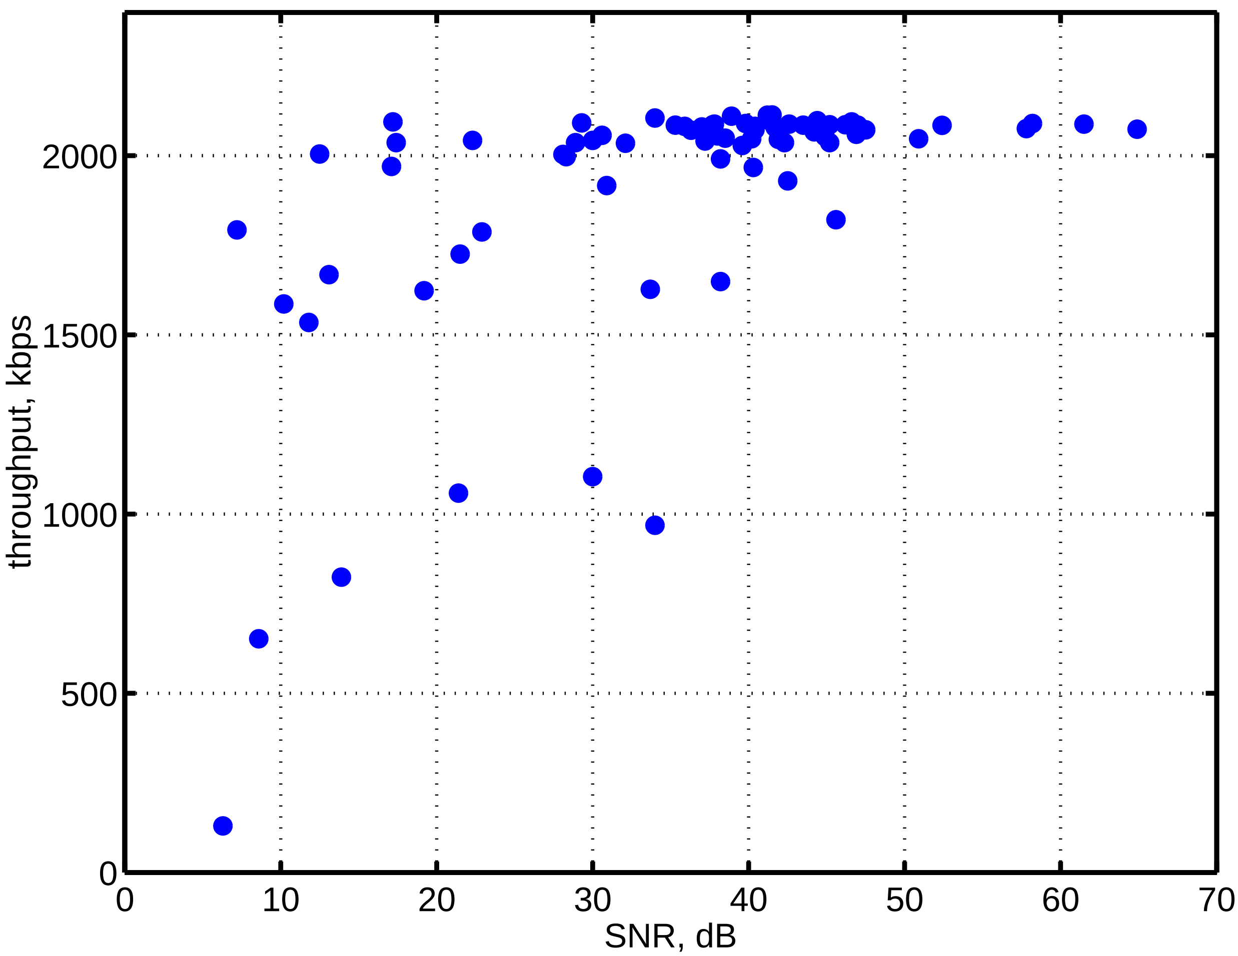

Figure 5.4: All measurements for Scenario A, 3Com System

5.2.5 3Com, Scenario A, Spatially Averaged Data

Figure 5.5 presents a clearer view of the measured data than Figure 5.4. The trend of a steady drop-off in throughput as the SNR drops below a critical point is visible, but the maximal throughput is not always reached when the SNR is above the critical SNR value.

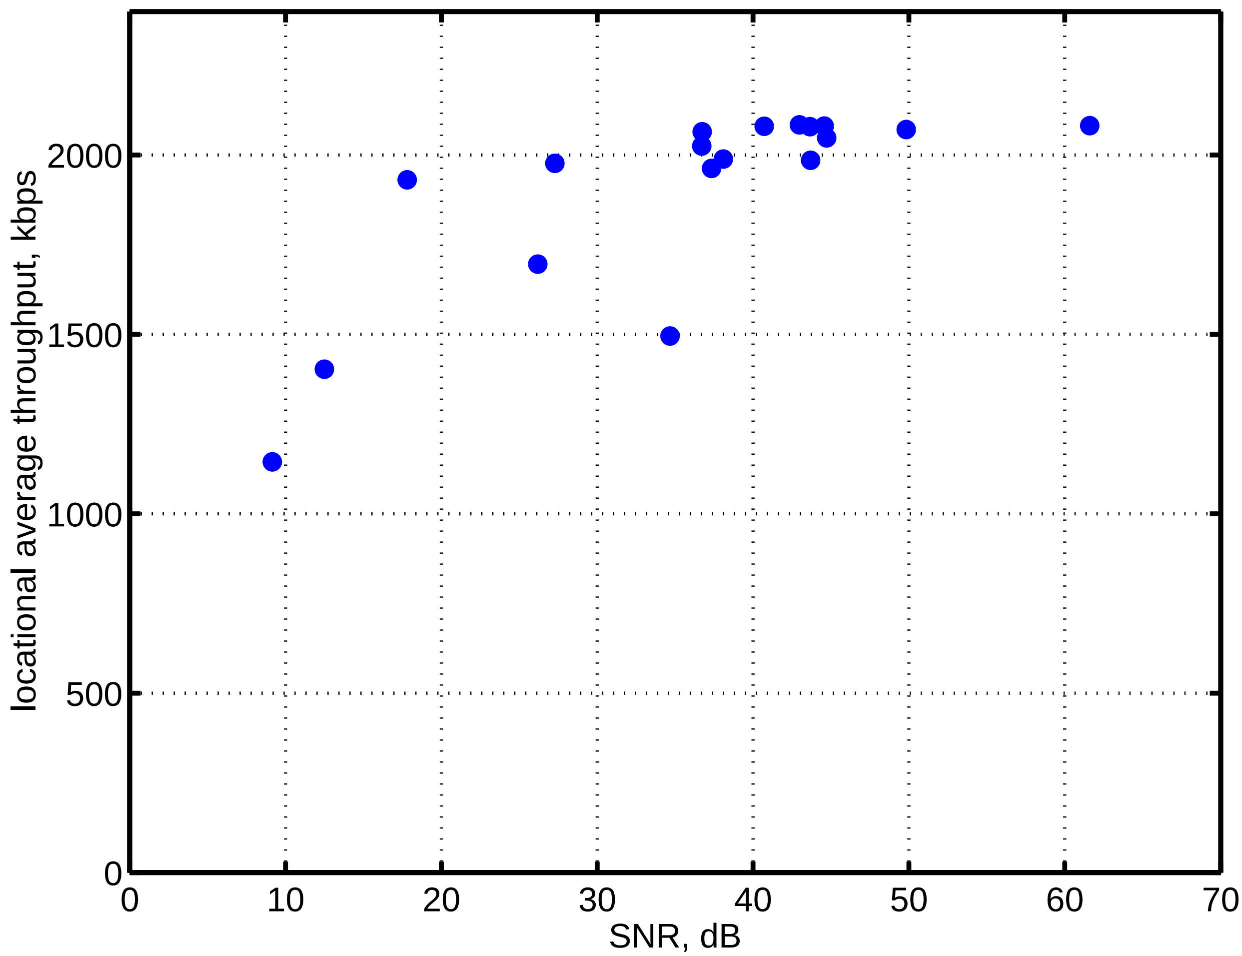

Figure 5.5: All measurements averaged for all directions at each measurement location for Scenario A, 3Com System

5.2.6 3Com, Scenario A, Window Averaged Data

The window averaging used in Figure 5.6 very clearly shows the validity of a exponential model of throughput. The 3Com data has a slight drop-off in throughput above the critical SNR value, but drops of more quickly below the critical value, making an exponential model of throughput a good choice for this data.

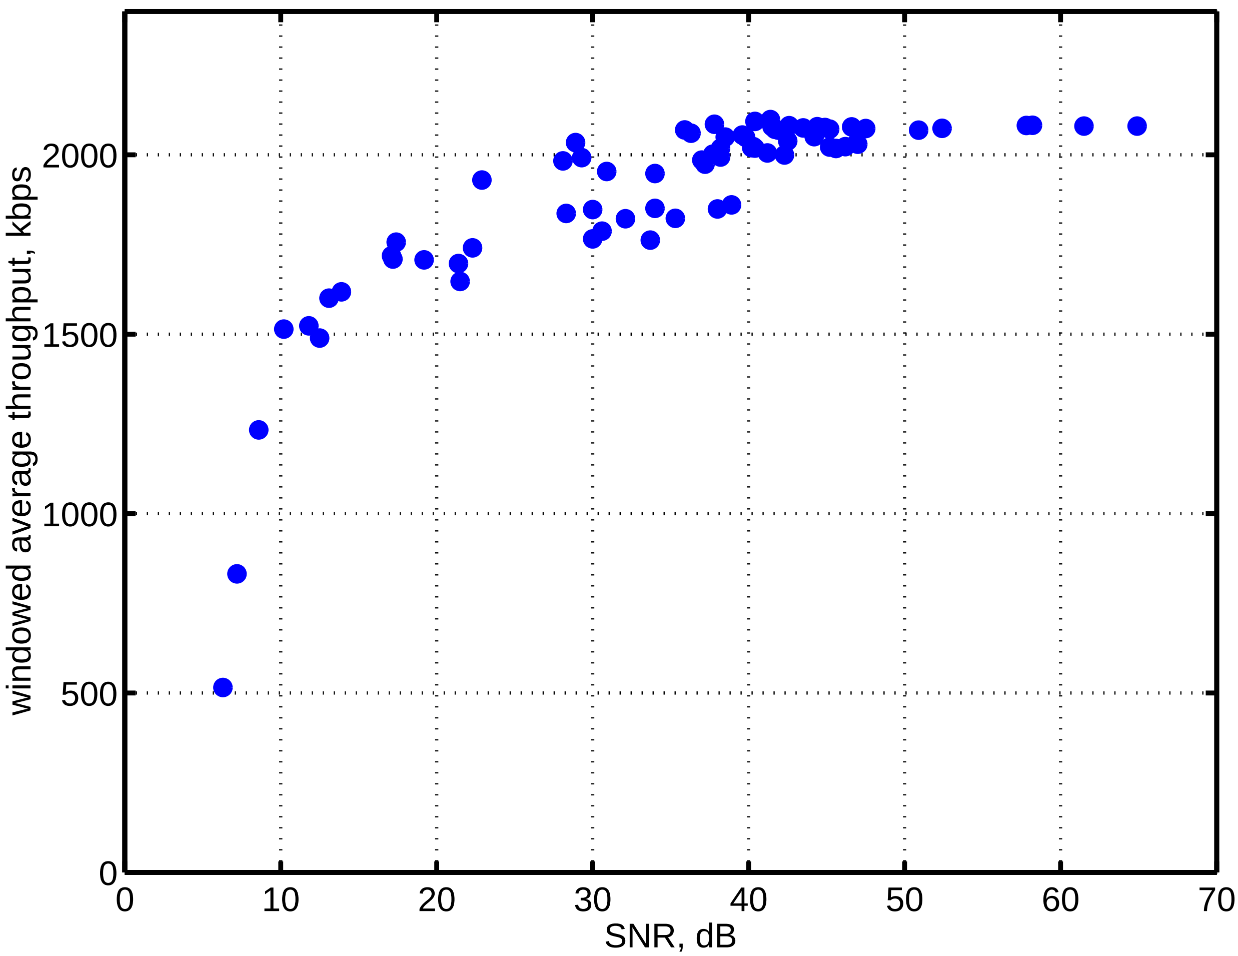

Figure 5.6: All measurements averaged over similar signal strength (i.e. RSSI) locations for Scenario A, 3Com System

5.3 Scenario B Data - Two Clients

5.3.1 WaveLAN, Scenario B, All Data

Figure 5.7 shows all of the two client measurements made with the WaveLAN system. Both client's individual throughput is plotted in the graph. The figure shows a similar trend as the single client WaveLAN data. However, the two client data shows that a number of points show a lower than expected throughput and the overall throughput is lower than that of the single client data.

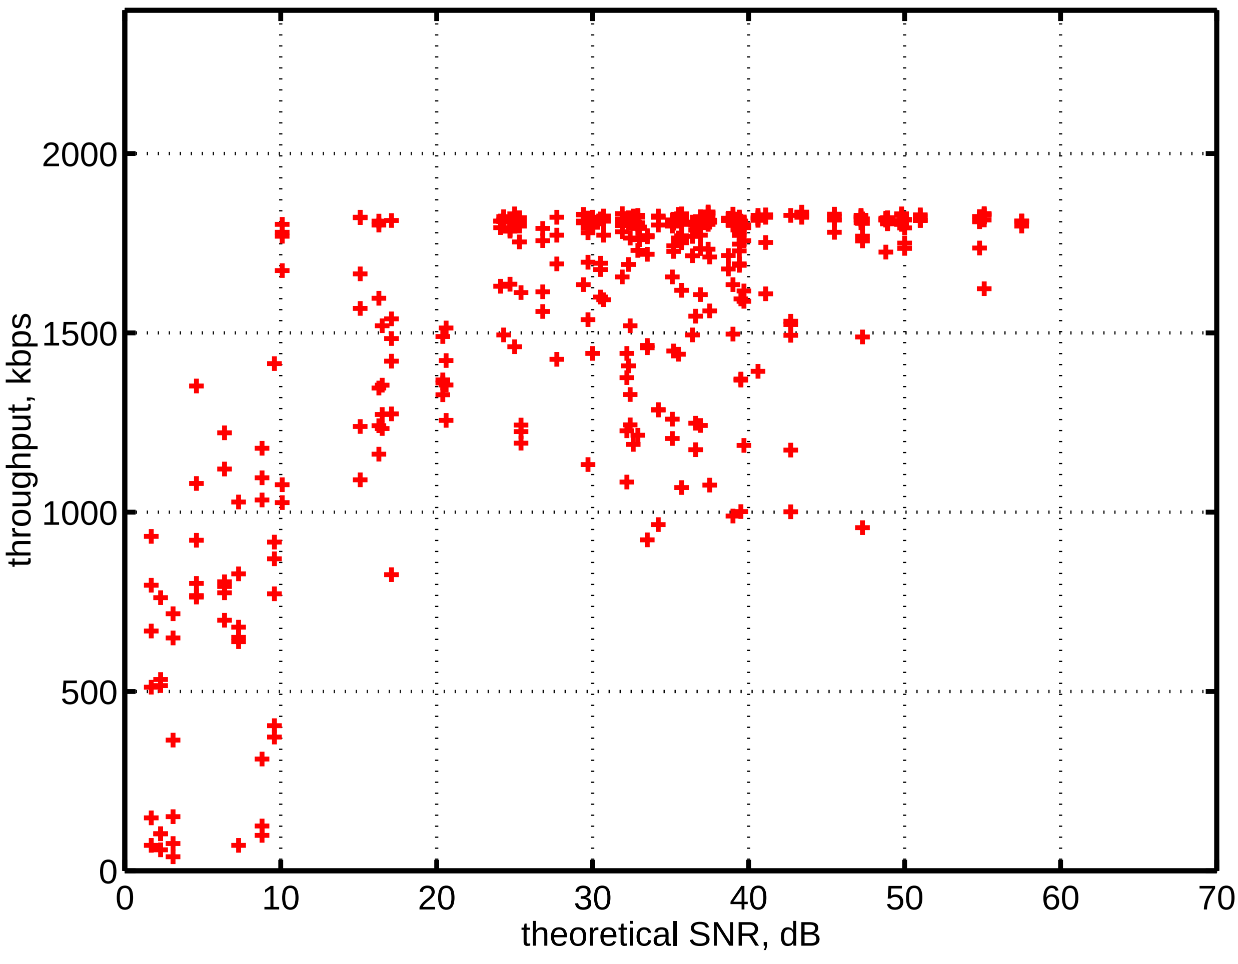

Figure 5.7: All measurements for Scenario B, WaveLAN System

5.3.2 WaveLAN, Scenario B, Spatially Averaged Data

Figure 5.8 shows a similar trend as Figure 5.7 but with a little less variation due to the spatial averaging.

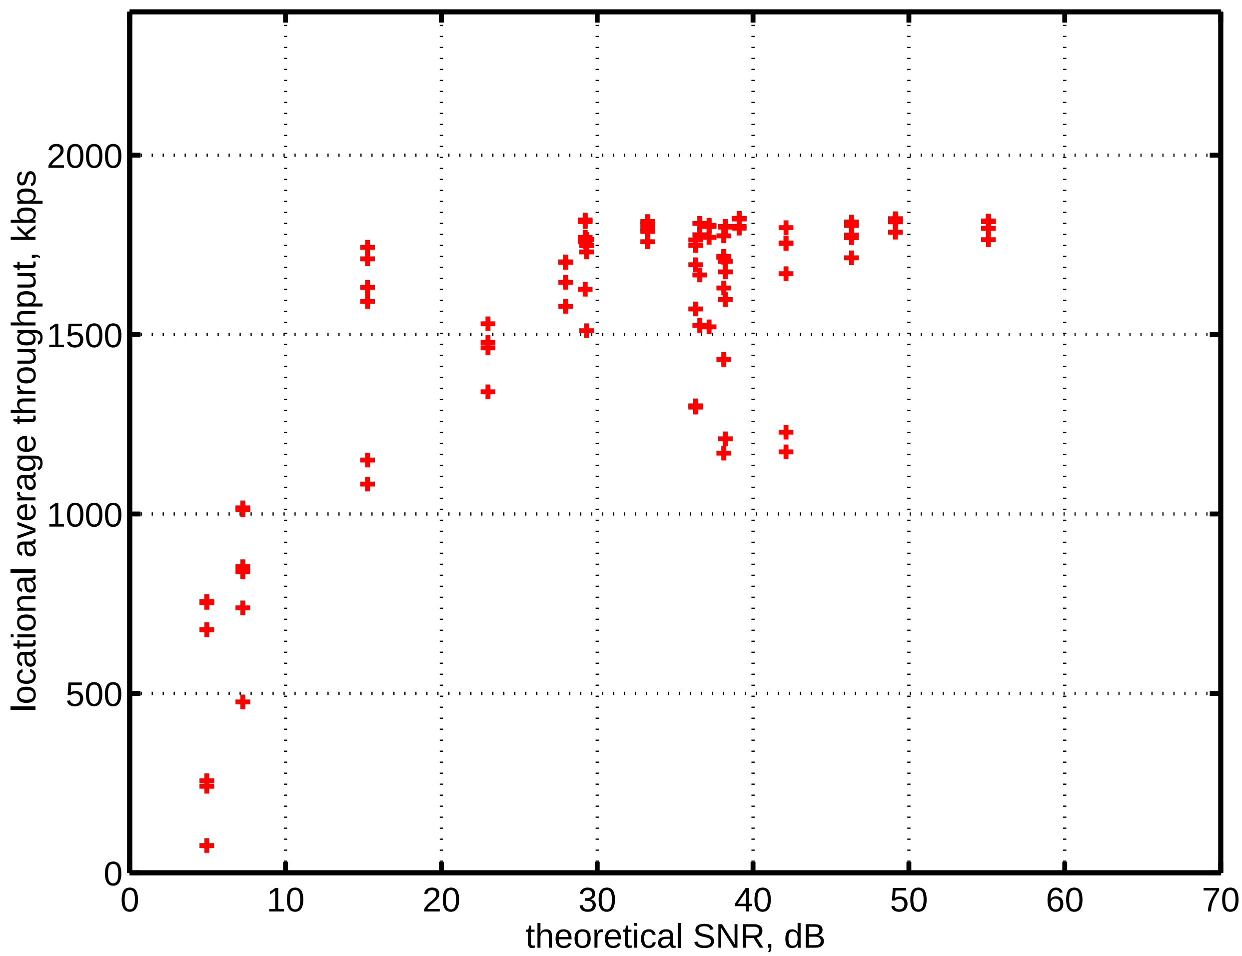

Figure 5.8: All measurements averaged for all directions at each measurement location for Scenario B, WaveLAN System

5.3.3 WaveLAN, Scenario B, Window Averaged Data

Figure 5.9 clearly shows the presence of a critical SNR level for the data. Additionally, the figure shows the same overall trend of a loss in throughput as was indicated in Figures 5.7 and 5.8.

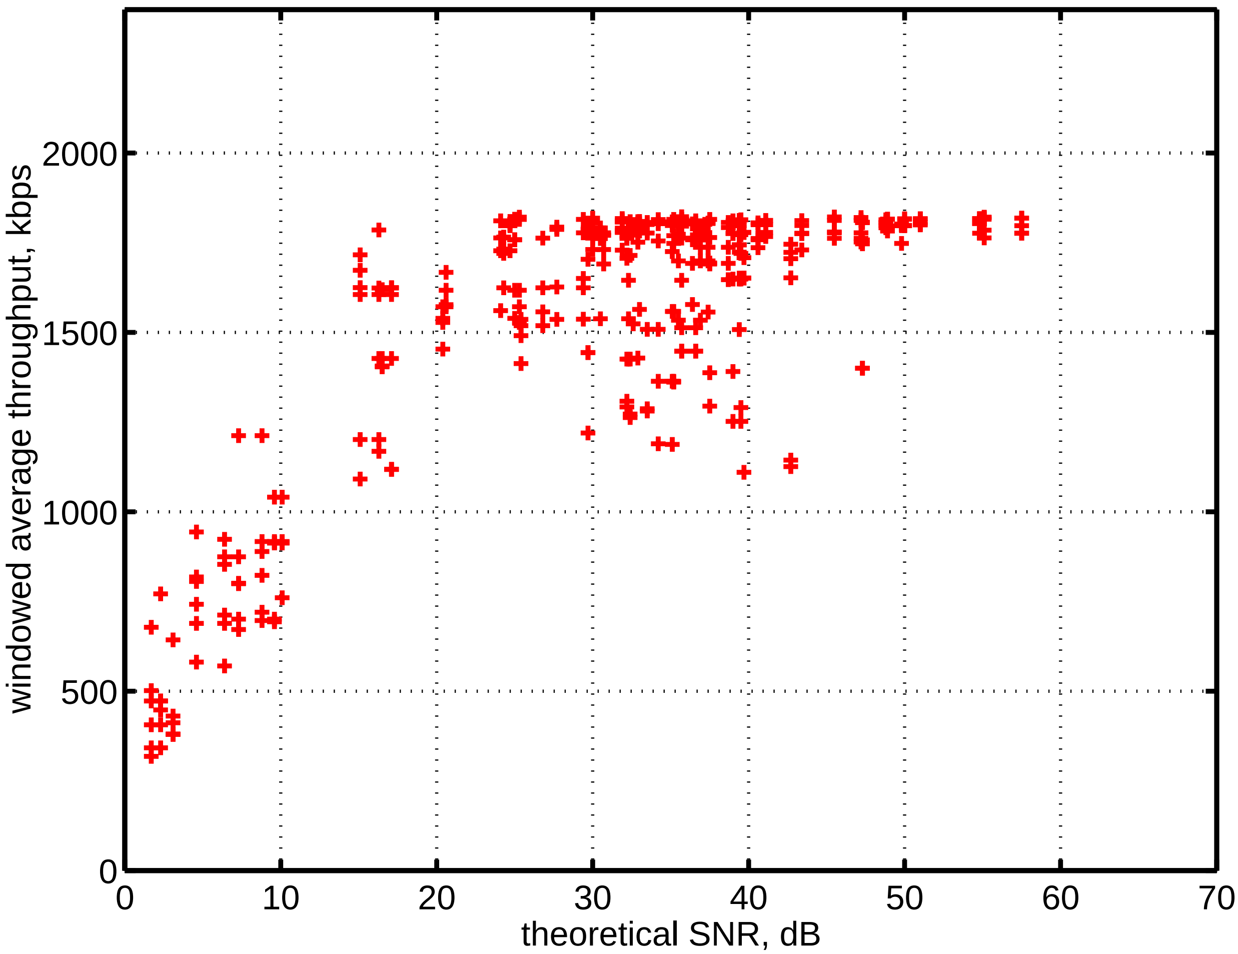

Figure 5.9: All measurements averaged over similar signal strength (i.e. RSSI) locations for Scenario B, WaveLAN System

5.3.4 3Com, Scenario B, All Data

Figure 5.10 shows all of the two client data points for both 3Com clients. In the figure, the drop-off in throughput below a certain critical SNR point is evident, however, there is a wide range in the throughput performance for all measured points.

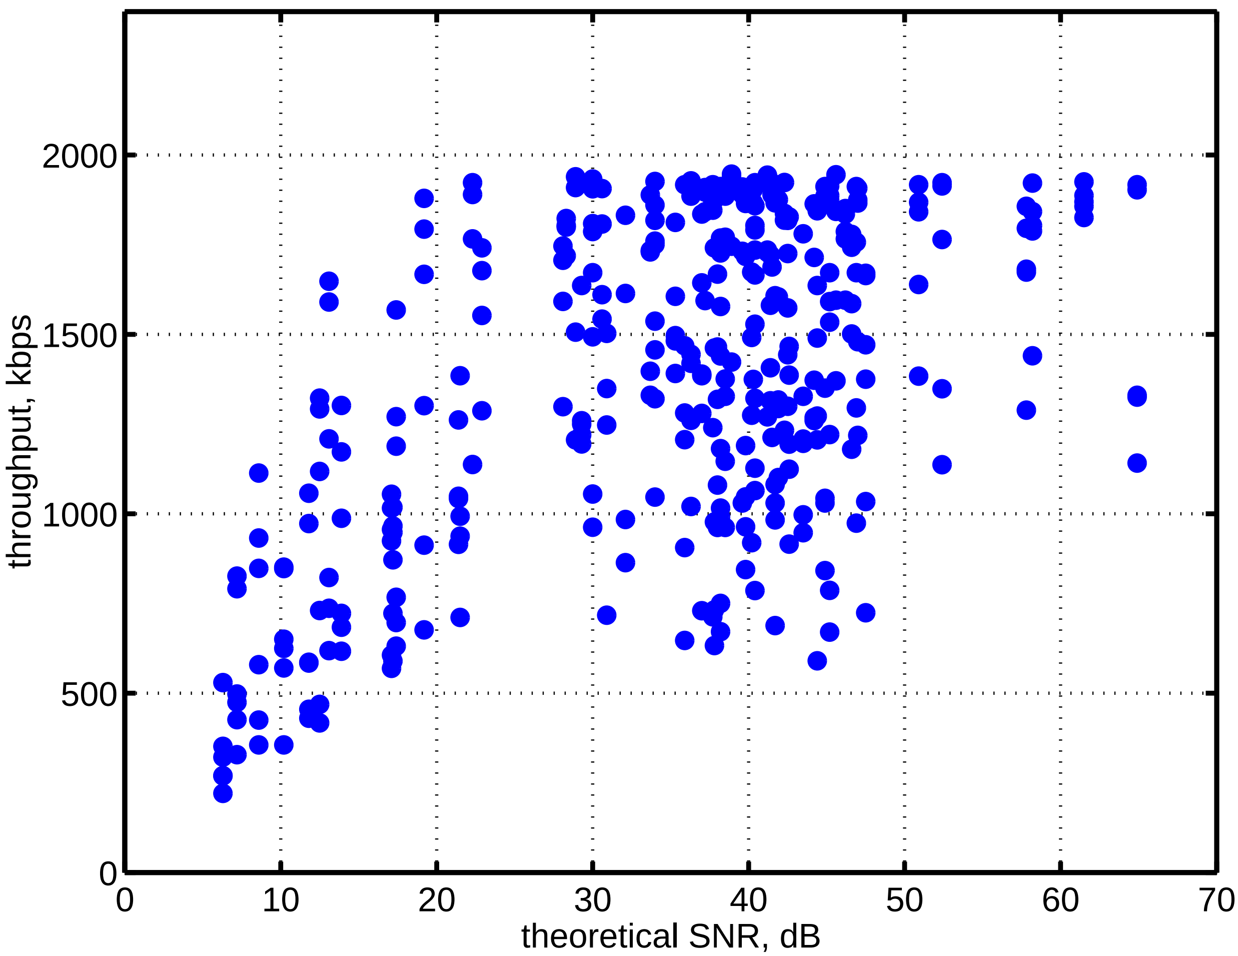

Figure 5.10: All measurements for Scenario B, 3Com System

5.3.5 3Com, Scenario B, Spatially Averaged Data

Figure 5.11 shows a smaller range in the variation of the throughput experienced than shown in Figure 5.10. The same overall trend found in the single client data is also present in Figure 5.11.

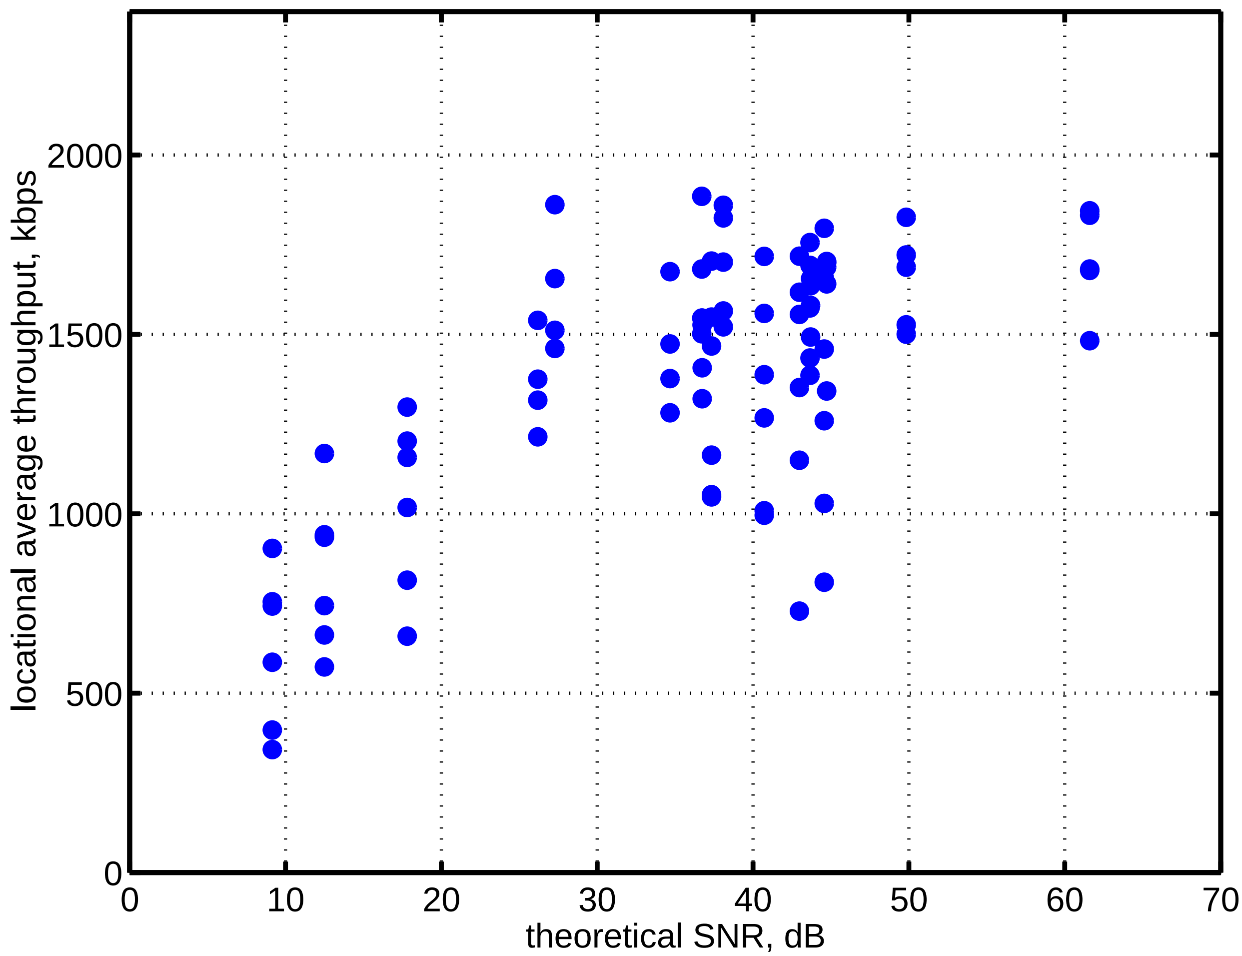

Figure 5.11: All measurements averaged for all directions at each measurement location for Scenario B, 3Com System

5.3.6 3Com, Scenario B, Window Averaged Data

In Figure 5.12, the variation in throughput is still present. However, the exponential trend in the throughput versus the SNR seen in Figure 5.6 is also present.

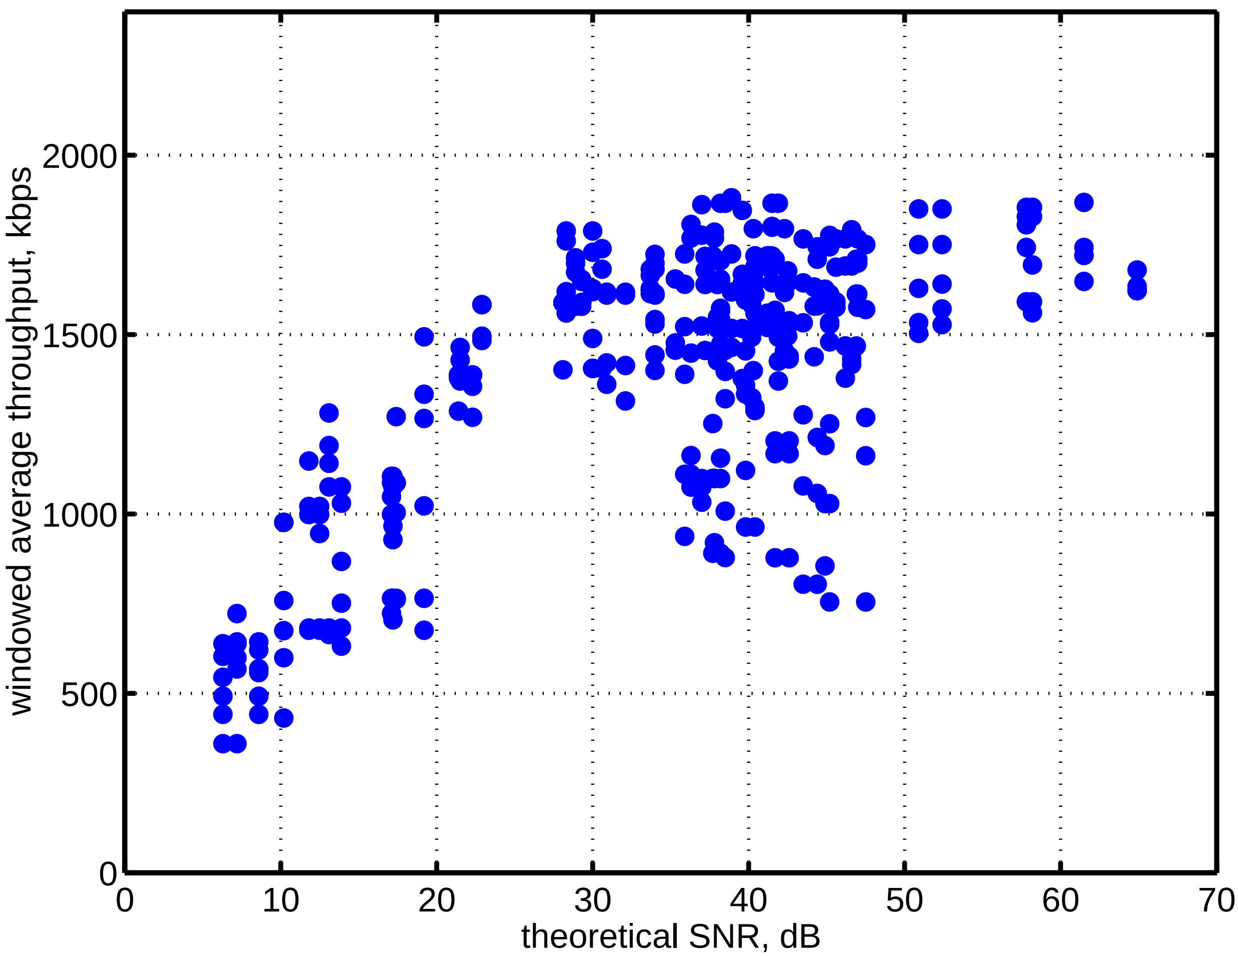

Figure 5.12: All measurements averaged over similar signal strength (i.e. RSSI) locations for Scenario B, 3Com System

5.4 Scenario C Data - One Client with Interference

5.4.1 WaveLAN, Scenario C, All Data

Figure 5.13 presents all of the single client with interference data for the WaveLAN system. A definite loss in throughput can be seen in the figure. However, a new tendency to operate at a single throughput level can be seen in the data. This is probably due to the wireless LAN dropping to the 5.5 Mbps transmission setting. Other losses could be due to bit errors and transmission delays.

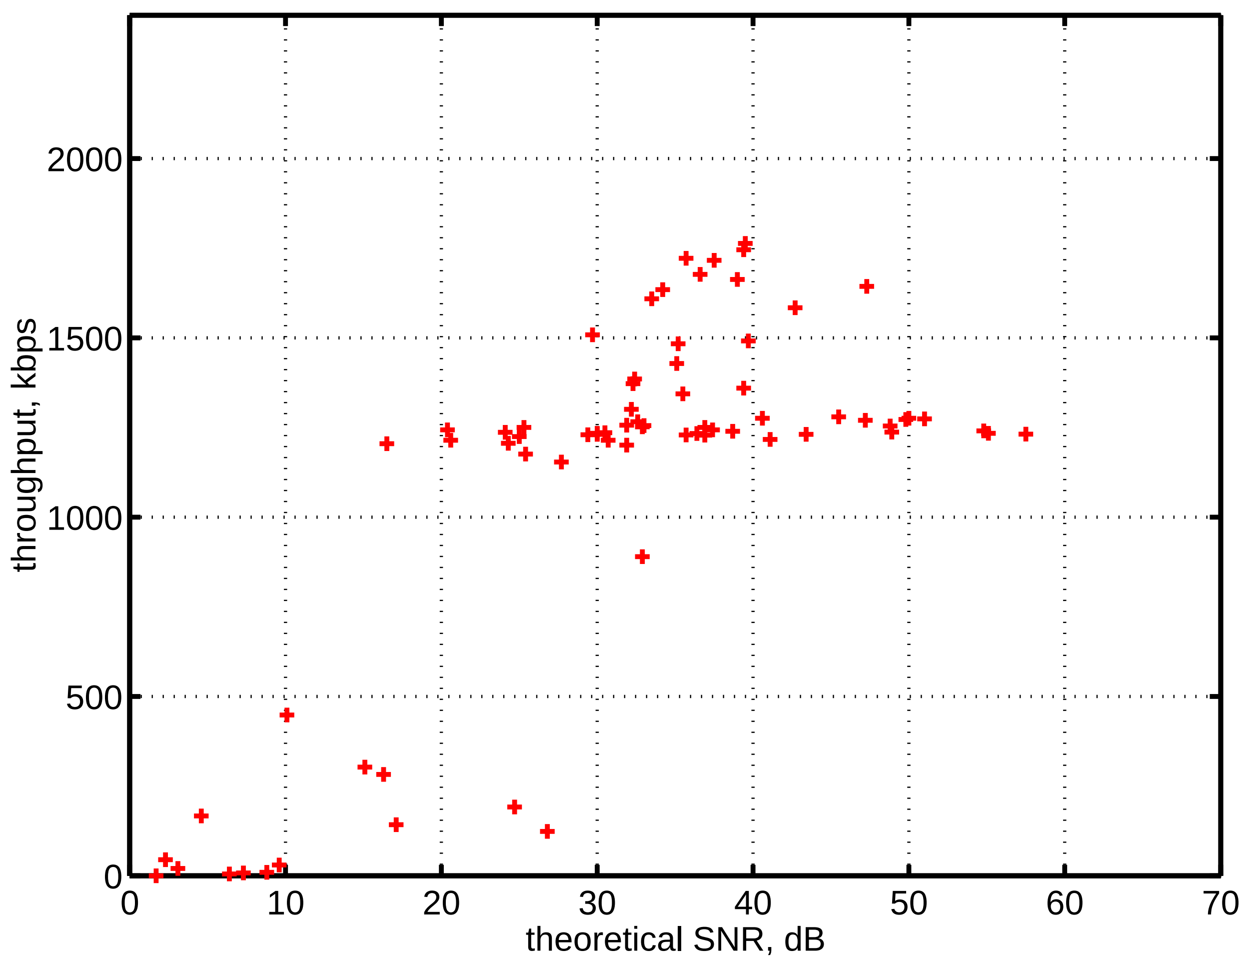

Figure 5.13: All measurements for Scenario C, WaveLAN System

5.4.2 WaveLAN, Scenario C, Spatially Averaged Data

Figure 5.14 shows the overall trend in the WaveLAN data shown in Figure 5.13. However, the tendency to operate at a single throughput level is slightly more evident in Figure 5.14.

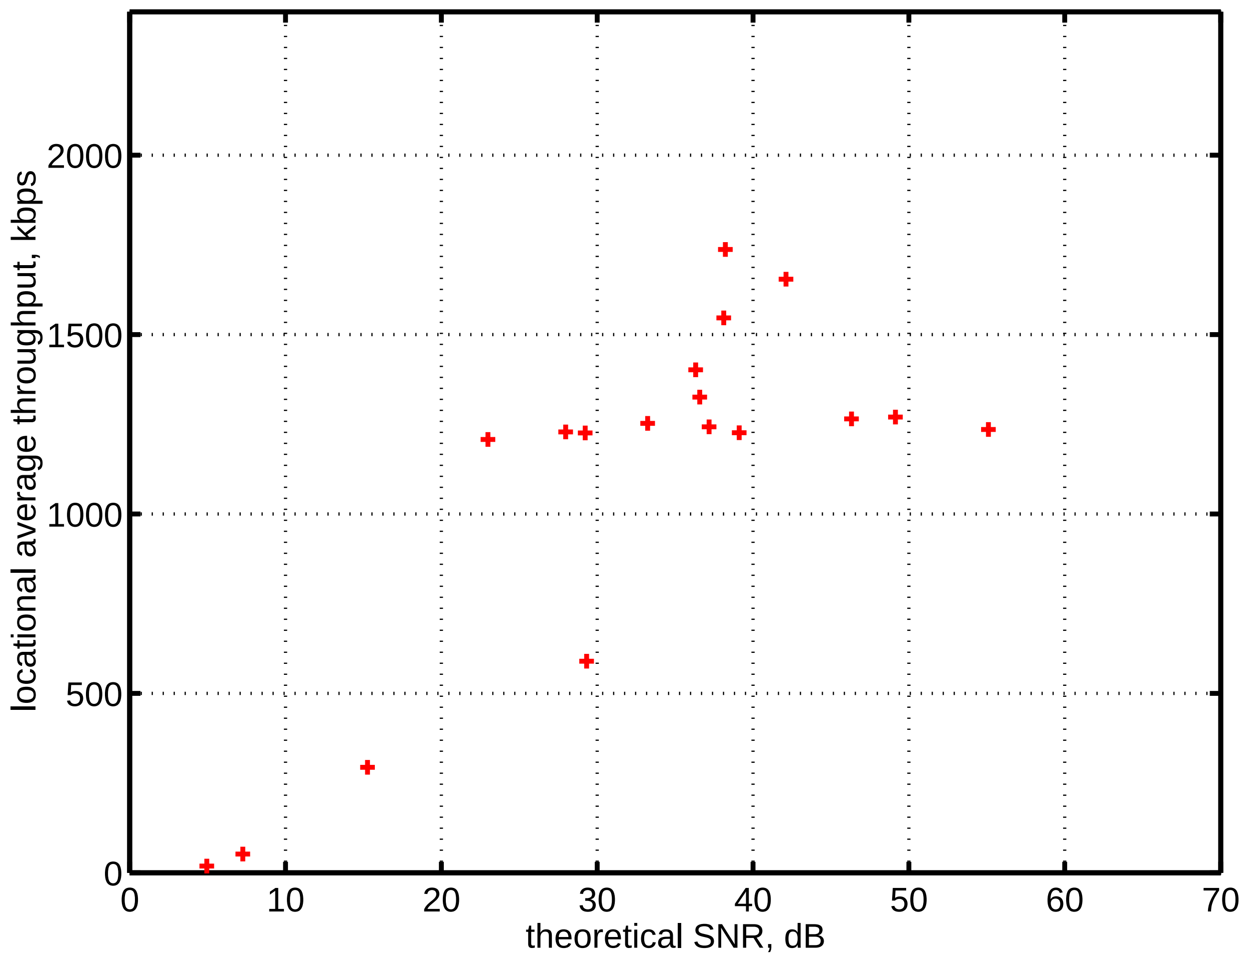

Figure 5.14: All measurements averaged for all directions at each measurement location for Scenario C, WaveLAN System

5.4.3 WaveLAN, Scenario C, Window Averaged Data

The window averaging has somewhat distorted the overall trend in the WaveLAN data. Figure 5.15 has a higher throughput for mid-range SNR levels due to the low SIR levels measured at these points. This results in the misleading trend of the data in Figure 5.15.

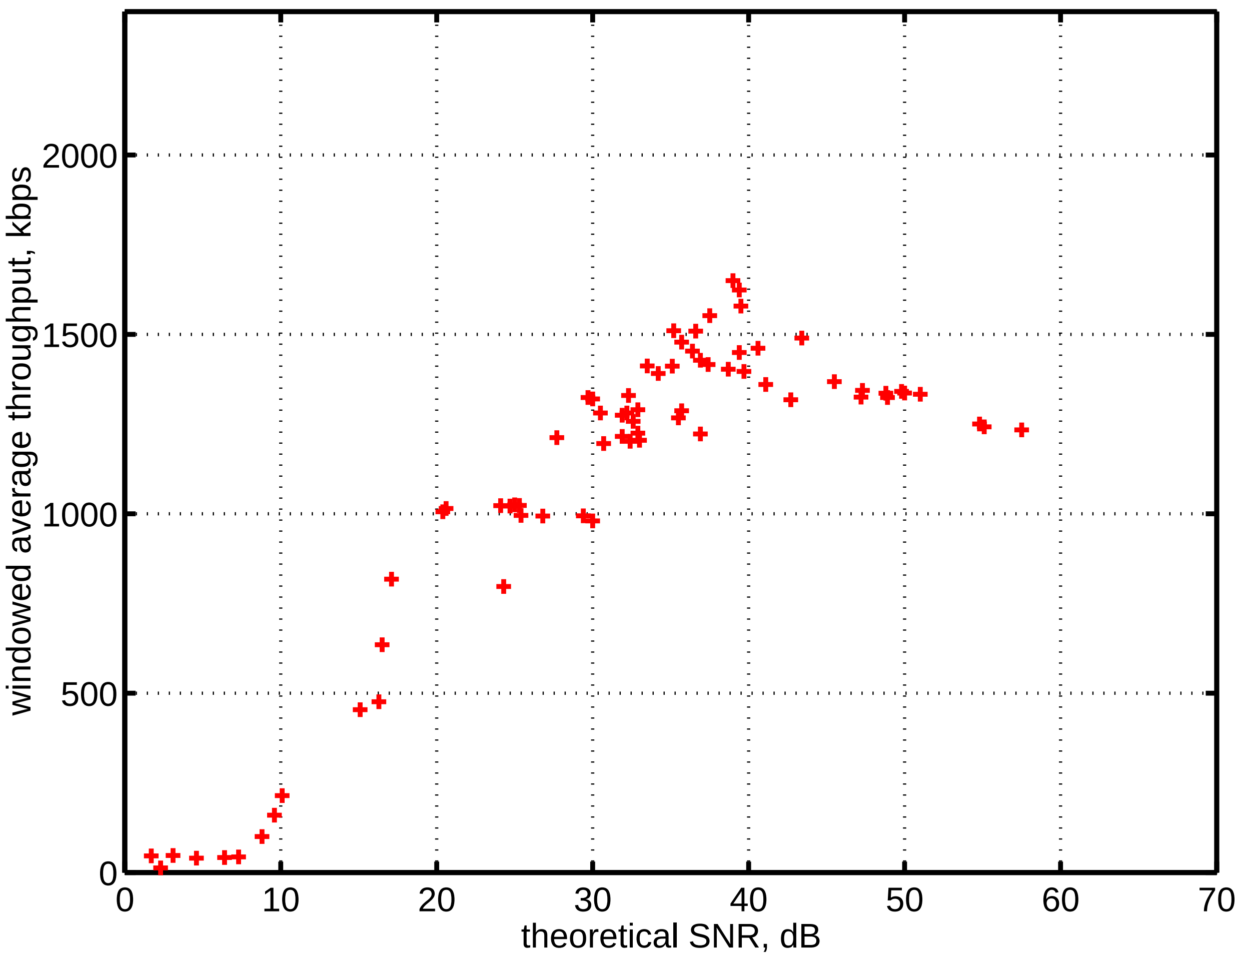

Figure 5.15: All measurements averaged over similar signal strength (i.e. RSSI) locations for Scenario C, WaveLAN System

5.4.4 3Com, Scenario C, All Data

Figure 5.16 presents all of the single client with interference data measured for the 3Com system. The data shown in the figure indicates that all measured locations appear to have been influenced by the interference. All measurements have dropped in throughput while maintaining the same basic trend as was found in the single client (without interference) data shown in Figure 5.4.

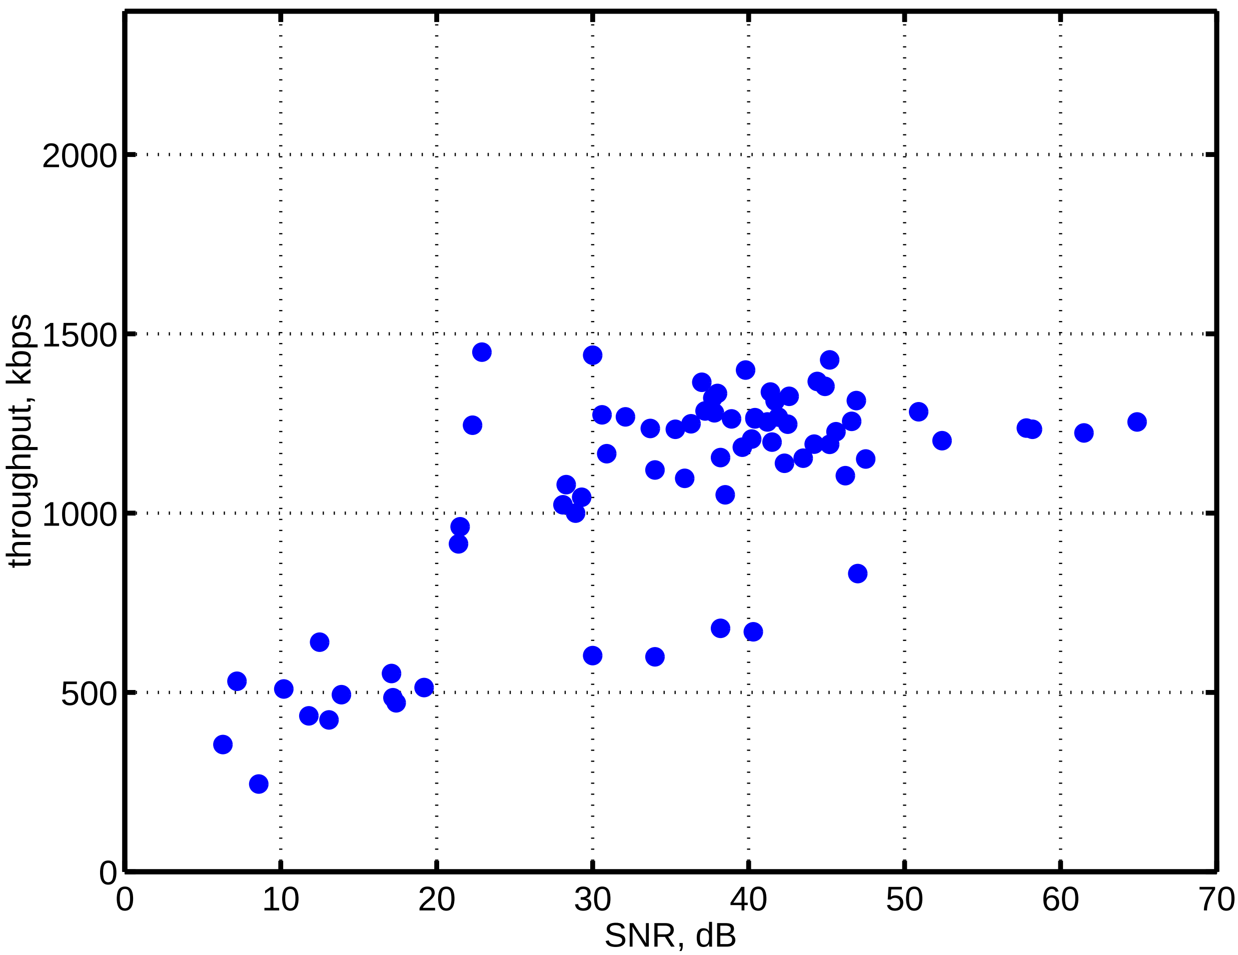

Figure 5.16: All measurements for Scenario C, 3Com System

5.4.5 3Com, Scenario C, Spatially Averaged Data

Figure 5.17 shows the data tends to be centered about two different levels. This is probably due to the wireless LAN operating at the 5.5 and 2 Mbps transmission levels.

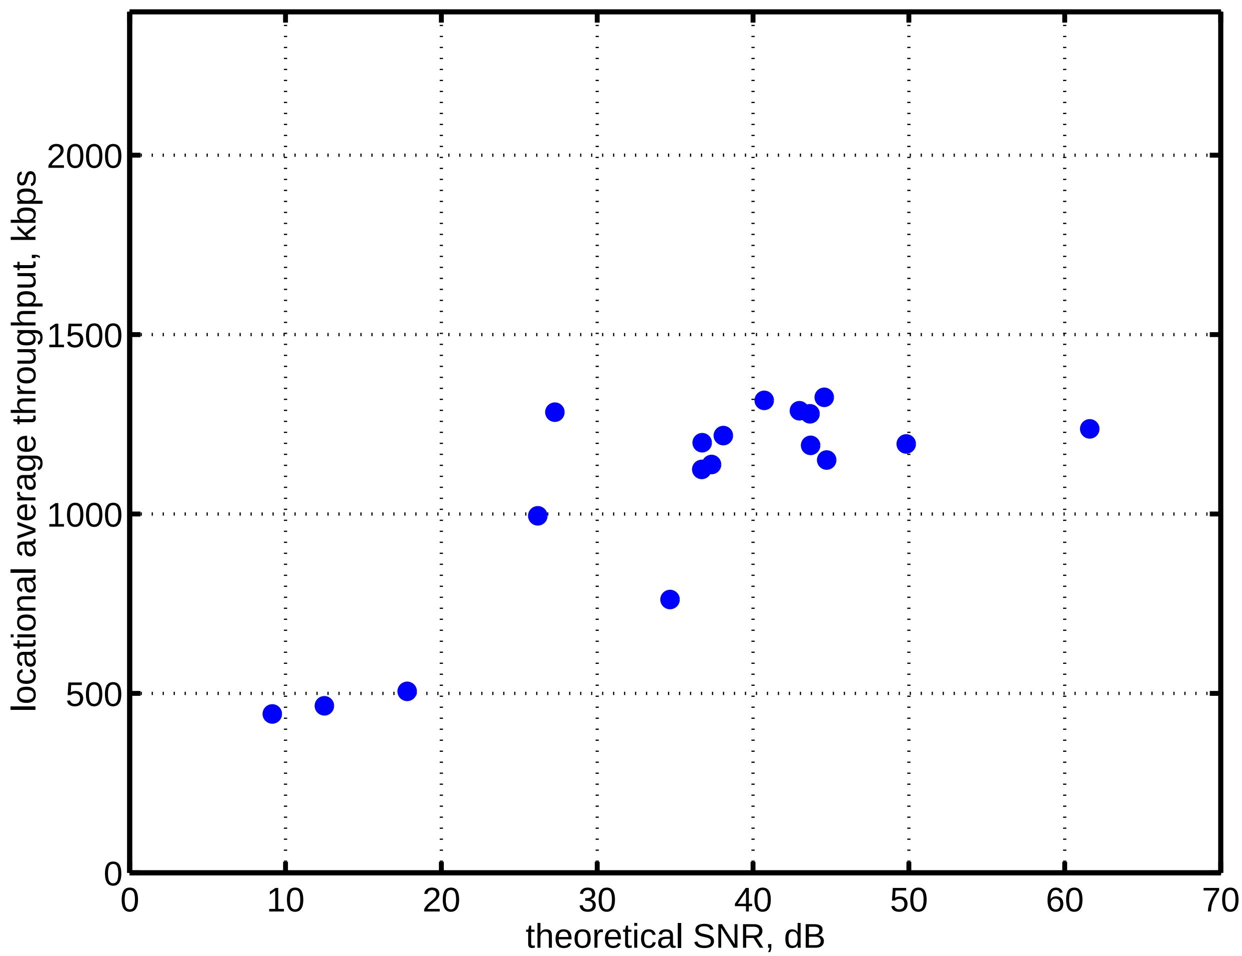

Figure 5.17: All measurements averaged for all directions at each measurement location for Scenario C, 3Com System

5.4.6 3Com, Scenario C, Window Averaged Data

Figure 5.18, like Figure 5.17, shows the throughput dropping to two different operating levels. In Figure 5.18, the data can be seen to be trending downward in steps, probably as the wireless LAN system drops to the 5.5 and 2 Mbps transmission settings.

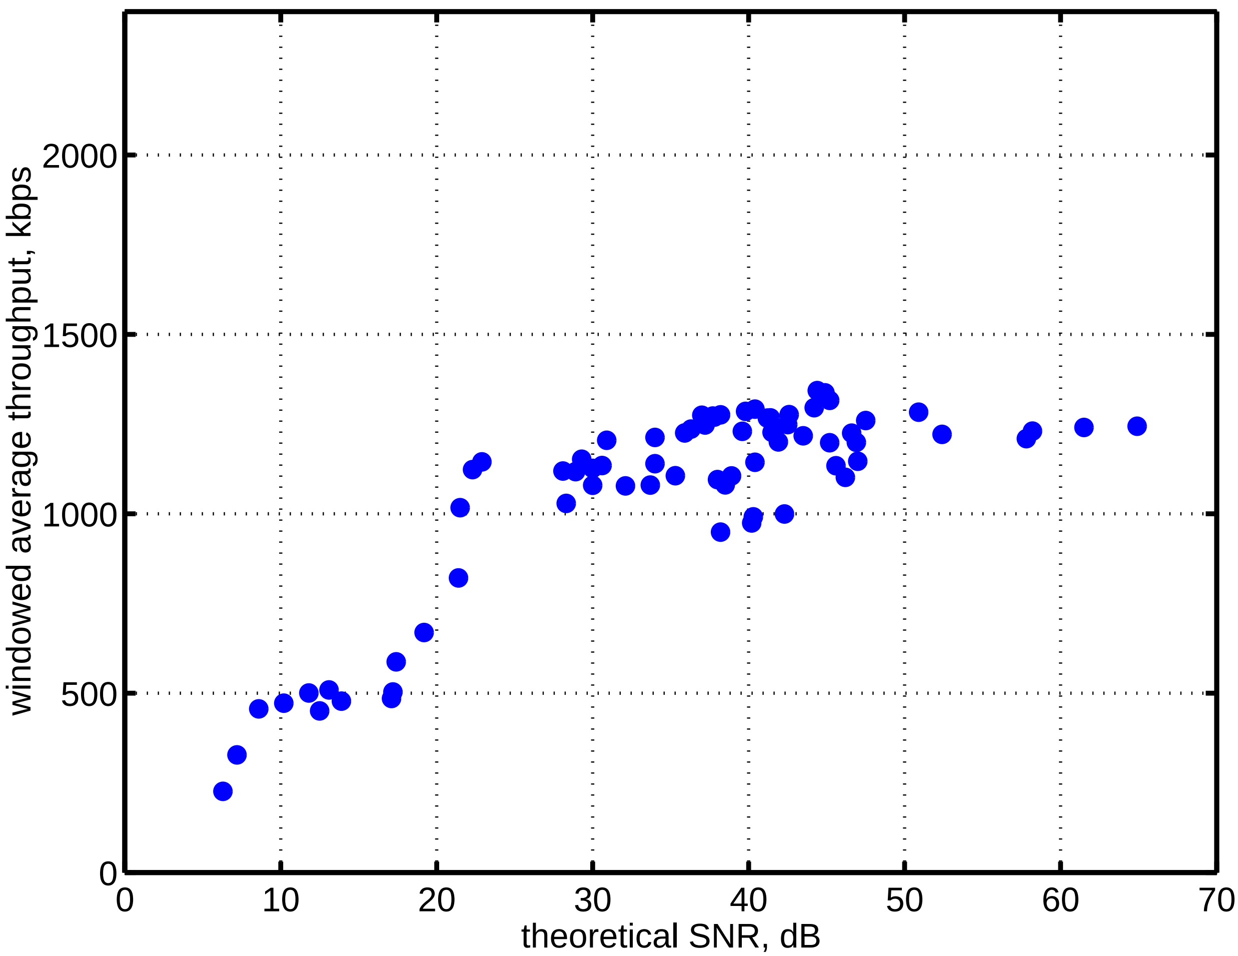

Figure 5.18: All measurements averaged over similar signal strength (i.e. RSSI) locations for Scenario C, 3Com System

5.5 Scenario D Data - Two Clients with Interference

5.5.1 WaveLAN, Scenario D, All Data

Figure 5.19 shows all of the WaveLAN two client with interference measurements for both clients. This data shows the variation in throughput found in Figure 5.7 but also shows the tendency to operate at two different throughput levels as was seen in Figure 5.13. These two throughput levels are probably due to the wireless LAN using the 11 and 5.5 Mbps transmission settings.

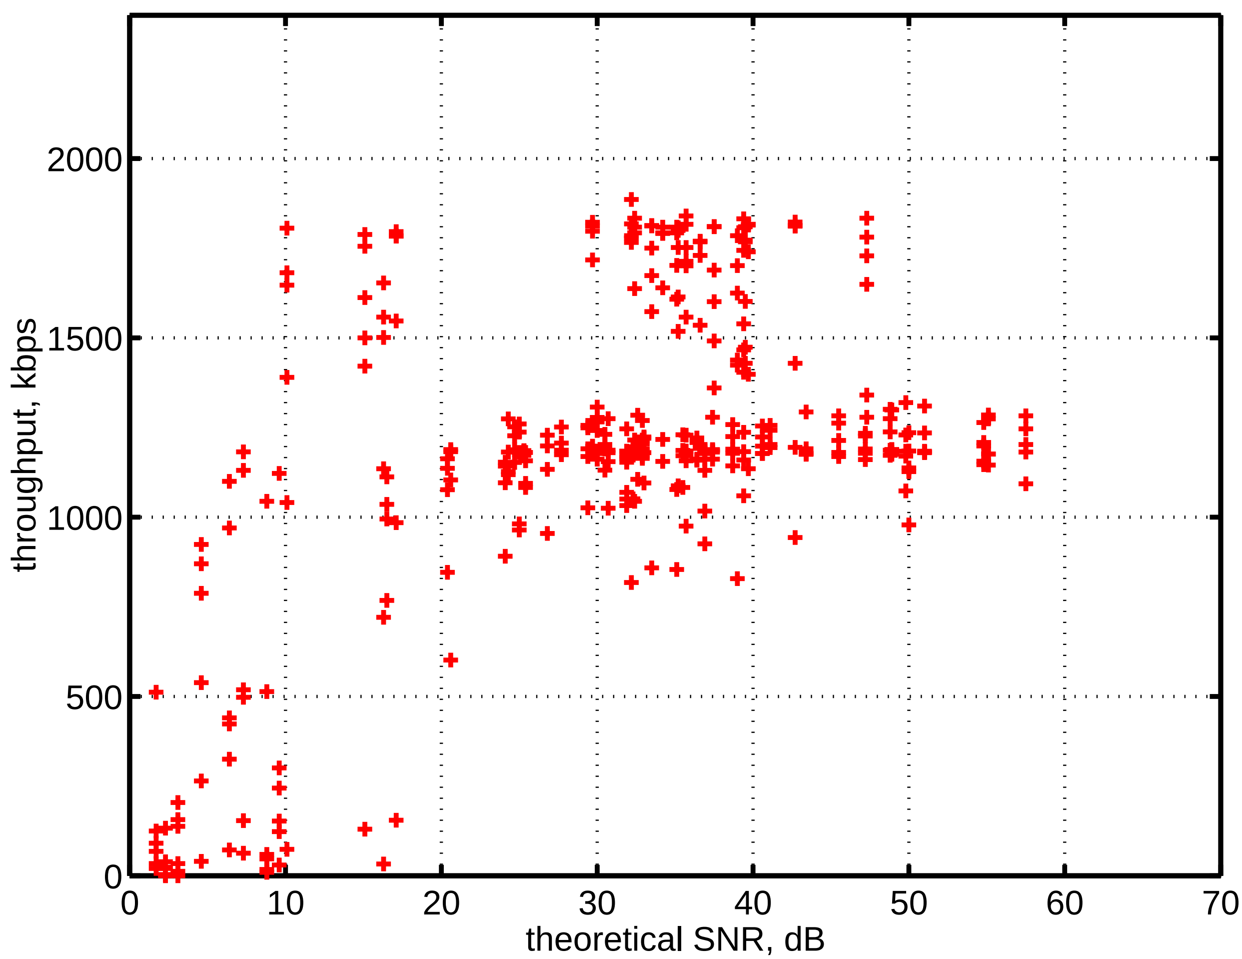

Figure 5.19: All measurements for Scenario D, WaveLAN System

5.5.2 WaveLAN, Scenario D, Spatially Averaged Data

Figure 5.20 indicates that the throughput has the similar trend as the single client data in which the throughput is constant above a critical SNR value and drops steadily off below that level. This is only the basic trend, as several points do not follow this pattern. Further the constant throughput reached by the WaveLAN data probably corresponds to the wireless LAN operating at the 5.5 Mbps transmission level.

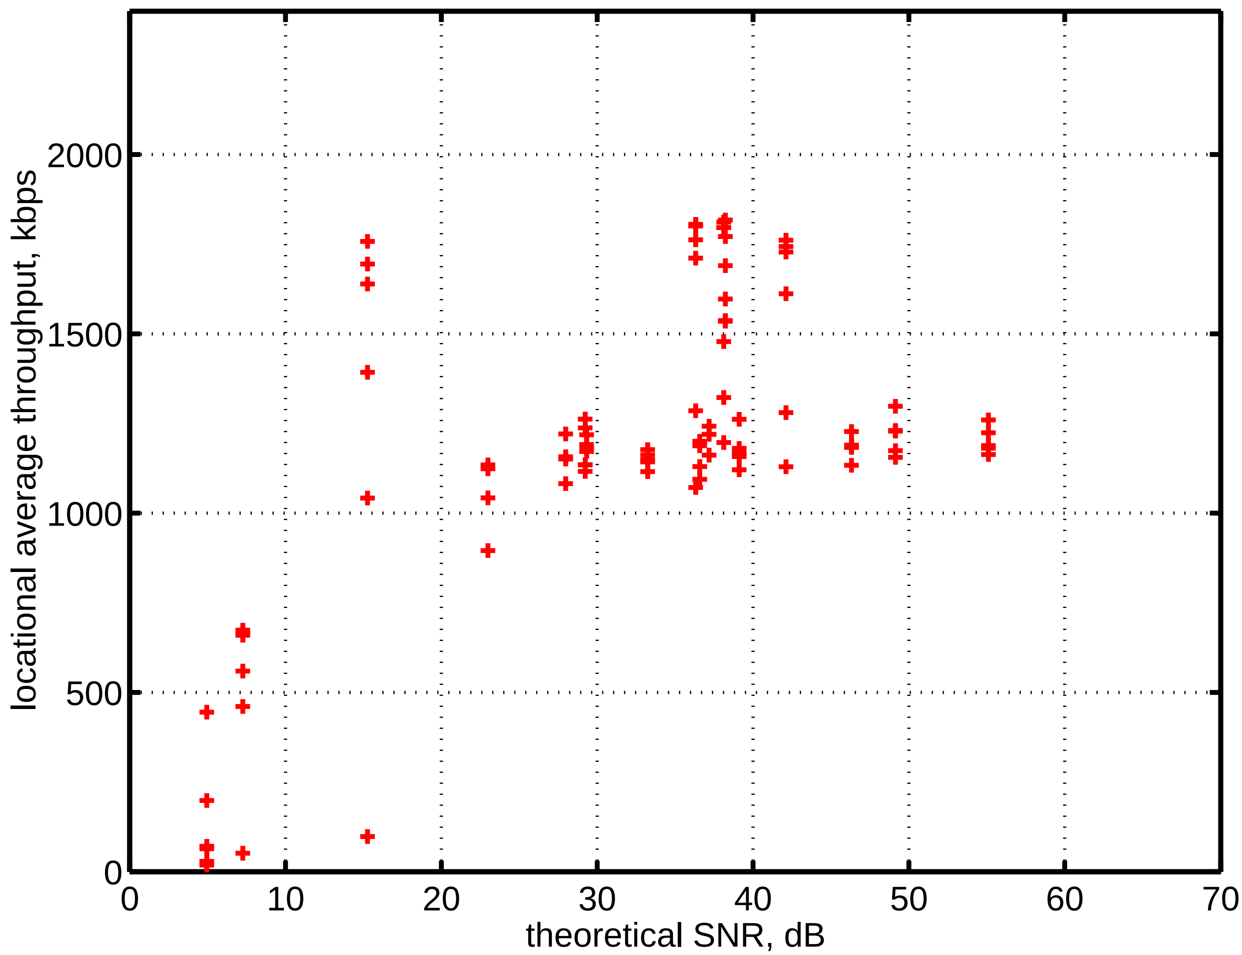

Figure 5.20: All measurements averaged for all directions at each measurement location for Scenario D, WaveLAN System

5.5.3 WaveLAN, Scenario D, Window Averaged Data

Figure 5.21 has tended to smear the overall features of the performance of the WaveLAN system that was presented in Figure 5.19. However, the data does have similar aspects to Figure 5.20, though the trends are not as clear.

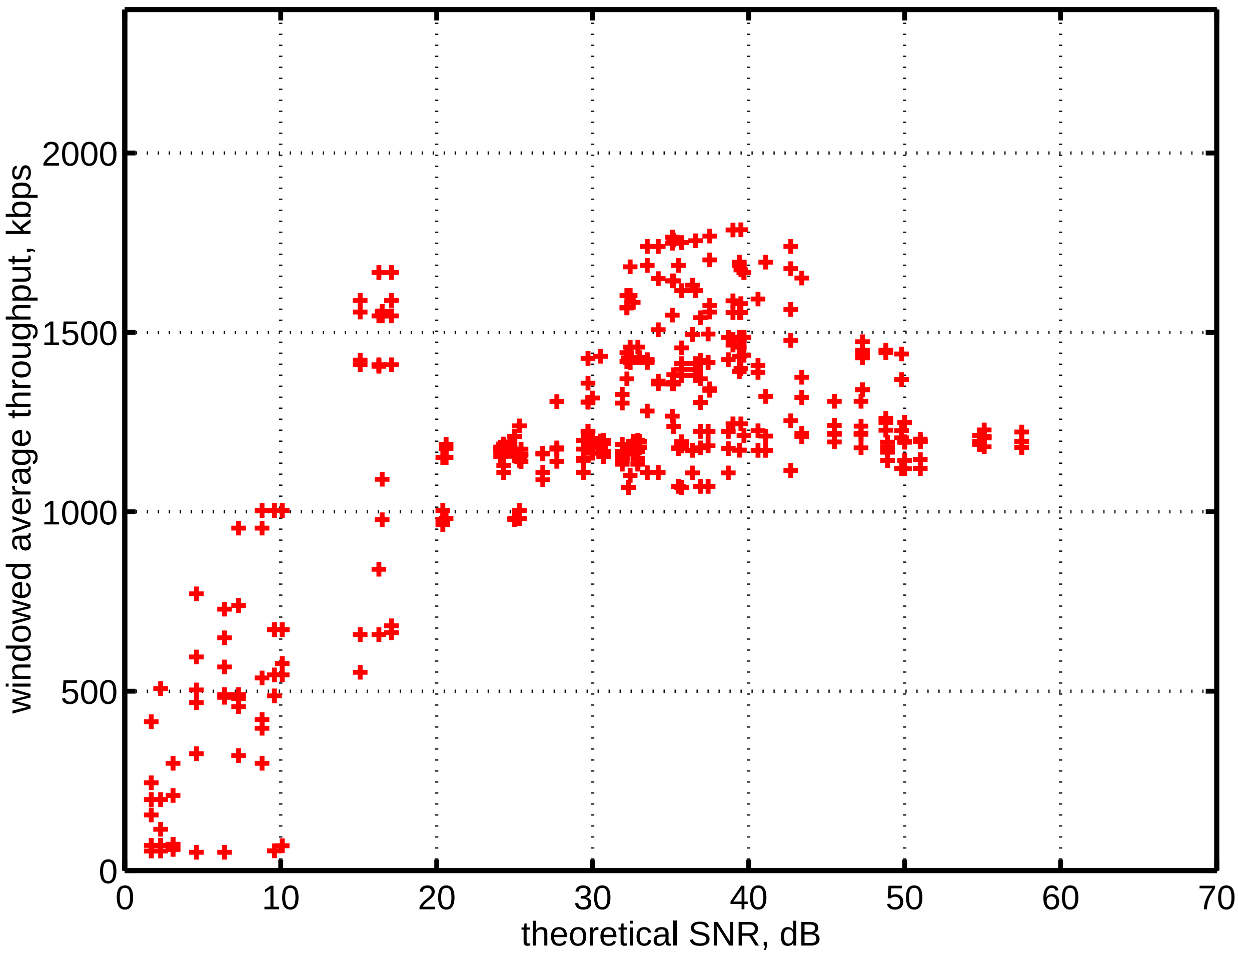

Figure 5.21: All measurements averaged over similar signal strength (i.e. RSSI) locations for Scenario D, WaveLAN System

5.5.4 3Com, Scenario D, All Data

Figure 5.22 presents the 3Com two client with interference data for both clients. The data appears to be a logical combination of the trends found in Figures 5.10 and 5.16. Figure 5.22 indicates the data has a variation in the throughput due to two users and a drop to lower throughput levels due to the presence of interference.

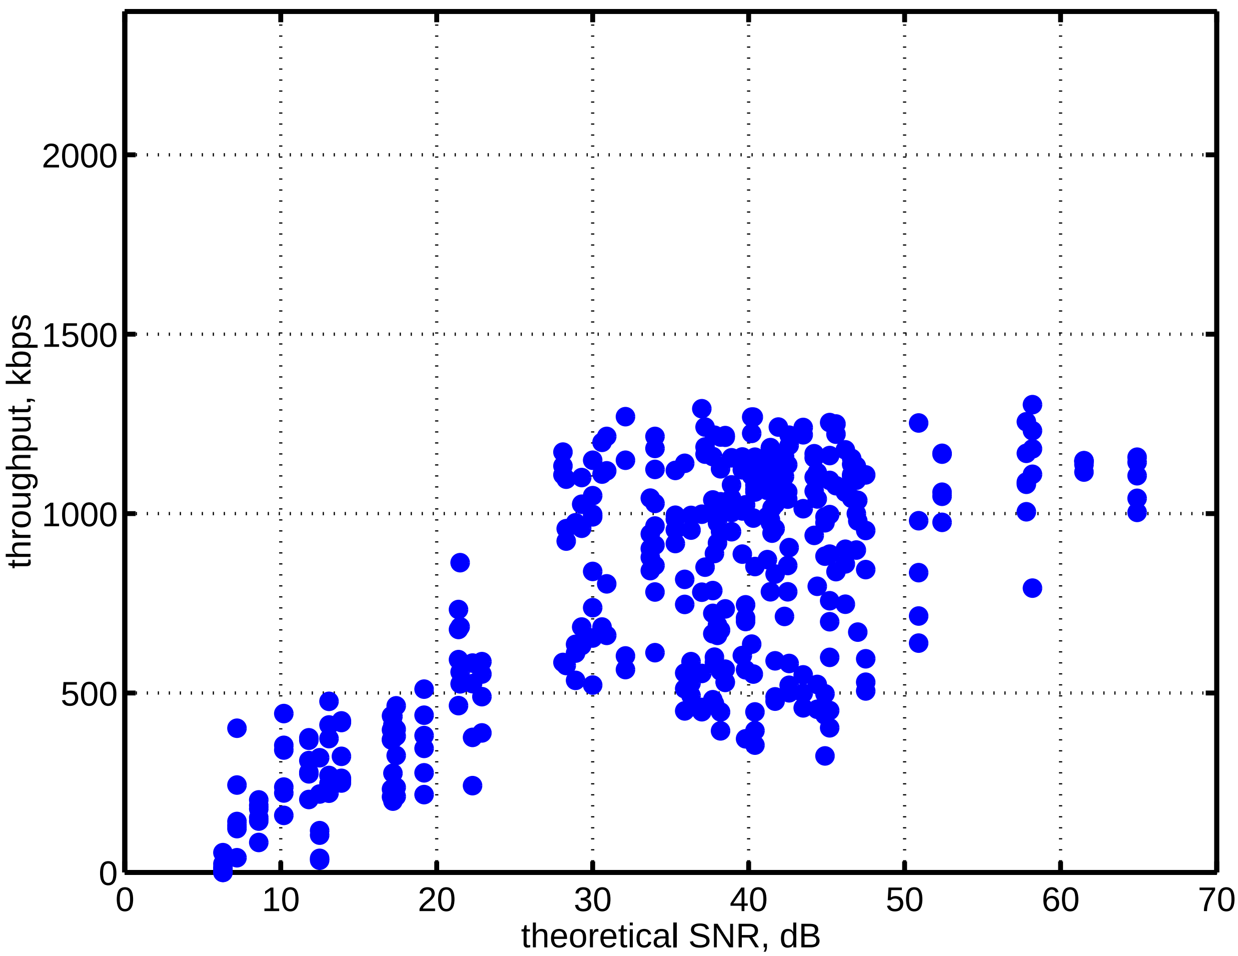

Figure 5.22: All measurements for Scenario D, 3Com System

5.5.5 3Com, Scenario D, Spatially Averaged Data

Figure 5.23 shows that the spatial averaging has smoothed out the trend of the measurements into a single smooth drop-off in throughput below a critical point and fairly constant above the critical point. The data, thus, has a similar trend as Figure 5.4 but with a lower constant throughput then the single client data.

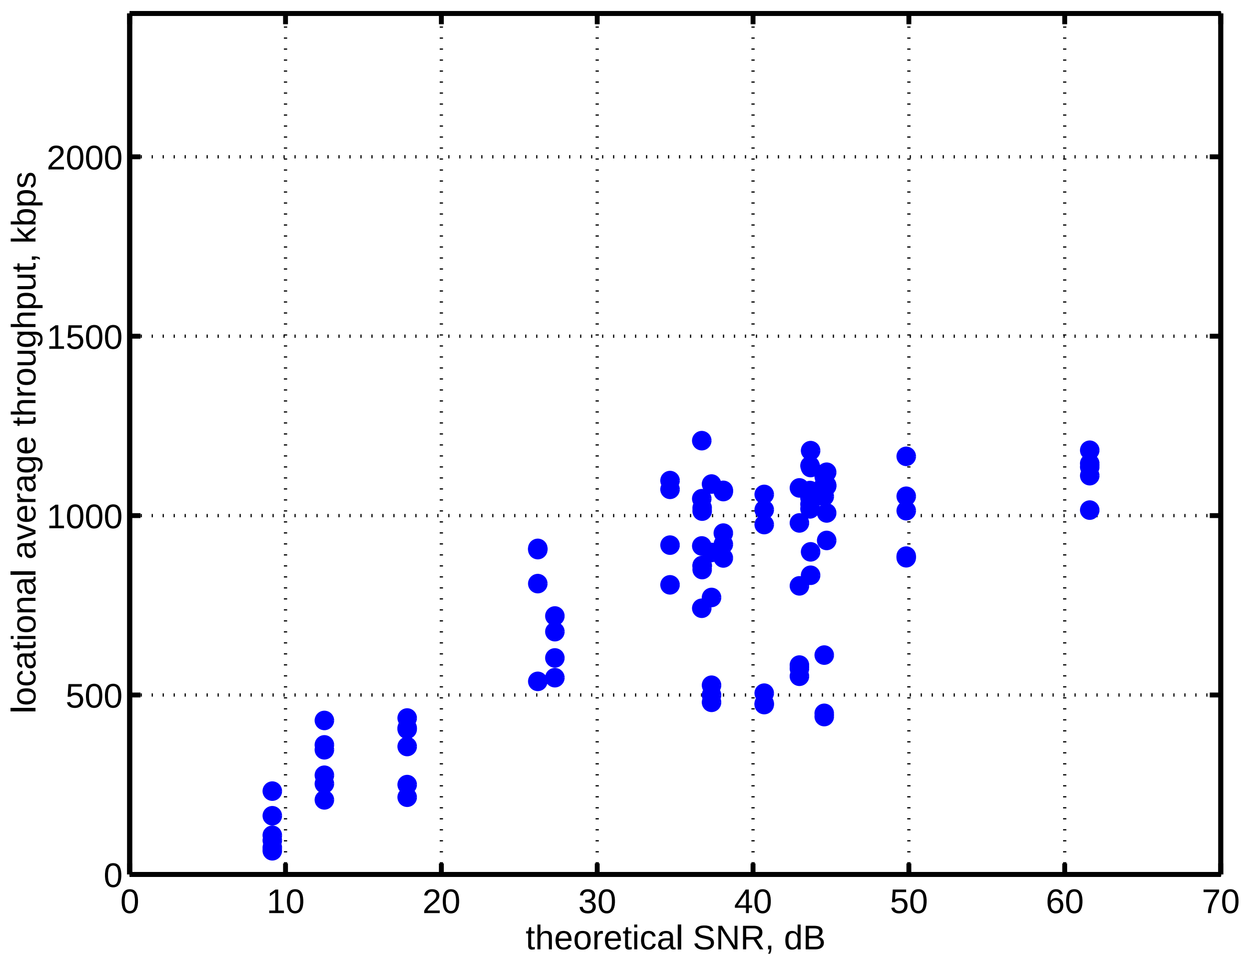

Figure 5.23: All measurements averaged for all directions at each measurement location for Scenario D, 3Com System

5.5.6 3Com, Scenario D, Window Averaged Data

Figure 5.24 appears to have also smoothed the data as in Figure 5.23. However, the data still appears to be centered about two different throughput levels, rather than smoothed into a single trend as was the case for Figure 5.23.

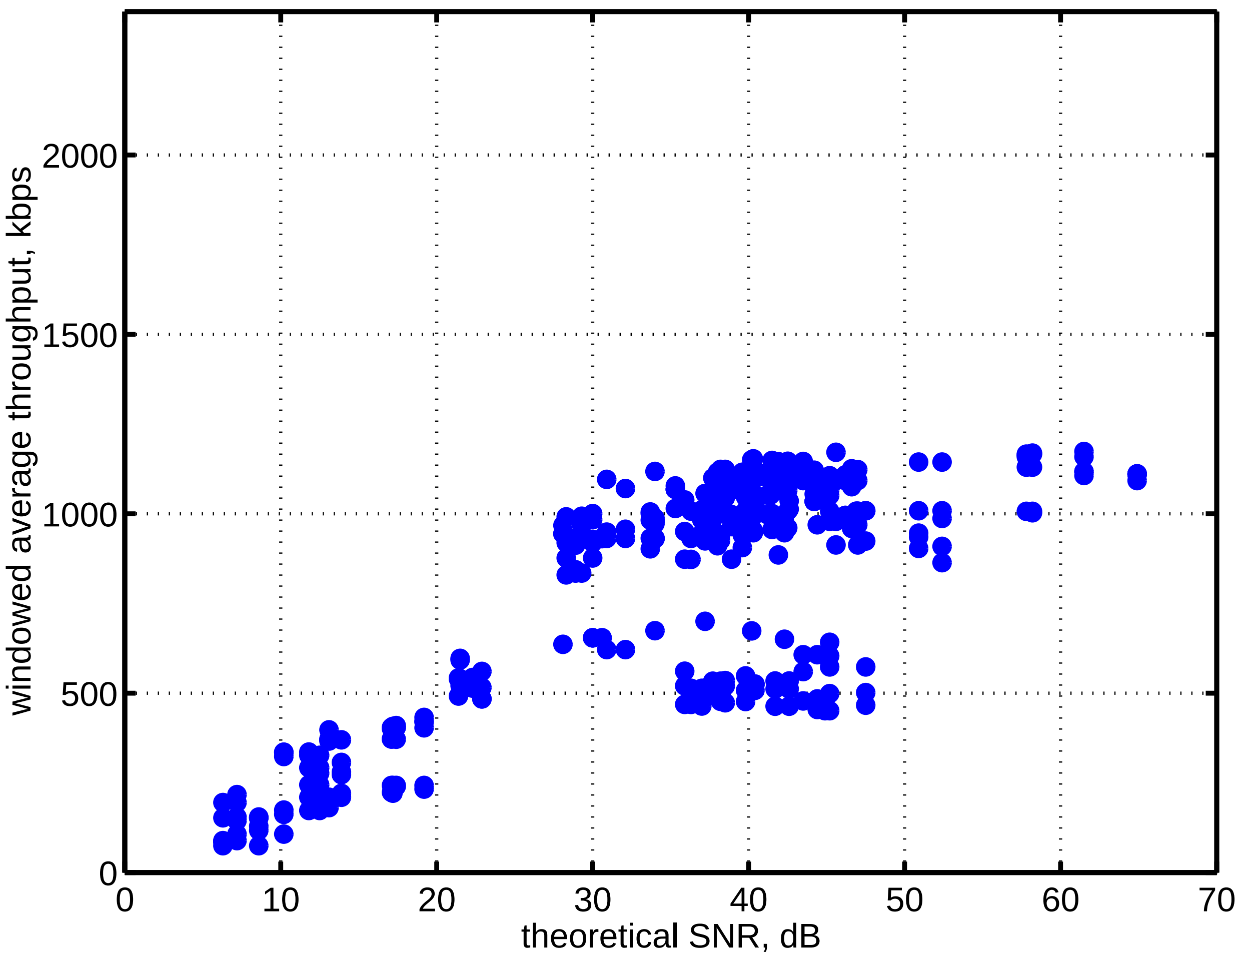

Figure 5.24: All measurements averaged over similar signal strength (i.e. RSSI) locations for Scenario D, 3Com System

5.6 Analysis of Throughput and SNR Measurements

As can be seen in the above figures, a distinct relationship between the signal strength and throughput of a wireless LAN is present. This relationship can be seen clearest in the window averages shown in the third graphs of each of the presented sets. An overall trend of an exponential curve is seen throughout all of the measurements. Additionally, the interference measurements show a distinct switching between discrete data bandwidths for the interference Scenarios C and D, as in the two levels of different throughput present in Figure 5.18. This is expected because the wireless LAN standard defines different data bandwidths that allow the throughput to decrease gracefully as the RF environment degrades. These observations were applied to the empirical throughput prediction models which are discussed in the next chapter.

Notice that the Throughput measured at each location never exceeds a certain level. This maximum throughput is well below the 11 Mbps data bandwidth supposedly available with IEEE 802.11b products. It is normal for the throughput to be below the data bandwidth of a network connection because of the overhead in each data packet sent. Further, the manner in which the throughput is measured will affect the measured throughput. Since LANFielder measures the throughput for the non-optimistic case in which only a single packet is transmitted at any one time, the throughput measured is particularly low compared to the 11 Mbps data bandwidth available. However, the measured data can easily be compared to measurements made using different techniques by simply normalizing the throughput by the maximum measured throughput.

5.7 Interference Signal Strength Measurements and Predictions

Measurements of the signal strength from the interfering access point location and the interfering client locations to the desired client locations were measured. It would have been preferred to have also measured the interfering signal strengths at the desired access point locations, as well. This would have been extremely difficult because of the impossibility of accurately mimicking the desired access point locations with a measurement laptop. Additionally, the 3Com access point antenna was incompatible with any wireless

LAN card which could be used to measure signal strength. For this reason, all prior SNR measurements were used to derive a path loss model (of the form presented in [25] and given in equation 5.1) which could then be used to predict what the signal strength at the desired access point location from the interfering access point and client locations. This process is discussed further in Section 5.7.1.

5.7.1 Predicting the Interference at Signal Strength Locations

SitePlanner Model of Measurement Cases

In order to predict the needed interference, the wireless prediction tool SitePlanner was used. In SitePlanner, the environments from all three cases were modeled in a three dimensions. Partitions such as walls and cubicles were divided into four different types: Primary Partitions, Secondary Partitions, Doors and Metal [26]. Each partition has a different Attenuation Loss associated with it. Primary partitions were used to model the cinder block hallways and drywall partitions. Secondary partitions were used to model the cubicle partitions. The door partitions are self-explanatory and the metal partitions were used to model the elevators. SitePlanner has facilities for using eight different partition types but these were not all used. All partitions extended from floor to ceiling except for the cubicle partitions which were modeled to extend to 1.7 meters as was true of the actual partitions. Figure 5.25 shows the environment as it was modeled as well as some of the measurements that were entered into SitePlanner. The different partition types are represented using different colors. The primary, secondary, door, and metal partitions are blue, gray, purple, and orange, respectively.

Inputting SNR Measurements into SitePlanner

Using InFielder, SNR measurements made using the custom RSSI measurement software, were entered into the drawing at the precise three dimensional coordinates that they were measured at. In order to do this, the RSSI measurements were first linearly averaged over the four cardinal directions for each measurement point and then placed into a SitePlanner logfile format and imported. The measurements were imported in such a way that the measurements were associated with the correct access point location in an automatic

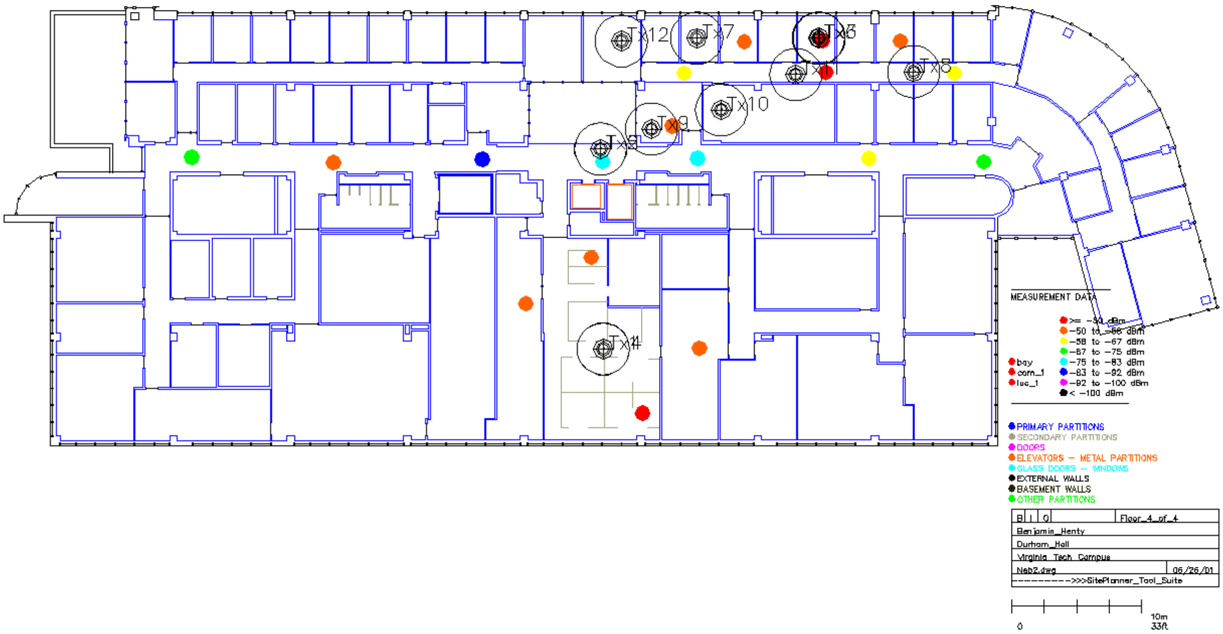

Figure 5.25: The entire measurement environment modeled in SitePlanner.

manner.

Path Loss Prediction in SitePlanner

SitePlanner has several different prediction models available [26] The model used for these predictions was the Wall/Floor Attenuation Factor, Single Path Loss Exponential model. Since all measurements were taken on a single floor, this model uses Equation 5.1 for all predictions.

$$PL(d) = PL(d_0) + 10nlog\left(\frac{d}{d_0}\right) + \sum_i [(P_i)(AF_i)] \tag{5.1}$$

This equation predicts the path loss, PL at distance d. d_0 is a reference distance, $PL(d_0)$ is the path loss at the reference distance, n is the path loss exponent, P_i is the number of partitions of type i and AF_i is the attenuation factor for partitions of type i. SitePlanner uses the Friis Equation [24] to calculate the path loss at the reference distance d_0 [26]. This formula for the value of $PL(d_0)$ in dB is given in equation 5.2.

$$PL(d_0) = 10 * log\left[\left(\frac{4\pi d_0}{\lambda}\right)^2\right] \tag{5.2}$$

Optimizing Path Loss Prediction Model Parameters

After all of the measurements were imported into SitePlanner and the access points were correctly modeled, the Optimatic module of SitePlanner was used to calculate optimal values for n and AF_i using the RSSI measurements previously entered into SitePlanner. SitePlanner uses a basic Mean Squared Error fit algorithm to find an ideal parameters based on known measured data. Using all of the 3Com and WaveLAN measured data from the desired access point to the desired client locations, the following parameters resulted in a mean error of 0.05 dB and a standard deviation of 3.27 dB. Further, 100measurements were within 10 dB of the predicted value. The calculated parameters were 3.02 for the path loss exponent, 1.70 dB for the primary partition attenuation factor, 2.41 dB for the secondary partition attenuation factor and 2.50 dB for the door partition attenuation factor. There was not enough measurements to compute an ideal value for the elevator partition attenuation factor, so the default value of 5.00 dB was used. These parameters and all of the settings used in SitePlanner are summarized in Table 5.1. It should also be noted that these results are in excellent agreement with the measurements made in [11].

Table 5.1: Summary of All Configured Parameters Used in SitePlanner

Parameter Name	Parameter Value
3Com Access Point Transmit Power	18 dBm
WaveLAN Access Point Transmit Power	15 dBm
BayNetworks Access Point Transmit Power	15 dBm
BayNetworks Client Transmit Power	15 dBm
Access Point Heights	2.5 meters
Receiver Transmit Power	15 dBm
Receiver Antenna Gain	2 dBi
Receiver Height	0.91 meters
Transmit/Receive Frequency	2.412 GHz
Reference Distance (d_0)	1 meter
Path Loss at the Reference Distance, $PL(d_0)$	40.0 dB
Path Loss Exponent, n	3.02
Primary Partition Loss Factor, AF_1	1.70 dB
Secondary Partition Loss Factor, AF_2	2.41 dB
Door Partition Loss Factor, AF_3	2.50 dB
Metal Partition Loss Factor, AF_4	5.00 dB

Modeling Access Points in SitePlanner

The desired access points are modeled in SitePlanner as base stations and can be seen in Figure 5.25 as concentric black circles and are labeled with "TX" names. The access points were correctly modeled to have 15 or 18 dBm transmit powers for the WaveLAN or 3Com access points, respectively. The access points were modeled at the precise three dimensional locations that they were positioned in the actual environment. This includes the correct antenna height. Each access point had a 2 dBi omnidirectional antenna attached. This is the correct maximum gain, though not terribly accurate in terms of the actual radiation pattern of the access points. It was, however, felt that since measurements were averaged over the four cardinal directions, the omnidirectional pattern would result in an accurate representation of the measurements. In addition to the desired access point, the interfering access point and interfering client were also modeled as base stations and are shown in Figure 5.25. These base stations were also modeled in the correct three dimensional location and as having a 2 dBi omnidirectional antenna and a transmit power of 15 dBm, which matches the actual specifications of the BayNetworks products.

Additional Parameters of the SitePlanner Model

It was necessary to set several additional parameters in addition to the locations and transmit powers of the access points. All base stations were configured to be IEEE 802.11b access points with the DSSS channels available and channel 1 the active channel. That is, all predictions were made using a center frequency of 2412 MHz which corresponds to the center frequency of channel 1 in the DSSS portion of the IEEE 802.11 standard [30]. Several receiver parameters were also configured in SitePlanner. The mobile receiver was configured to have a height of 0.91 meters and a receiver antenna gain of 2 dBi which was included in all path loss calculations.

Predicting Interference Signal Strength

Once an effective prediction model was calculated using Equation 5.25 in SitePlanner, the model was then used to predict the signal strength at the desired access point location from the interfering access point location and the interfering client locations. To do this, base stations were modeled in SitePlanner as was previously described and as is outlined in Table 5.1. These base stations were then used to predict the signal strength at the base stations which represent the desired access points. Combining these RSSI predictions with the RSSI measurements of the interfering access point and interfering client at the desired client locations resulted in four interference measurements for every desired client location. These four measurements are used to calculate a single SIR value using the technique outlined in Section 5.7.2.

5.7.2 Computing Total Interference Signal Strength

The measured interference was found through a combination of measurement and through predictions based on measurements. As a result, each throughput measurement in Cases C and D is associated with four interference signal strength measurements. These four measurement locations are illustrated in Figure 5.26. The four path loss measurements, PL, are shown using a solid arrows in the figure. For reference, the SNR for a client is based on the path loss from the desired client and the desired access point location, shown in Figure 5.26 as a dashed line.

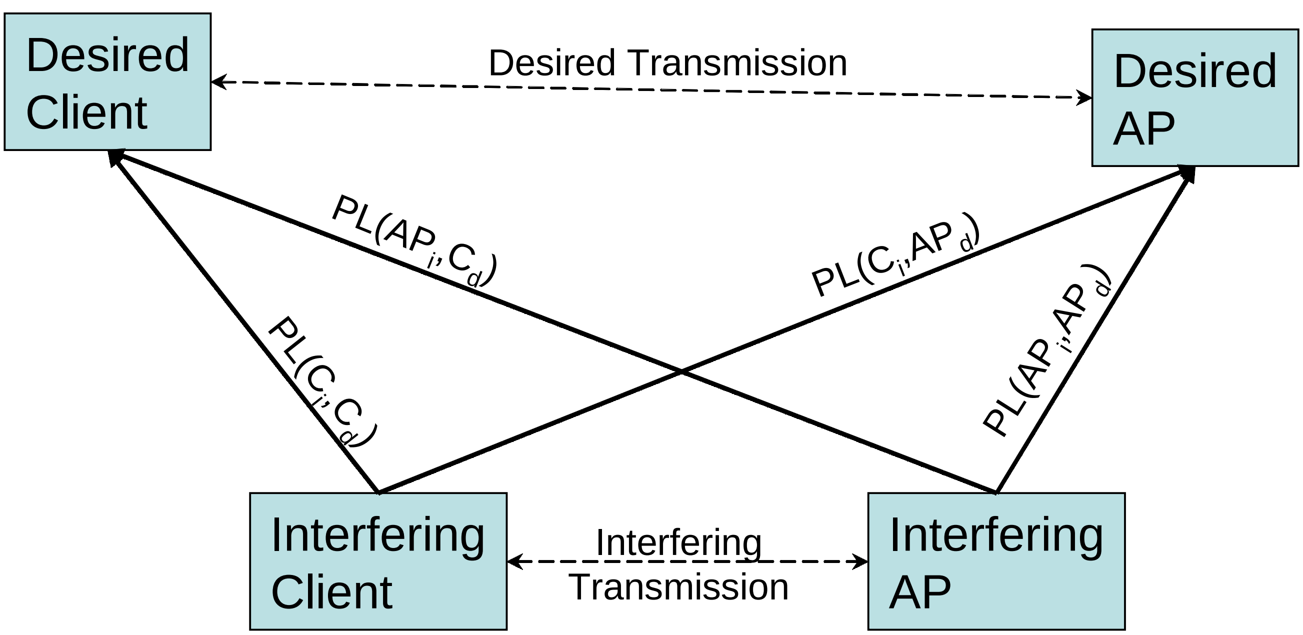

Figure 5.26: A diagram of the interference paths between the different wireless LAN users in Scenarios C and D.

The four path loss measurements, $PL(AP_i, AP_d)$, $PL(AP_i, C_d)$, $PL(C_i, AP_d)$ and $PL(C_i, C_d)$, correspond to the four respective $RSSI$ values used to compute a combined SIR value using Equation 5.3.

$$SIR_{net} = \frac{RSSI(AP_d, C_d)}{mean((RSSI(AP_i, AP_d) + RSSI(C_i, AP_d)), (RSSI(AP_i, C_d) + RSSI(C_i, C_d)))} \tag{5.3}$$

In this equation, $RSSI(X_a, Y_b)$ is the received signal strength at location Y_b from location X_a. X_a and Y_b are symbols in which X and Y are either AP or C or an access point or client location respectively. The subscript a or b is either d or i to represent the desired or interfering location respectively. Thus, $RSSI(AP_i, C_d)$ is the signal strength at the desired client from the interfering access point. Note that all RSSI values were summed and averaged by first converting to a linear value, performing the desired operation and then converting back into dB.

Equation 5.3 implements a linear average of the total interference power at the desired access point, AP_d, location with the total interference power at the desired client location, C_d. An average performed in this manner does not excessively decrease a strong interference signal when combined with a much weaker interference signal. An average is used because it is assumed that the access point will be receiving half the time and the client will be receiving half the time, rather than both simultaneously. This is the motivation behind Equation 5.3.

5.7.3 Advantages of SIR Calculation

This formula has the advantage that an unusually strong interfering signal is not excessively reduced by a unusually weak interfering signal. Further, Equation 5.3 is based on the average of the total interference power at the desired client and the total interference power at the desired access point. The two values are averaged because the client and access point never transmit simultaneously. It is theorized then that the throughput experienced will be related to the average interference power at the desired client and the desired server.

5.7.4 Resulting Plots of Measured and Predicted SIR Data

After applying Equation 5.3 to the measured and predicted interference power levels, the following plots of the data were generated. These plots use the measured throughput from Scenarios C and D and plot the throughput versus the computed net interference power and versus the SIR values. These plots are presented below.

Scenario C, WaveLAN Throughput versus Net Interference Power

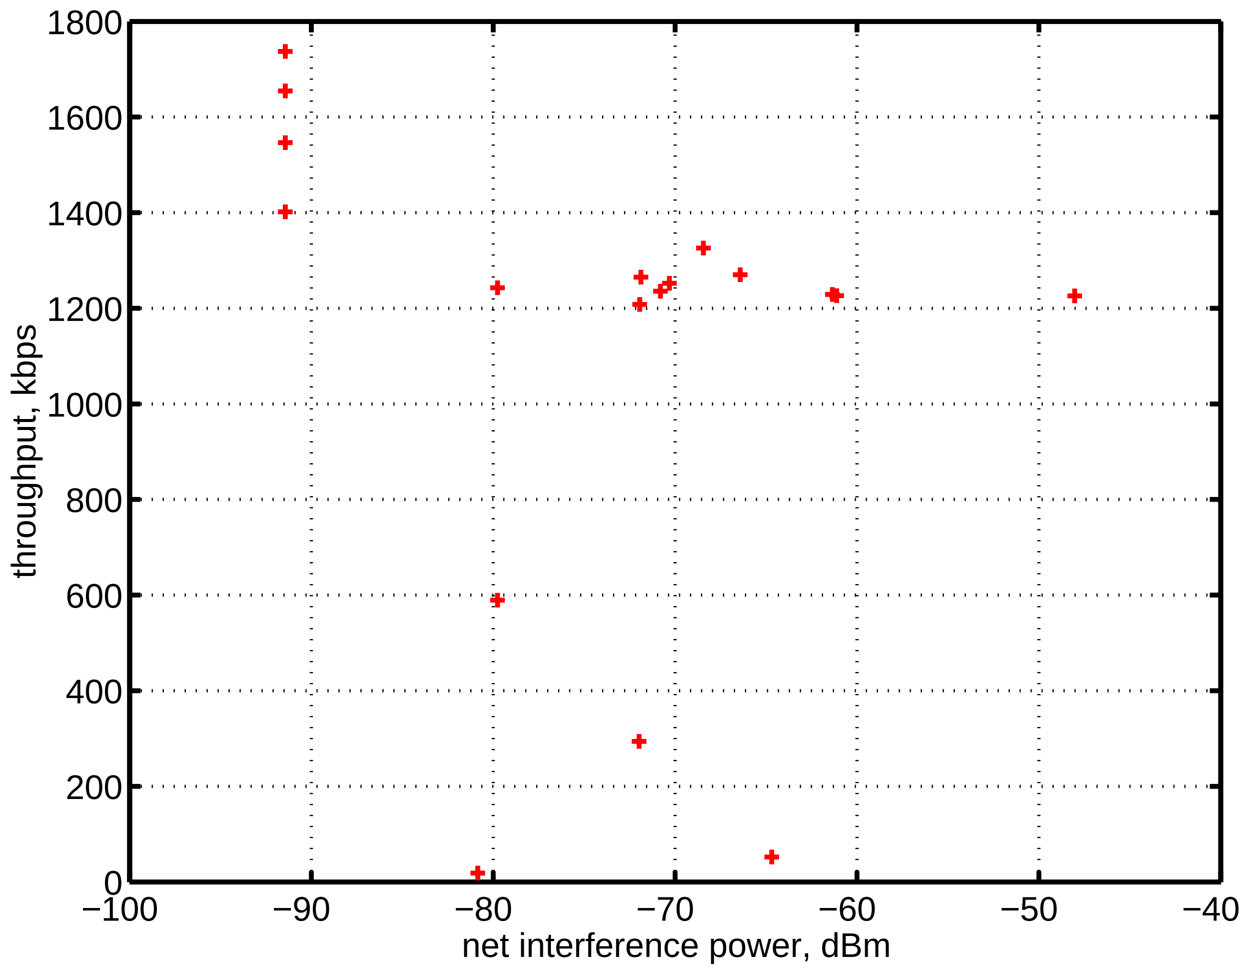

Figure 5.27: Plots of the WaveLAN measured throughput versus the calculated net interference power averaged over all cardinal directions for each physical location

Scenario C, WaveLAN Throughput versus SIR

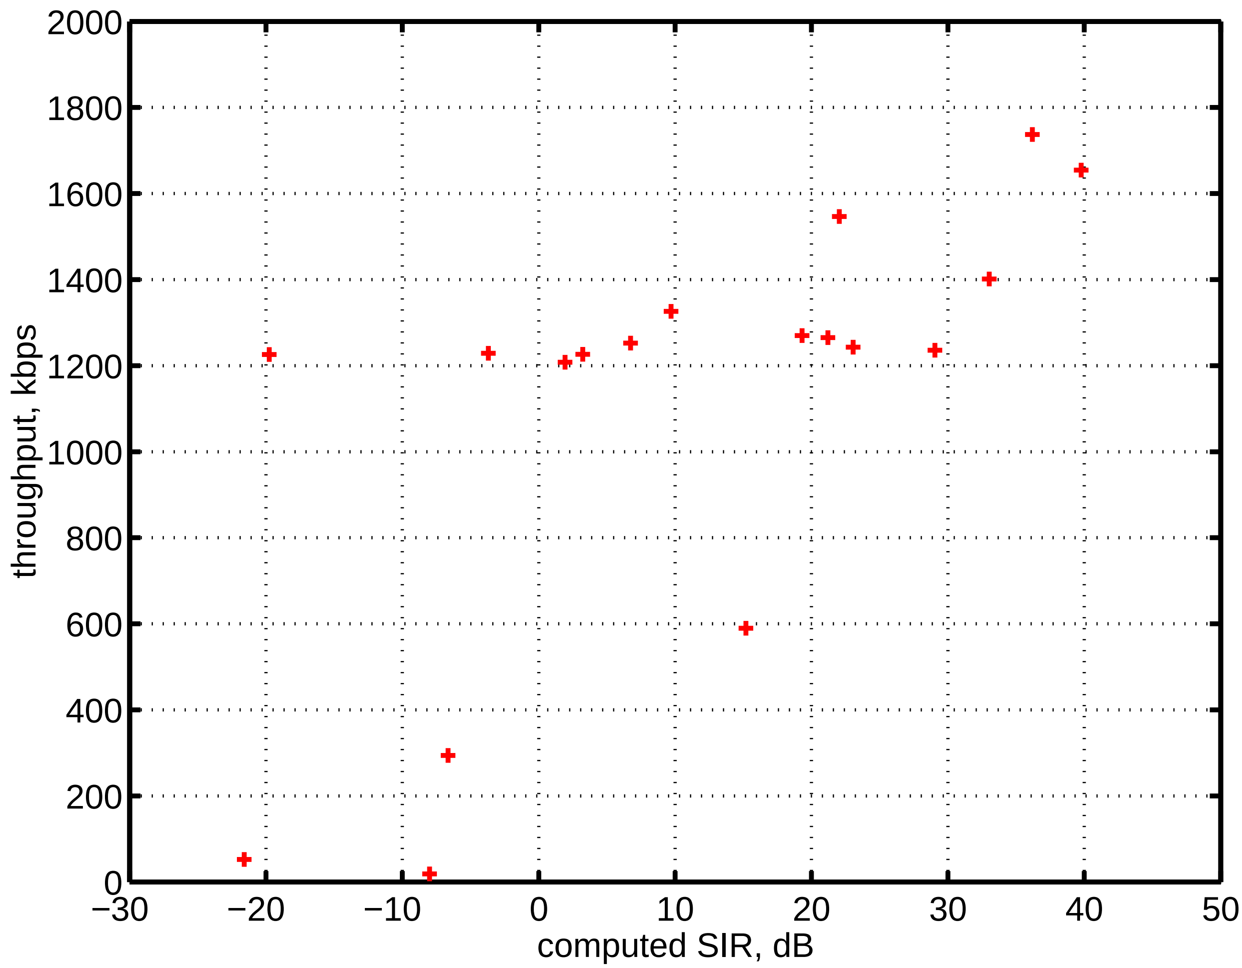

Figure 5.28: Plots of the WaveLAN measured throughput versus the calculated SIR averaged over all cardinal directions for each physical location

Scenario C, 3Com Throughput versus Net Interference Power

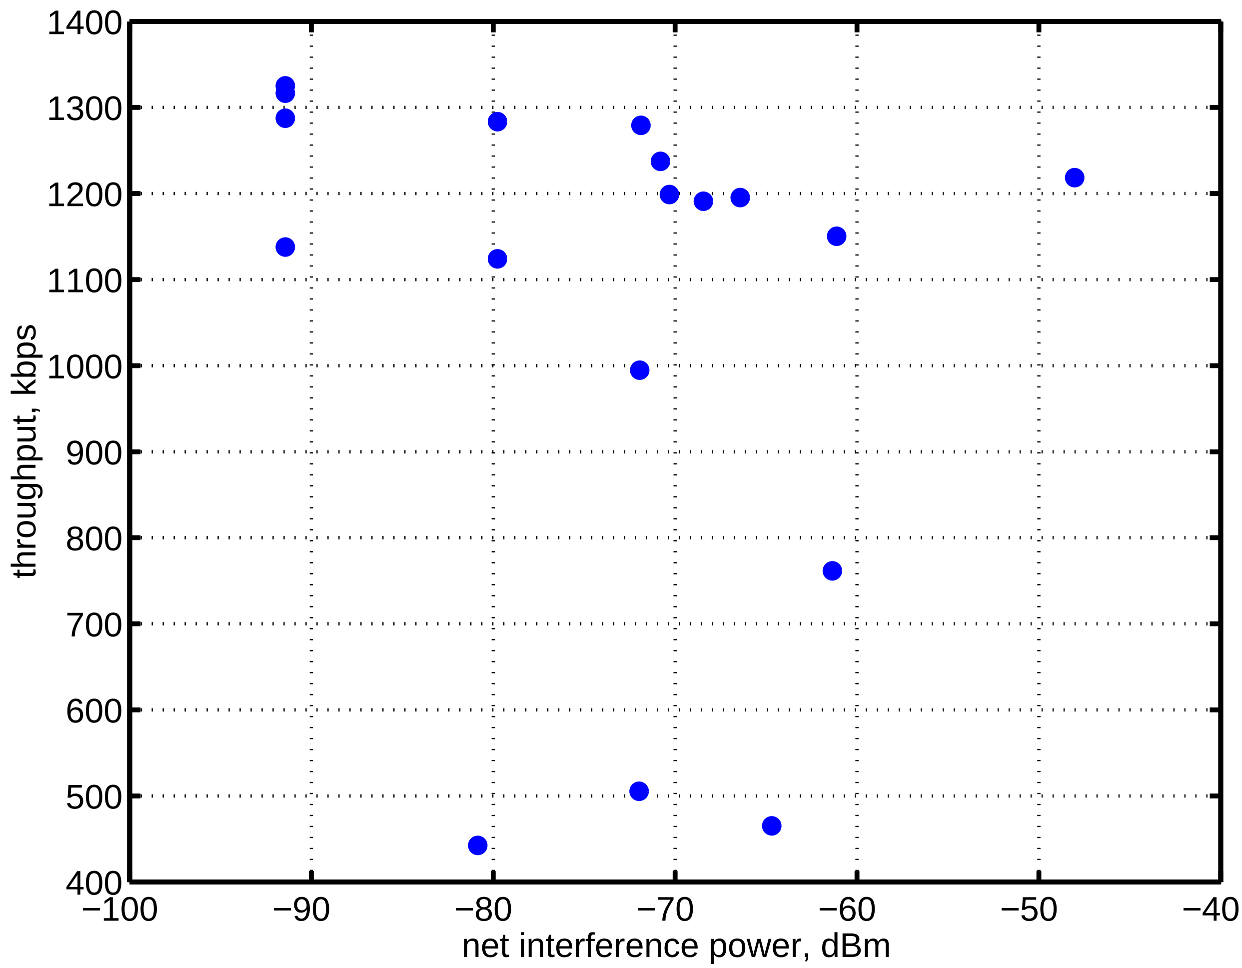

Figure 5.29: Plots of the 3Com measured throughput versus the calculated net interference power averaged over all cardinal directions for each physical location

Scenario C, 3Com Throughput versus SIR

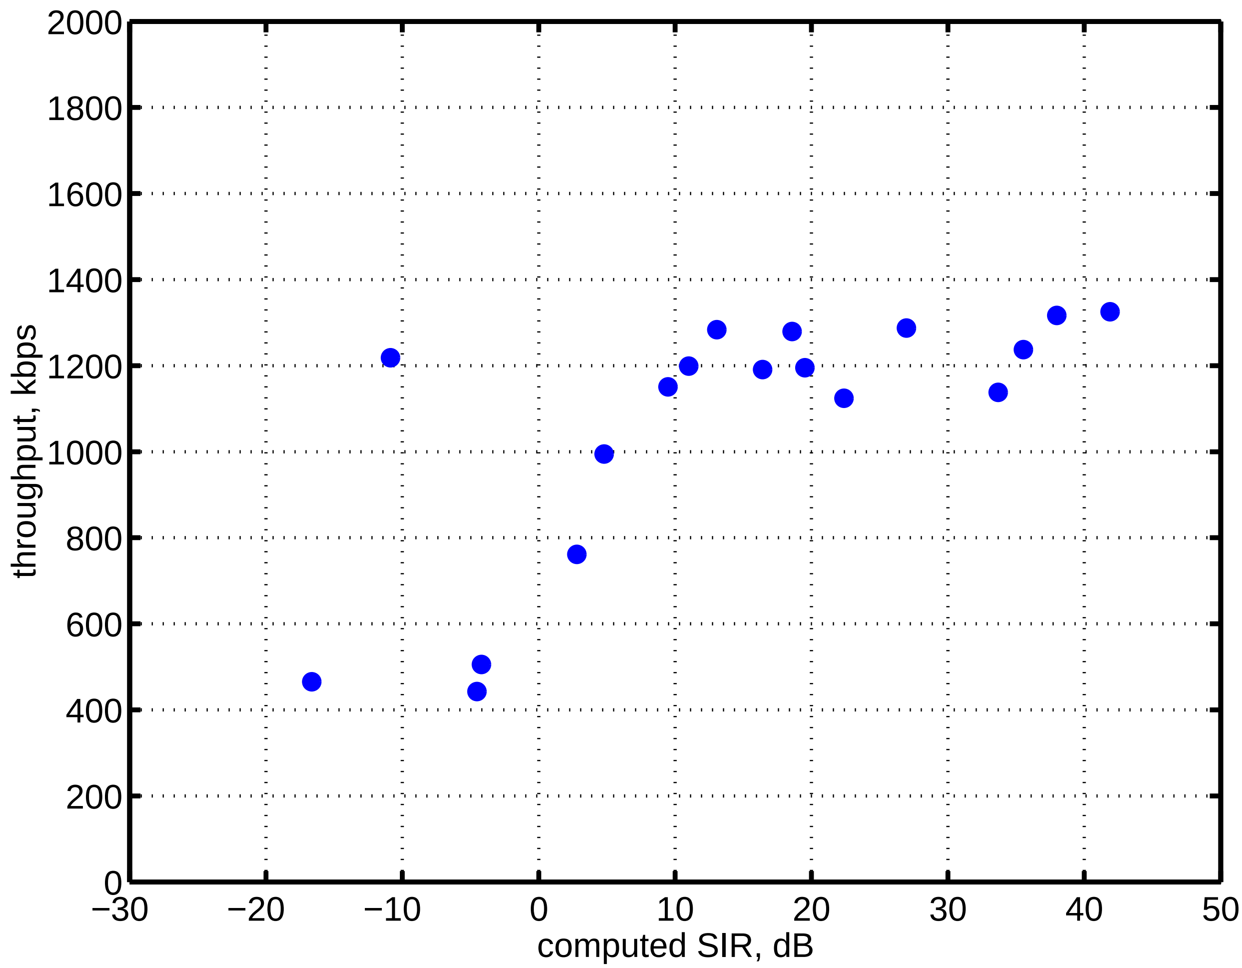

Figure 5.30: Plots of the 3Com measured throughput versus the calculated SIR averaged over all cardinal directions for each physical location

Scenario D, WaveLAN Throughput versus Net Interference Power

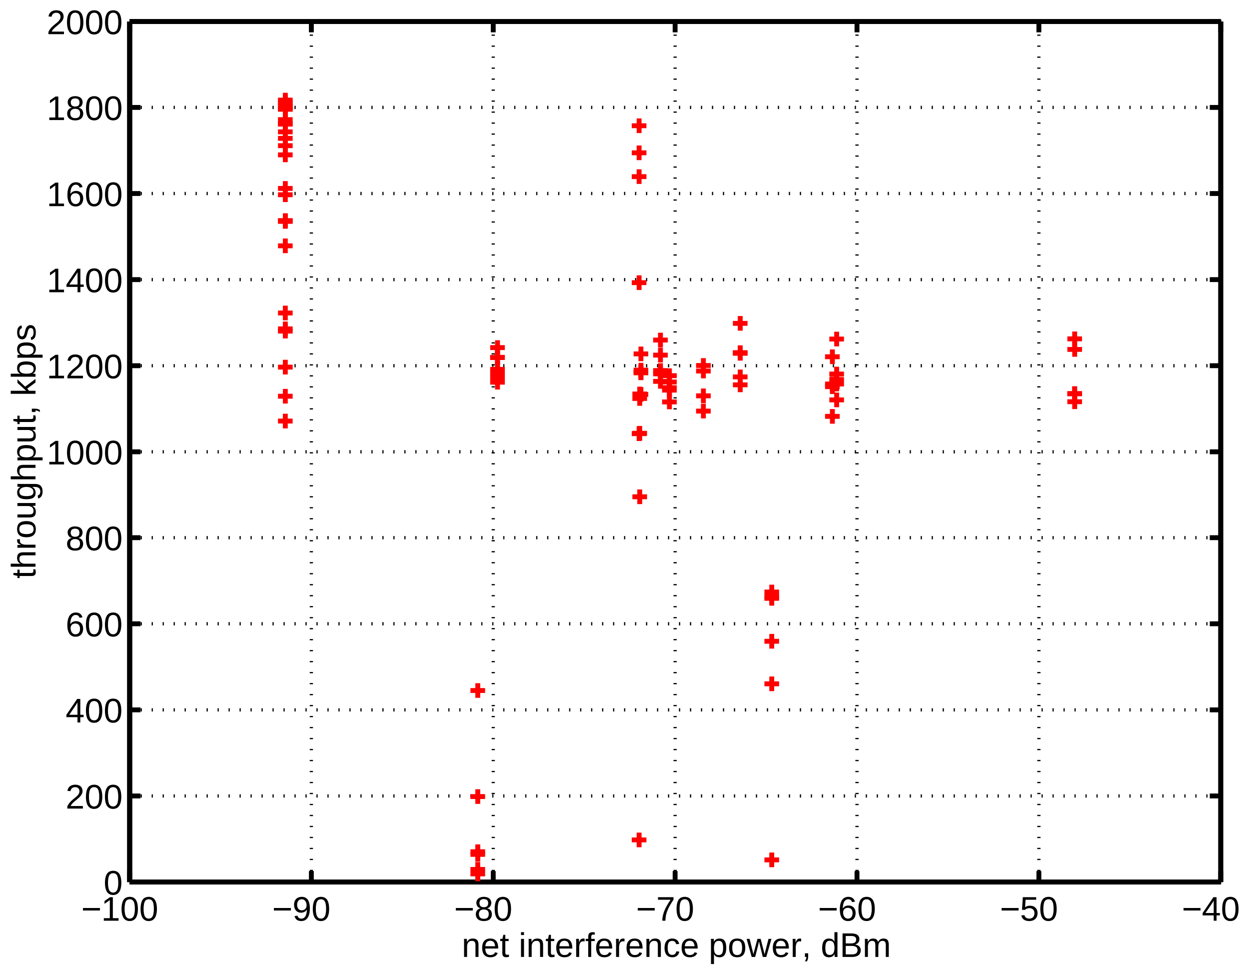

Figure 5.31: Plots of the WaveLAN measured throughput versus the calculated net interference power averaged over all cardinal directions for each physical location. All data points for both clients are shown in the graph.

Scenario D, WaveLAN Throughput versus SIR

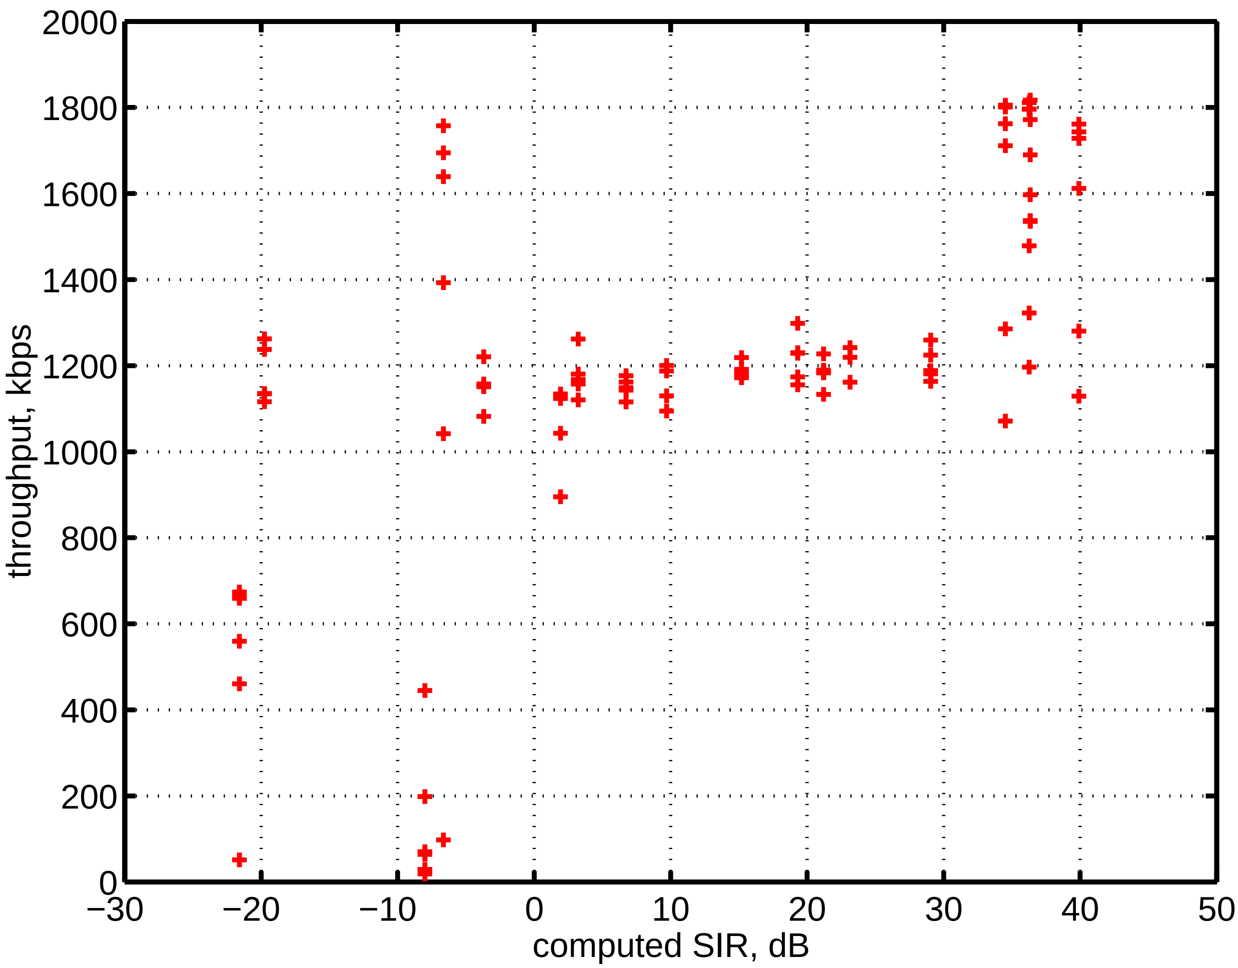

Figure 5.32: Plots of the WaveLAN measured throughput versus the calculated SIR averaged over all cardinal directions for each physical location. All data points for both clients are shown in the graph.

Scenario D, 3Com Throughput versus Net Interference Power

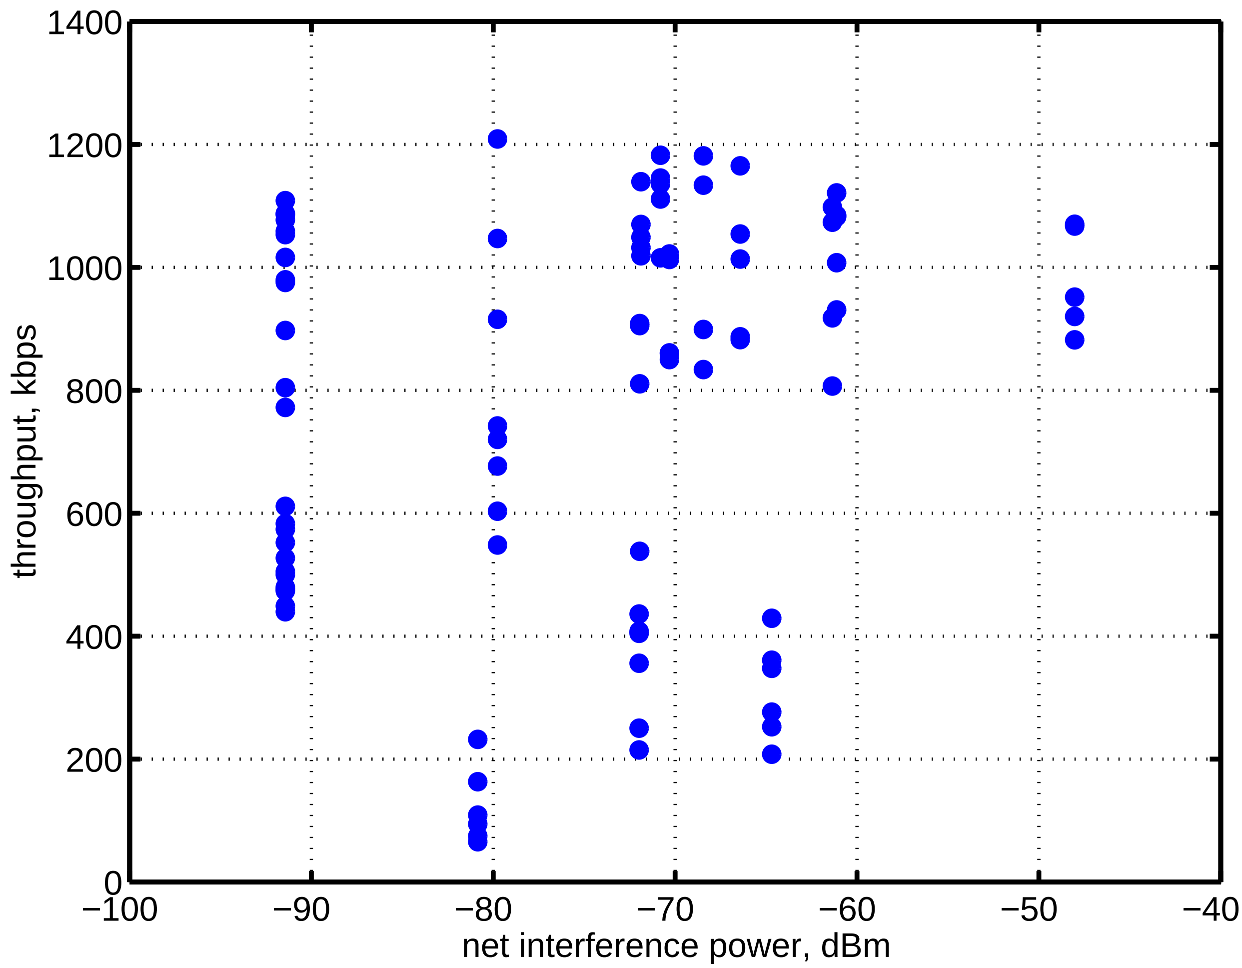

Figure 5.33: Plots of the 3Com measured throughput versus the calculated net interference power averaged over all cardinal directions for each physical location. All data points for both clients are shown in the graph.

Scenario D, 3Com Throughput versus SIR

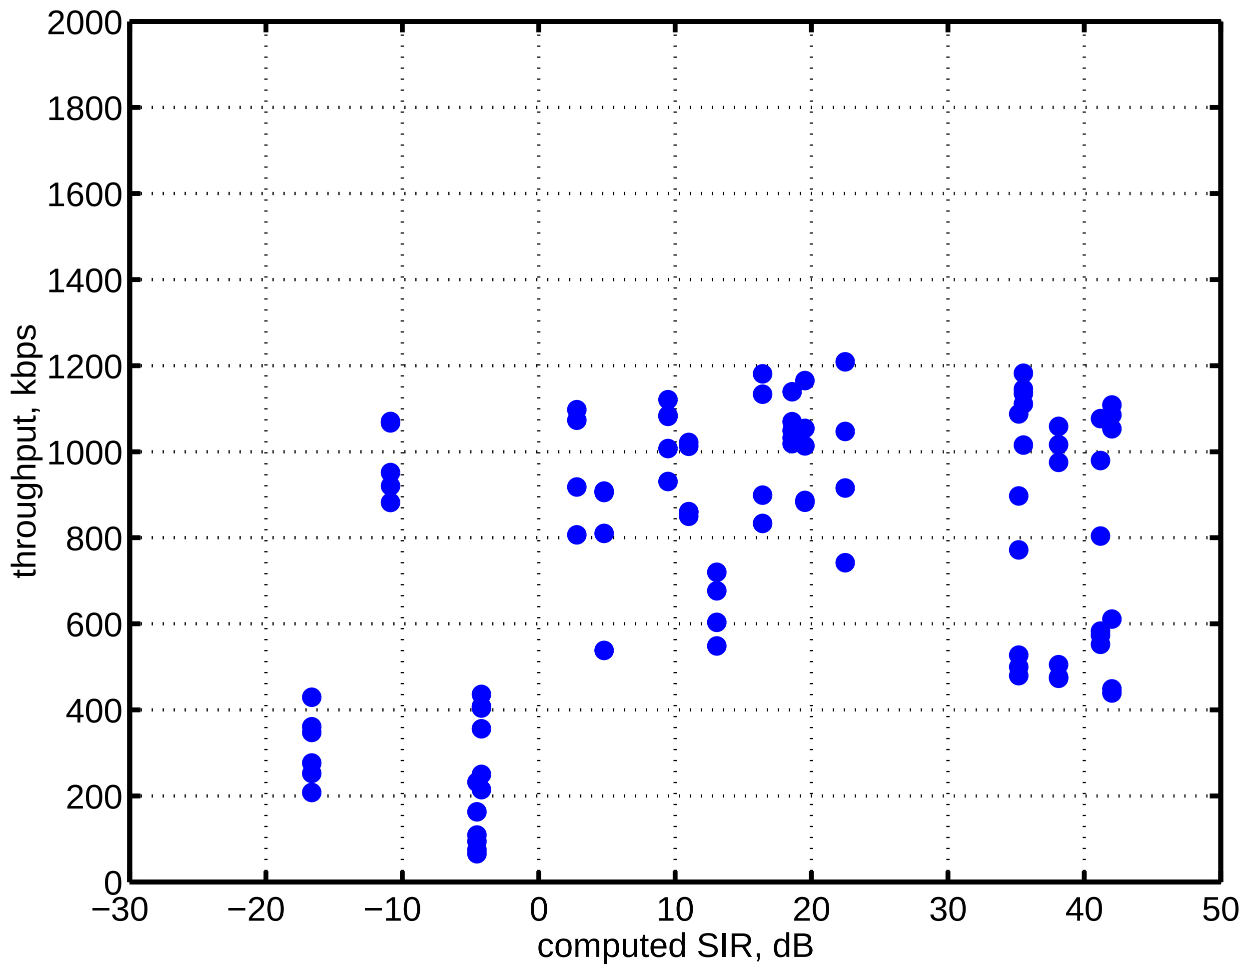

Figure 5.34: Plots of the 3Com measured throughput versus the calculated SIR averaged over all cardinal directions for each physical location. All data points for both clients are shown in the graph.

5.7.5 Analysis of Interference Measurements and Calculations

As can be seen in Figures 5.27 through 5.34, there is very little correlation between the net interference signal strength and the measured throughput measured, though there is some correlation between the SIR and the throughput. Further, by comparing scenarios C and D, it is clear that there is consistency from between the two scenarios. Thus, the throughput of a client can not be predicted directly from the net interference power, it does appear to be a useful parameter.

5.8 Measurement Results Conclusions

It has been seen from the measurement results presented in this chapter that there is definitely promise for creating empirical throughput prediction models. The presented plots of the measurements show a definite correlation between the signal to noise ratio, and to some small degree, to the net interference power. These relationships are clearest in the window averaged plots that have been presented. As a result, the measurements are used in Chapter 6 to create empirical throughput prediction models for wireless LANs.

Chapter 6

Empirical Throughput Prediction Models

6.1 Development of Wireless LAN Performance Models

After analyzing the data presented in the previous chapter, it appears that the throughput of a wireless LAN can be modeled and predicted. The goal of this chapter is to present models that have been developed to predict the throughput performance, T, of wireless LAN networks. The models that have been developed are presented in this chapter, along with the reasoning behind the models and the accuracy of the models, when compared to the measurements. All of the models presented were developed empirically (that is, models were tuned using the measured data), using the data that was averaged per each measurement location. That is, each data point used in the models was the average of four measurements (north, east, south, west) at a single physical location.

6.2 Optimizing Prediction Model Parameters

To fit the prediction models to the data, a technique similar to the MMSE fitting is used [20]. The difficulty with a MMSE error fit is that it requires a linear or polynomial equation to function. Since this thesis attempts to fit to nonlinear models, an exact MMSE fitting procedure is not valid. Therefore, the MATLAB command *nlinfit*, which is part of the Statistics Toolbox, was used. This command implements the Gauss-Newton method to find a local minimum of a given equation using provided data points [16]. The technique requires an initial guess to function, and does not guarantee the “best” fit, but only a local minima. However, a good initial guess typically provides a result which is an extremely good fit to the data.

6.3 Prediction Models for Scenario A – One User

Two primary models were considered for this scenario. The first of these models is a piece-wise linear fit to the measured data. The second model is an exponential curve fit to the measured data. Each of these models have three parameters that need to be specified:

1. T_{max}, the throughput experienced in ideal circumstances (i.e. under maximum SNR and without interference)
2. A or α, a “slope” or rate of drop-off in throughput as the SNR decreases
3. T_0 or SNR_0, intercept points for the model

Note that unless otherwise specified, all throughputs are in units of kbps and all SNR values are in dB. All calculations were carried out in MATLAB after having imported all of the measurement data into MATLAB from the LANFielder logfile format. Details of these models are now presented in the sections below.

6.3.1 Scenario A, Linear Model

Formula for the Scenario A Linear Model

A linear model was first considered because in looking at the single user measured data, without co-channel users, it appeared that the wireless LAN tended to operate with a very steady throughput until the signal to noise ratio (SNR) dropped below a certain critical level. Below this critical SNR, represented as SNR_c, the throughput appears to drop-off linearly. This drop-off can be clearly seen in Figures 6.1, 6.2, 6.3, and 6.4. The formula for a reasonable model is given in equation 6.1.

$$T = \begin{cases} T_{max} & \text{if} \quad SNR > SNR_c \\ A * SNR - T_0 & \text{if} \quad SNR <= SNR_c \end{cases} \tag{6.1}$$

In the linear prediction equation, T_{max} is a critical throughput that is vendor and application specific. That is, different wireless LAN cards (e.g. WaveLAN, 3Com, etc.) running different network applications (e.g. telnet, VoIP, FTP, etc.) will have a different value for T_{max}. T_{max} is the throughput that the wireless LAN will provide in perfectly ideal circumstances. The parameters A and T_0 are also vendor and application specific, and are used to predict the throughput of the wireless LAN when the performance is not ideal. That is, when the SNR drops below SNR_c, the throughput drops of with a slope of A and reaches a throughput of T_0 when the SNR drops to 0 dB. Note that the value of SNR_c is the intersection of the two lines of the model. Thus, SNR_c can be derived from T_{max}, A and T_0 using Equation 6.2.

$$SNR_c = \frac{T_{max} + T_0}{A} \tag{6.2}$$

Calculating the Linear Model Parameters

The procedure explained in Section 6.2 was used to fit the parameters of Equation 6.3. The initial guesses selected for the WaveLAN data were T_{max} equal to 1800, A equal to 100 and T_0 equal to 0. The initial guesses for 3Com were T_{max} equal to 2100, A equal to 100 and T_0 equal to 0. After obtaining the optimum values for T_{max}, A and T_0, the value of SNR_c was calculated using Equation 6.2. The optimized model parameters are given in Table 6.1. Further, the resulting model and the locationally averaged data are shown in Figures 6.1 and 6.2.

The equations used to generate the model are shown on the graphs along with the mean and standard deviation in the error. These error statistics are also presented for this and all prediction models in this Chapter in Table 6.2. As can be seen from the statistics, both the WaveLAN and the 3Com data fit to the linear model quite well. Both had mean errors of 0 and standard deviations of less than about 150 kbps. This is quite good considering LANFielder and SiteSpy have a measurement precision of about 10 kbps.

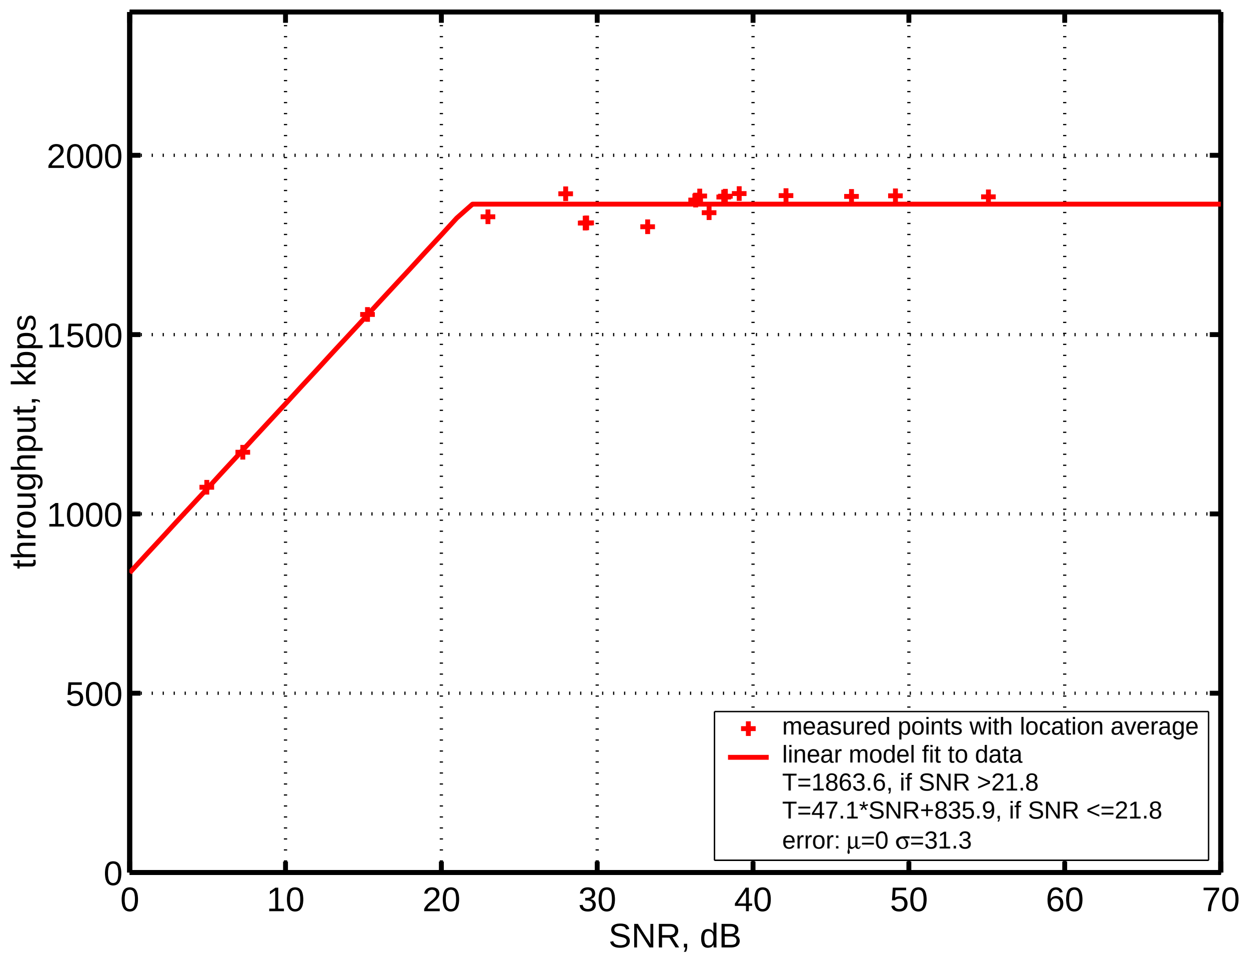

Figure 6.1: WaveLAN Scenario A, Location Averaged Data and Linear Throughput Prediction Model

Constraining the Linear Model

As can be seen from Figures 6.1 and 6.2, the linear model of throughput does not predict zero throughput at a 0 dB SNR. For this reason, the Linear models were refit to the data, except with T constrained to be equal to zero at zero SNR which is equivalent to forcing T_0 to be zero. In this manner the prediction models are simpler, more intuitive and more realistic, since throughput drops to zero when the SNR reaches zero. The model was again fit using the MATLAB *nlinfit* command as described in Section 6.2 and was performed previously in Section 6.3.1 except that T_0 was not considered an adjustable model parameter.

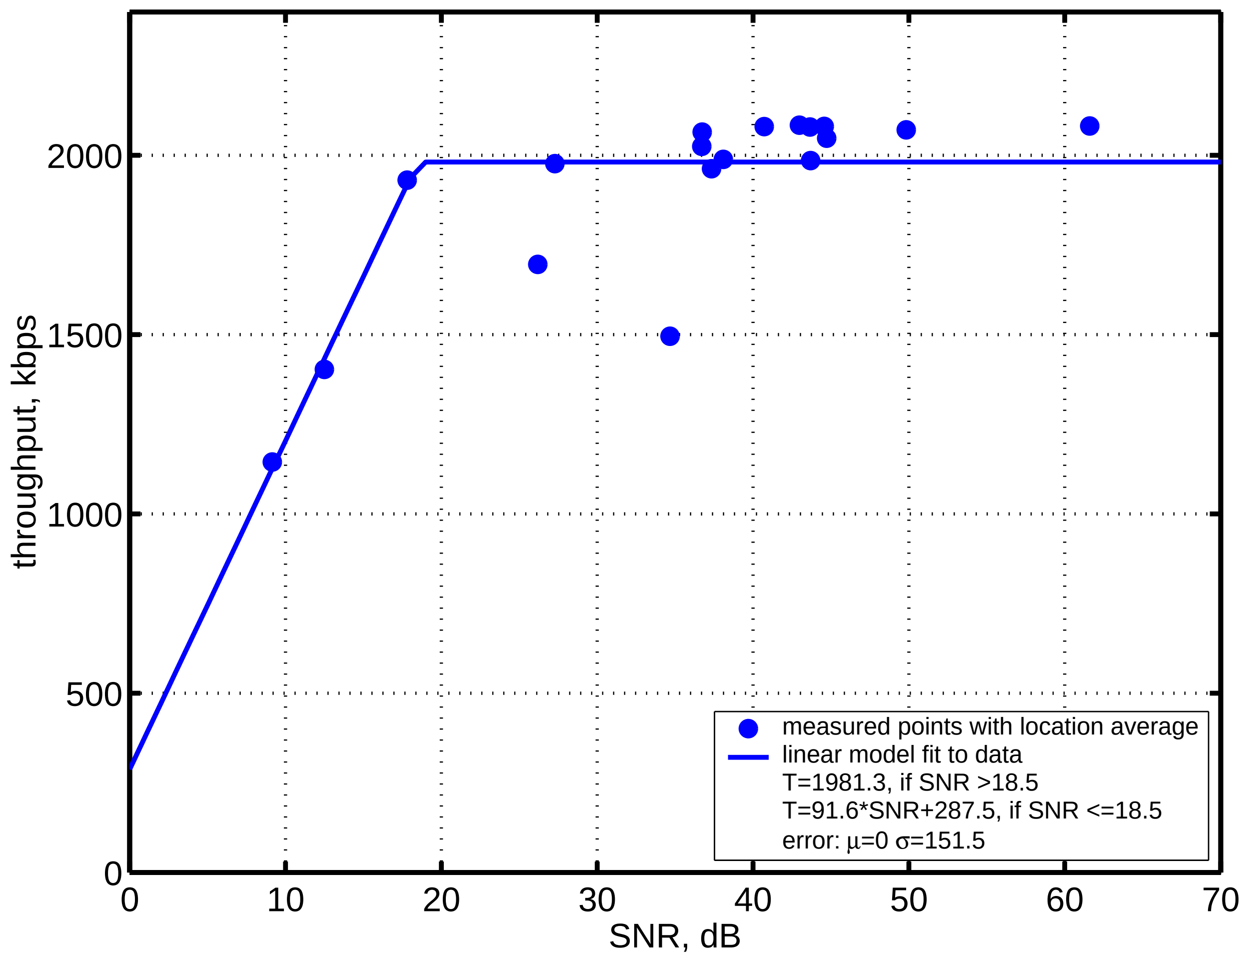

Figure 6.2: 3Com Scenario A, Location Averaged Data and Linear Throughput Prediction Model

The resulting error statistics are given in Table 6.2 and the calculated model parameters are given in Table 6.1. resulting constrained models and the measured data are plotted in Figures 6.3 and 6.4.

Table 6.2 and Figures 6.3 and 6.4 show that the mean error, μ, increased from 0 to 26.29 kbps for the WaveLAN data and from 0 to 1.12 kbps for the 3Com data. The error standard deviation, σ, increased from 31.3 kbps to 153.3 kbps for the WaveLAN data and from 151.5 kbps to 153.4 kbps for the 3Com data. As can be seen from these results, constraining the model introduces a mean error and the error standard deviation increases. Thus, the more intuitive, simpler model that results from constraining T to be zero at 0 SNR comes with a large penalty in accuracy.

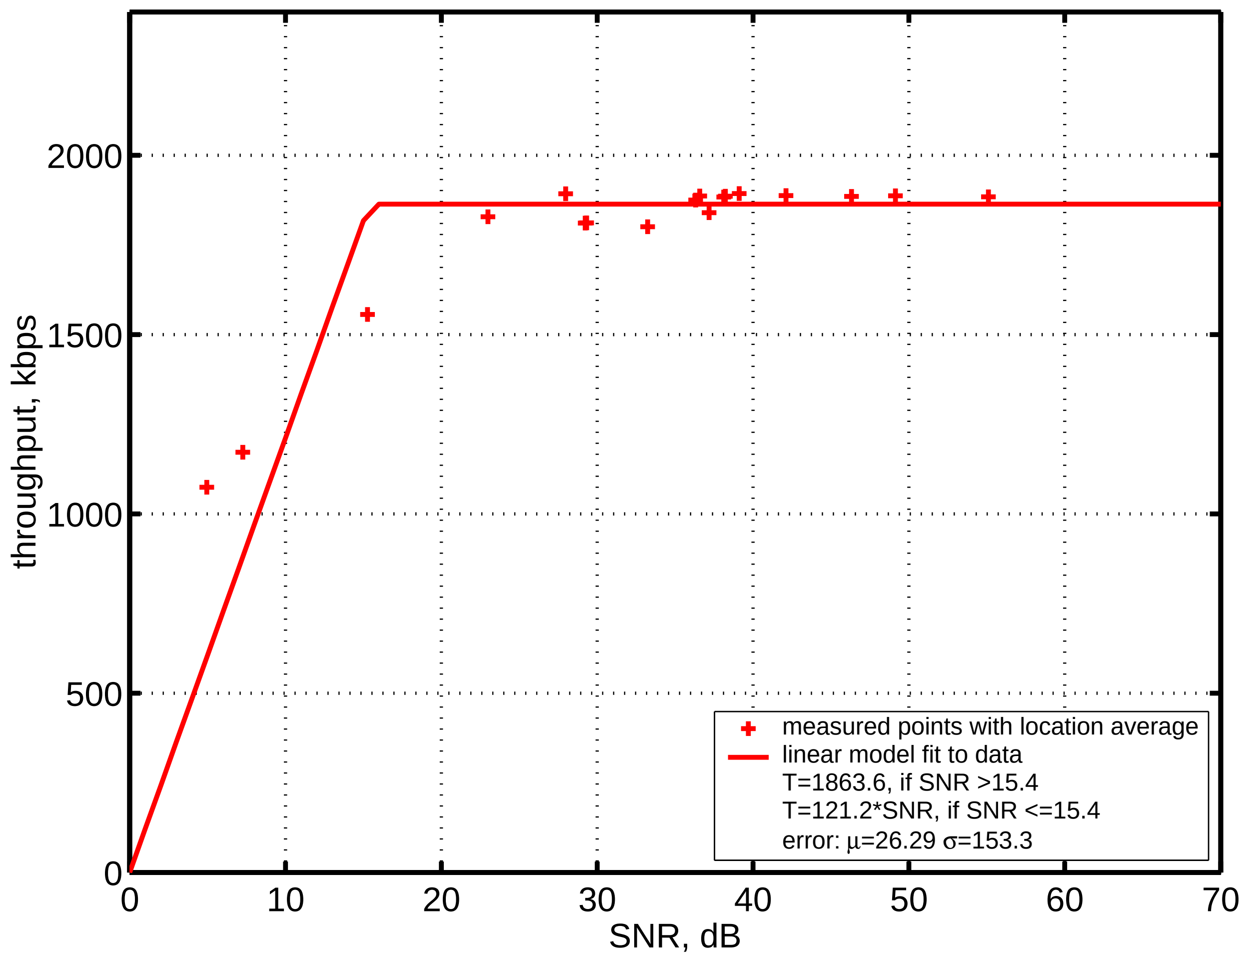

Figure 6.3: WaveLAN Scenario A, Location Averaged Data and Linear Constrained Throughput Prediction Model

6.3.2 Scenario A, Exponential Model

Formula for the Scenario A Exponential Model

After considering a piece-wise linear model, an exponential model was also considered. The advantage of an exponential model is that it is a simple model of the behavior of the wireless LAN throughput which, above a critical SNR value is close to a maximum throughput, T_{max}. Below the critical SNR value, the throughput drops off quickly to zero. However, the exponential model should provide a better match to the data because it provides a more gradual transition around the critical SNR value. The formula used for the exponential model is give in Equation 6.3

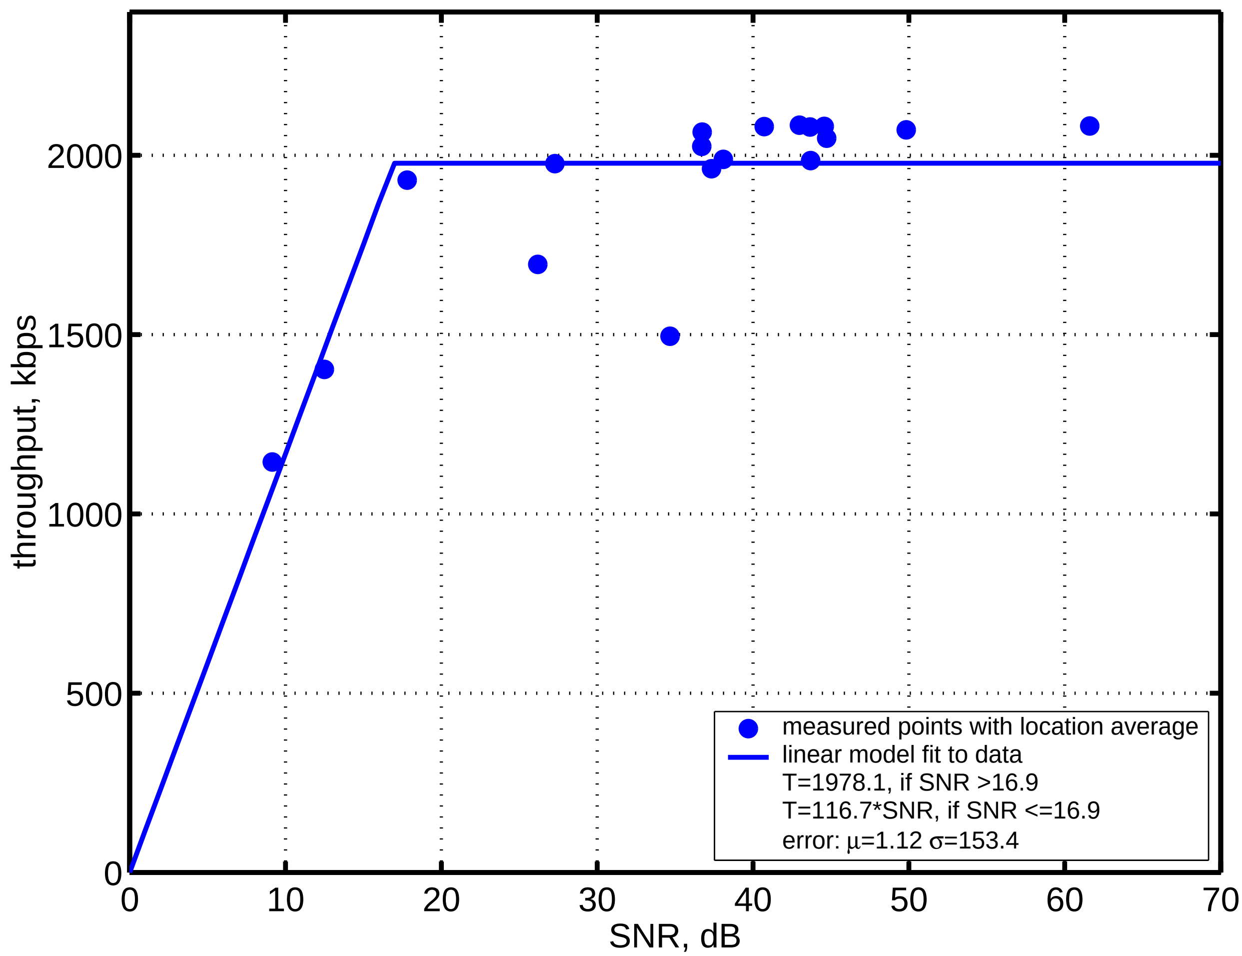

Figure 6.4: 3Com Scenario A, Location Averaged Data and Linear Constrained Throughput Prediction Model

$$T_1 = T_{max}(1 - exp(-\alpha(SNR - SNR_0))) \tag{6.3}$$

As was the case for the Linear Model, the values of T_{max}, α and SNR_0 are application and vendor specific, but can be tuned by measurement. T_{max} is the throughput that the wireless LAN connection will experience in ideal circumstances. α can be thought of as the slope of the exponential curve and governs the rate at which the exponential drops off. Lastly, the parameter SNR_0 is the SNR value at which the wireless LAN provides zero throughput.

Table 6.1: The Values of the Calculated Parameters and Error Statistics for the Scenario A Throughput Prediction Models, as defined in Equations 6.1 and 6.3

Model Name	T_{max}	A or α	T_0 or SNR_0	SNR_c	μ	σ
WaveLAN Linear	1,863.6 kbps	47.1 kbps/dB	835.9 kbps	21.8 dB	0 kbps	31.3 kbps
WaveLAN Linear, Constrained	1836.3 kbps	121.2 kbps/dB	0 kbps	15.4 dB	26.29 kbps	153.3 kbps
WaveLAN Exponential	1907.2 kbps	0.099	-3.0 dB	N/A	0 kbps	34.6 kbps
WaveLAN Exponential, Constrained	1874.7 kbps	0.144	0 dB	N/A	1.71 kbps	48.7 kbps
3Com Linear	1981.3 kbps	91.6 kbps/dB	287.5 kbps	18.5 dB	0 kbps	151.5 kbps
3Com Linear, Constrained	1978.1	116.7 kbps/dB	0 kbps	16.9 dB	1.12 kbps	153.4 kbps
3Com Exponential	2048.6 kbps	0.097	0.2 dB	N/A	0 kbps	147.3 kbps
3Com Exponential, Constrained	2050.9 kbps	0.095	0	N/A	-0.2 kbps	147.3 kbps

Calculating the Exponential Model Parameters

The procedure explained in Section 6.2 was used to fit the parameters of Equation 6.3. The initial guesses selected for the WaveLAN data were T_{max} equal to 1800, α equal to 0.1 and SNR_0 equal to 0. The initial guesses for 3Com were T_{max} equal to 2100, α equal to 0.1 and SNR_0 equal to 0. The model parameters resulting from this fitting process are given in Table 6.1 and the error values are given in Table 6.2. The exponential provided a fairly good fit to the data as can be seen in Figures 6.5 and 6.6.

The equations used to generate the model are shown on the graphs along with the mean and standard deviation in the error. As can be seen from the error statistics and the figures, the exponential provides a good fit to the data with the error statistics very similar to those for the linear model. Again, the mean error has been reduced to 0 kbps and the error standard deviation is less than about 150 kbps. The exponential, however, is closer to intersecting 0 kbps throughput when the SNR is zero than the linear models were.

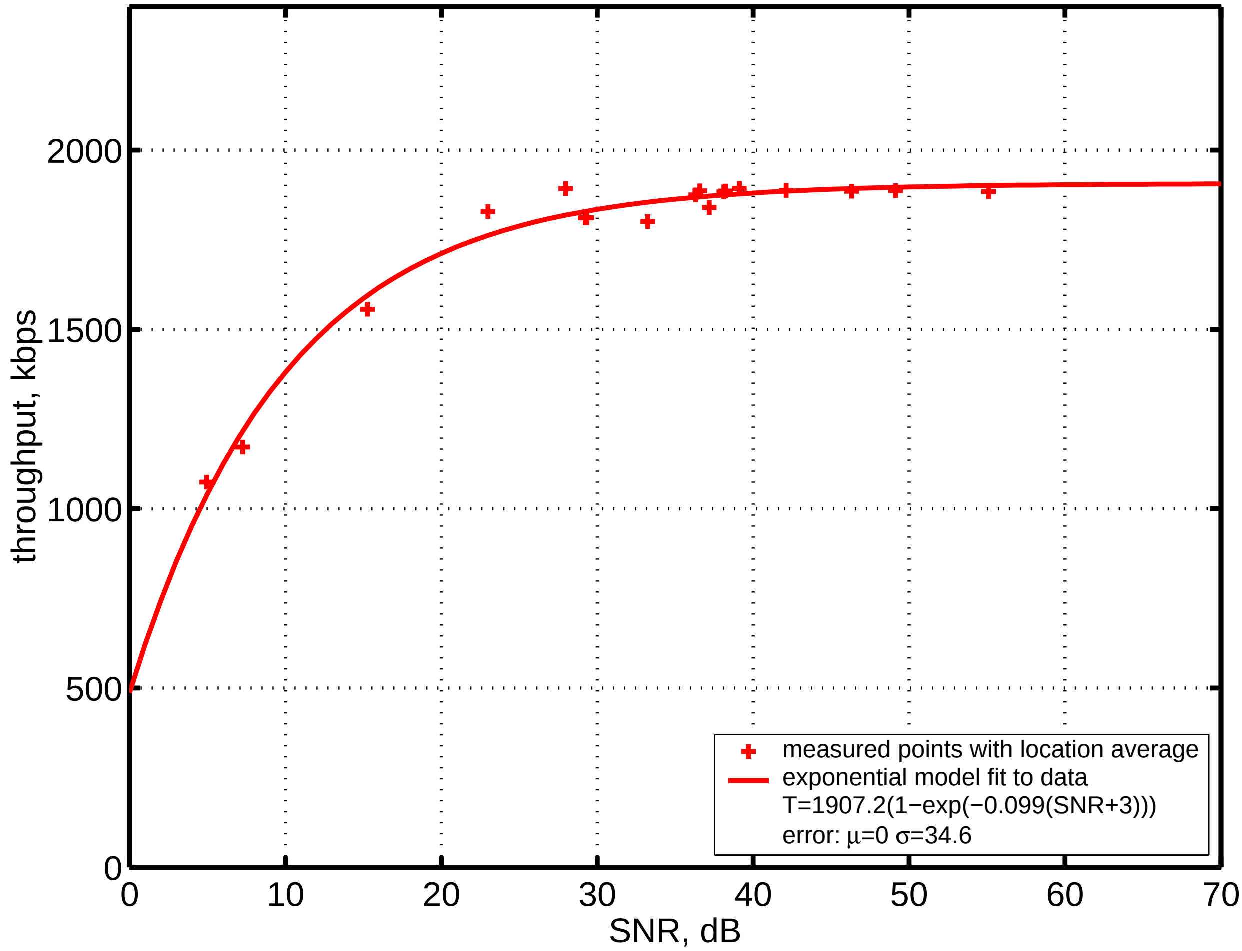

Figure 6.5: WaveLAN Scenario A, Location Averaged Data and Exponential Throughput Prediction Model

Constraining the Exponential Model

As can be seen from Figures 6.5 and 6.6, the exponential model of throughput does not predict zero throughput at a 0 dB SNR. For this reason, the Exponential models were refit to the data, except with SNR_0 forced to be equal to zero. In this manner the prediction models are simpler, more intuitive and more realistic. The model was again fit using MATLAB *nlinfit* command as described in 6.2 and was performed previously in Section 6.3.2 except that SNR_0 was not considered an adjustable model parameter. The resulting error statistics are given in Table 6.2 and the calculated model parameters are given in Table 6.1. Additionally, the resulting constrained models and the measured data are plotted in Figures 6.3 and 6.4.

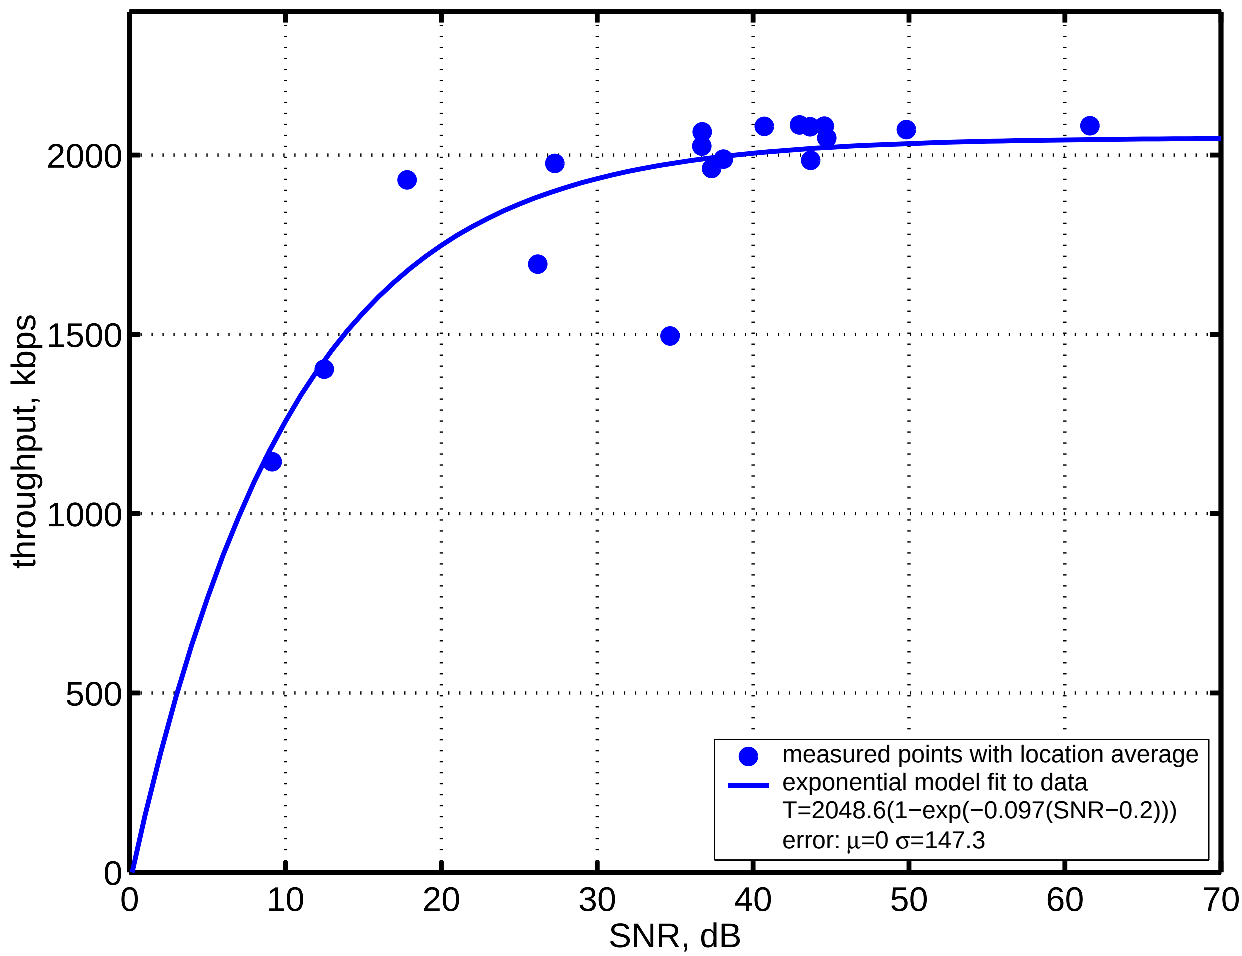

Figure 6.6: 3Com Scenario A, Location Averaged Data and Exponential Throughput Prediction Model

As can be seen from the results of constraining the exponential model, there is a slight mean error introduced and a slight increase in the error standard deviation. Again, the use of this, simpler more intuitive constrained model has resulted in a loss in prediction accuracy.

6.3.3 Foundations for Other Scenario Models

Scenario A has provided a foundation model for use Scenarios B, C and D. Considering all of the results of the model optimizations presented, as summarized in Table 6.2, the Exponential prediction model has proven to be superior to the Linear prediction model.

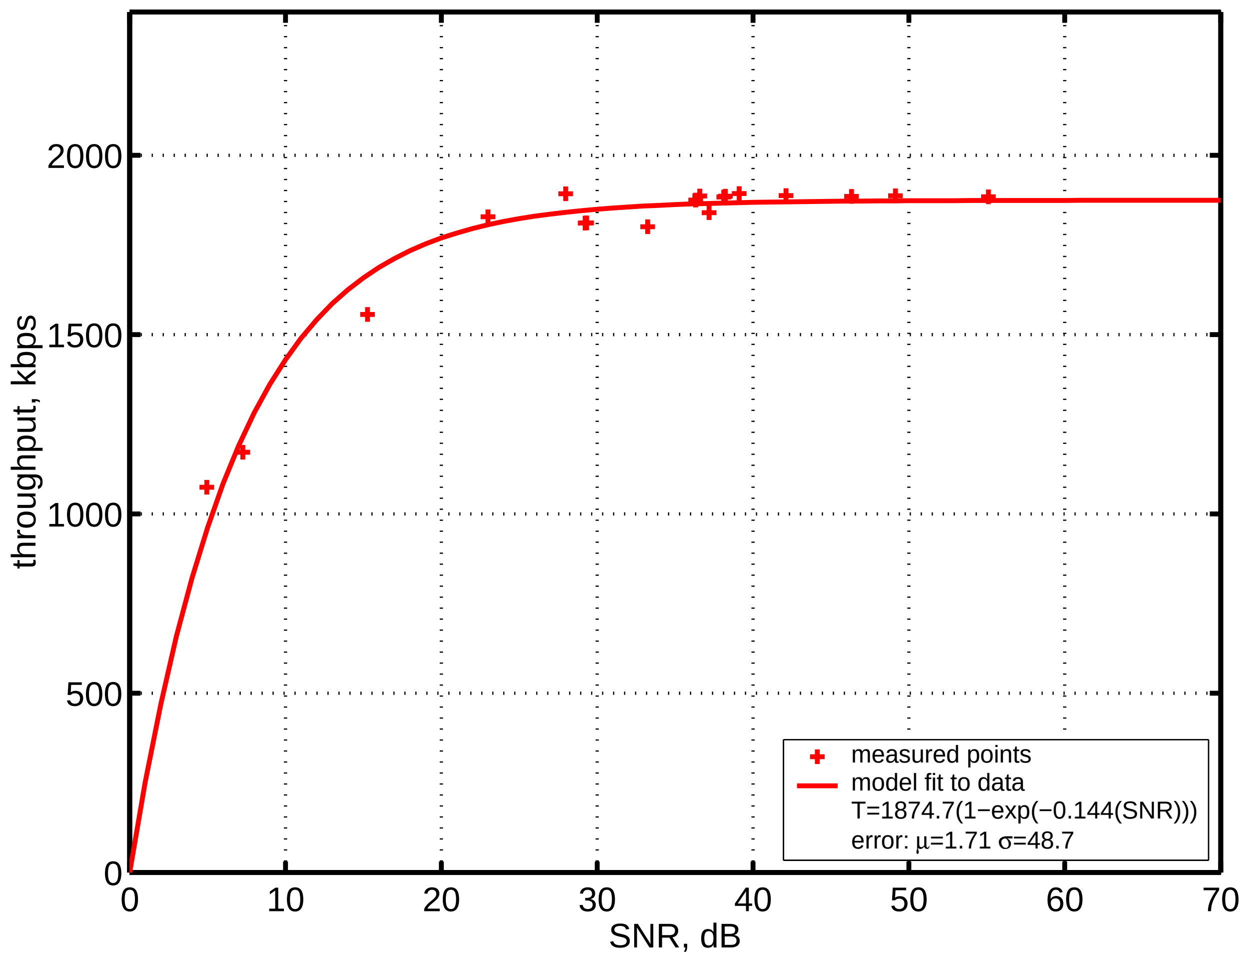

Figure 6.7: WaveLAN Scenario A, Location Averaged Data and Exponential Constrained Throughput Prediction Model

The Exponential model is conceptually simpler to implement than linear model and has fewer associated parameters. In addition, the exponential model provides slightly lower error statistics, especially when constrained and especially for the WaveLAN data. As a result the exponential model will be used as a basis for further prediction models for Scenarios B, C and D.

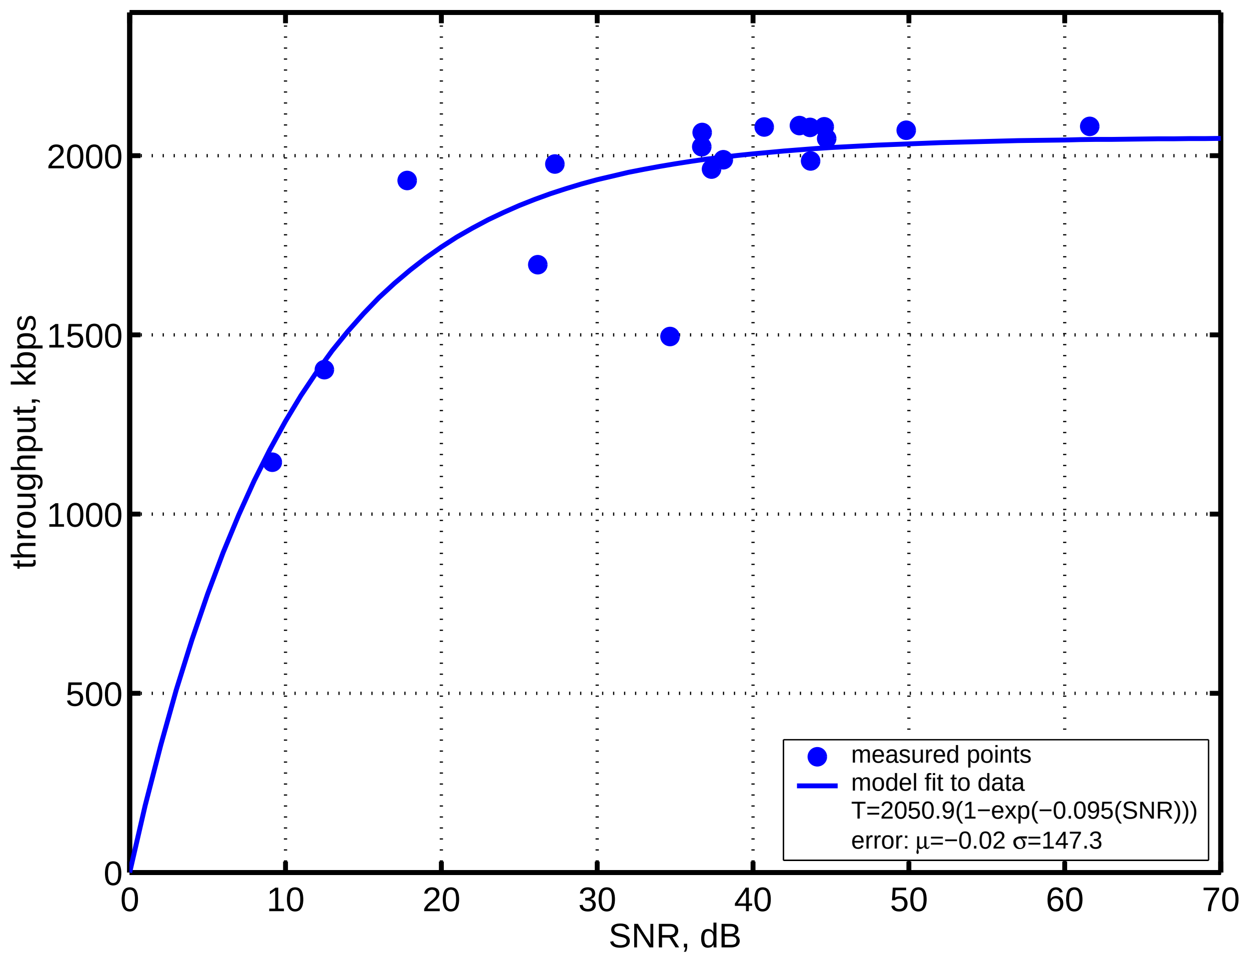

Figure 6.8: 3Com Scenario A, Location Averaged Data and Exponential Constrained Throughput Prediction Model

Table 6.2: A Summary of the Accuracy of the Empirical Throughput Prediction Models.

Scenario	Manufacturer	Model	μ	σ
A	WaveLAN	Linear, Equation 6.1	0 kbps	31.3 kbps
	WaveLAN	Linear, Constrained, Equation 6.1	26.29 kbps	153.3 kbps
	WaveLAN	Exponential, Equation 6.1	0 kbps	34.6 kbps
	WaveLAN	Exponential, Constrained, Equation 6.1	1.71 kbps	48.7 kbps
	3Com	Linear, Equation 6.3	0 kbps	151.5 kbps
	3Com	Linear, Constrained, Equation 6.3	1.12 kbps	153.4 kbps
	3Com	Exponential, Equation 6.3	0 kbps	147.3 kbps
	3Com	Exponential, Constrained, Equation 6.3	-0.2 kbps	147.3 kbps
B	WaveLAN	Two Users, Equation 6.4	0 kbps	185.2 kbps
	WaveLAN	Two Users, Constrained , Equation 6.4	0 kbps	178.9 kbps
	3Com	Two Users, Equation 6.4	0 kbps	225.2 kbps
	3Com	Two Users, Constrained, Equation 6.4	0 kbps	225.2 kbps
C	WaveLAN	One User with Interference, Equation 6.8	0 kbps	254.9 kbps
	WaveLAN	One User, with Interference, Constrained, Equation 6.8	0 kbps	258.3 kbps
	3Com	One User with Interference, Equation 6.8	0 kbps	136.2 kbps
	3Com	One User with Interference, Constrained, Equation 6.8	0 kbps	135.6 kbps
D	WaveLAN	Two User with Interference, Equation 6.10	0 kbps	268.7 kbps
	WaveLAN	Two User, with Interference, Constrained, Equation 6.10	2.65 kbps	253.0 kbps
	3Com	Two User with Interference, Equation 6.10	0.2 kbps	196.1 kbps
	3Com	Two User with Interference, Constrained, Equation 6.10	-10.2 kbps	195.7 kbps

6.4 Prediction Models for Scenario B – Two Users

Scenario B utilizes the basic one client exponential model derived for Scenario A as a base for its models. The two client model for throughput prediction introduces T_{loss} factors to adjust the ideal, one User throughput to match the actual throughput in the presence of another user. The exponential model given in Equation 6.3 was used as the basis for this model, rather than the piece-wise linear model of Equation 6.1, because it is simpler and more intuitive than the linear model. It also lends itself well to be modified by various situations that can cause a loss in throughput for two or more users of a shared wireless LAN.

6.4.1 Throughput Prediction Formula for Two Simultaneous Users

The calculation of the throughput of a client which is sharing the access to a wireless LAN is fairly complex. The basic concept of the throughput prediction model is to start with the normal, ideal wireless LAN throughput and subtract various losses due to known difficulties with wireless LAN systems. The basic formula for two clients is shown in Equation 6.4.

$$T = T_1(SNR_1) - T_{loss-compete}(SNR_2) - T_{loss-hidden}(PL(C_1, C_2)) - T_{offset} \tag{6.4}$$

As can be seen from Equation 6.4, the throughput of an individual client is based on the ideal one user throughput, T_1 from Equation 6.3, which represents the single client, exponential model developed for Scenario A. This single client throughput is modified in Equation 6.4 by two loss functions and a basic throughput offset value, T_offset. SNR_1 is the SNR of the considered client, SNR_2 is the SNR of the other client which is sharing the access point and is simultaneously transmitting data. $T_{loss-compete}(SNR_2)$, is a loss in throughput due to the client having to compete with other users for throughput to the desired access point. The formula for this loss parameter is give in Equation 6.5. $T_{loss-hidden}$ is a loss in throughput due to the hidden terminal problem [12]. The formula for this loss factor is given in Equation 6.6.

$$T_{loss-compete}(SNR_2) = A_{compete} * [T_{max} - T_1(SNR_2)] \tag{6.5}$$

$T_{loss-compete}$ is the loss in throughput due to the congestion caused by both users attempting to use the same access point at the same time. As shown in Equation 6.5, it is a function of the ideal throughput of the other user. Thus, the more the other client is experiencing poor throughput, the more likely that client is to interfere with, and decrease the throughput of the first client. In Equation 6.5, SNR_2 is the SNR of the other client. T_{max} is the maximum throughput under ideal circumstances, as defined by the results of the One Client Exponential Model from Scenario A. T_1 is the ideal exponential throughput calculation function, also defined by the results from Scenario A.

$$T_{loss-hidden}(N) = A_{hidden} * PL(C_1, C_2) \tag{6.6}$$

$T_{loss-hidden}$ is a loss in throughput due to the so called "hidden terminal" problem [12]. The hidden terminal problem is when two clients can not detect the transmissions of one another and so are unable to avoid collisions with one anther's packets. This loss factor is based on the path loss from the main, first client, C_1, to the second client, C_2. This path loss is represented as $PL(C_1, C_2)$. As the path loss between the two clients increases, the clients are less likely to detect the transmissions of one another and thus the probability that packets sent from different clients will collide at the access point increases. It should be noted that the value of A_{hidden} will change if the RTS/CTS functionality (see Appendix A for more information) of the wireless LAN is enabled. This functionality was not enabled for the measurement made in this thesis.

6.4.2 Calculating Parameters for the Two User Model

Due to the complexity of the two user model shown in Equation 6.4, it was imperative that the *nlinfit* command in MATLAB be used to compute the best values for the parameters of the model. However, before using the same MATLAB solution technique as was applied in Scenario A using *nlinfit*, it was first necessary to obtain the path loss between different client locations, $PL(C_1, C_2)$. These values are used in Equation 6.6. The path loss values were found using the SitePlanner RF Prediction model previously calculated and utilized in Section 5.7.1 and given in Equation 5.1. This prediction model allowed the accurate calculation of the path loss between different measurement locations. After the path losses between all client locations were known, it is possible to calculate the optimal values for the two client prediction models. The results of the MATLAB calculations to calculate the optimal model parameters are given in Table 6.3. The error statistics for the model given in Equation 6.4 with the optimized model parameters are give in Table 6.2. All this information is also shown on Figures 6.9 and 6.10. These figures provide graphical insight into the accuracy of this above two client throughput prediction model.

Table 6.3: The Values of the Calculated Parameters and Error Statistics for the Scenario B Throughput Prediction Models, as given in Equation 6.4

Model Name	$A_{compete}$	A_{hidden}	T_0	μ	σ
WaveLAN 2 Client Model	0.359	3.48 kbps/dB	-81 kbps	0 kbps	185.2 kbps
WaveLAN 2 Client Model, Constrained	0.309	3.84 kbps/dB	-91.8 kbps	0 kbps	178.9 kbps
3Com 2 Client Model	0.034	6.44 kbps/dB	60.6 kbps	0 kbps	225.2 kbps
3Com 2 Client Model, Constrained	0.072	6.46 kbps/dB	64.7 kbps	0 kbps	225.2 kbps

6.4.3 Constraining the Two User Throughput Measurement Model

In order to reduce the number of parameters needed for the two client model given in Equation 6.4, the model was constrained as was done for the Single client model. That is, the value of $T_1(SNR)$ is forced to be zero when SNR is zero. $T_1(SNR)$ is the single user model given in Equation 6.3 and can be constrained in the desired manner by simply setting the T_0 parameter of this equation to zero. The parameters of the two user model are then re-computed using the MATLAB *nlinfit* command. The results of this constrained optimization are given in Tables 6.3 and 6.2. Additionally, the success of this technique can be viewed in Figures 6.11 and 6.12 where the theoretical and actual throughput are compared.

6.4.4 Extending the Two User Model to N Users

The two user model presented in Equation 6.4 could be generalized to multiple users by calculating the $T_{loss-compete}$ and $T_{loss-hidden}$ values for several other operating clients rather

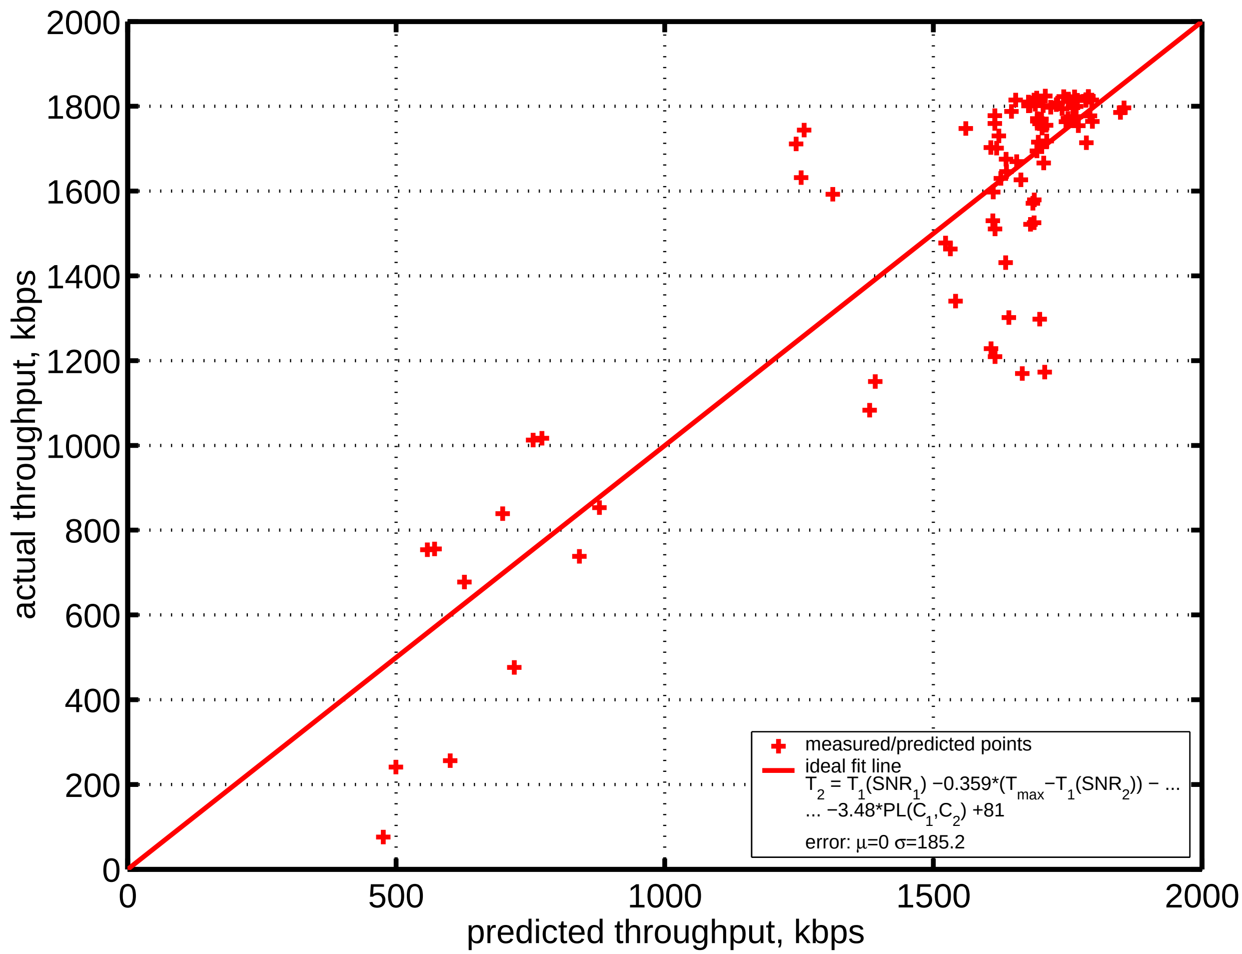

Figure 6.9: WaveLAN Scenario B, 2 User Prediction Model. This figure shows the predicted versus actual throughput results and an ideal fit line

than just one other client. A modified formula to predict the performance of a wireless LAN with N total users is presented in Equation 6.7. Further modifications of this formula may be necessary, but measurements need to be performed to verify this.

$$T = T_1(SNR_1) - \sum_{j=2}^{N} [T_{loss-compete}(SNR_j) + T_{loss-hidden}(PL(C_1, C_j))] - T_{offset} \quad (6.7)$$

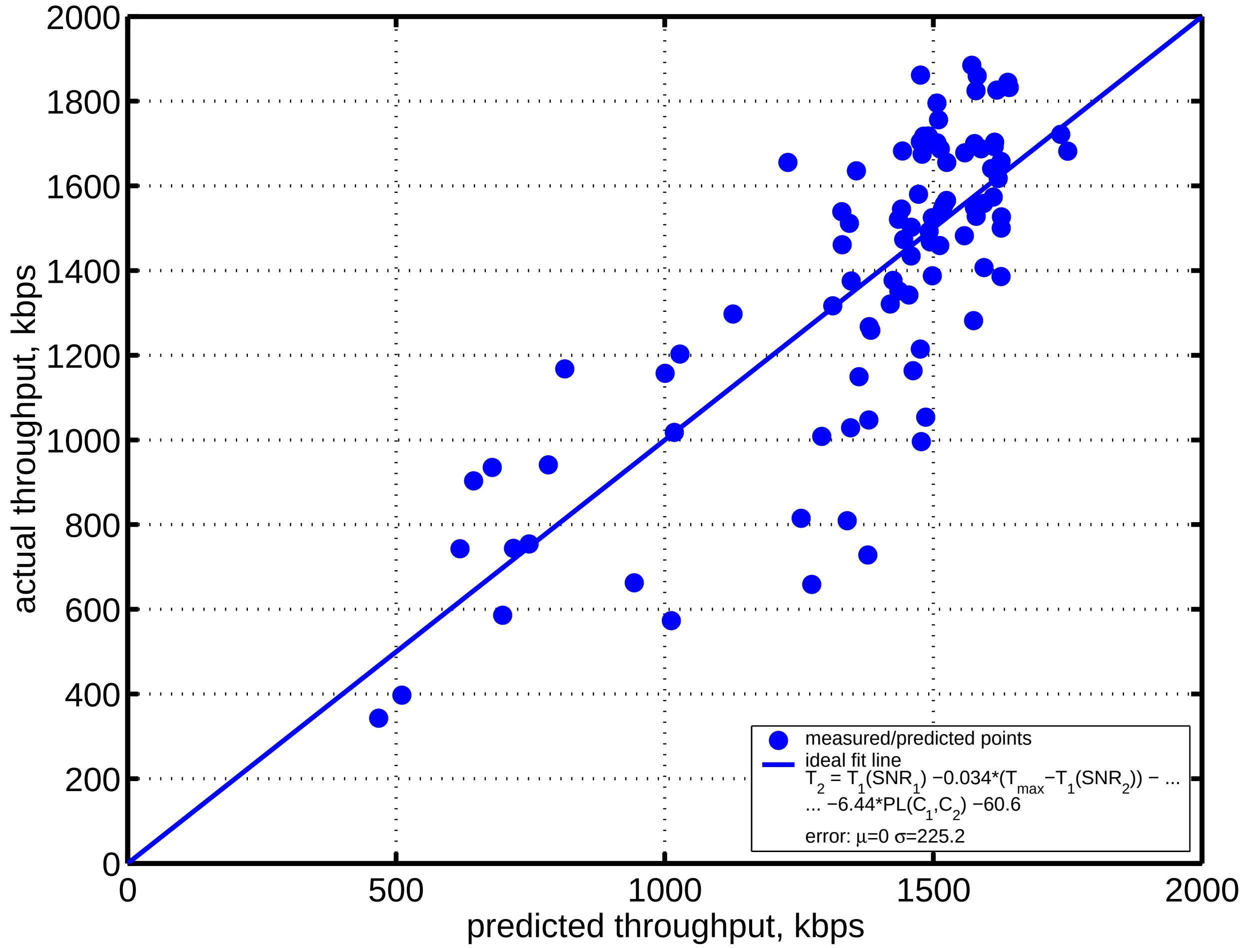

Figure 6.10: 3Com Scenario B, 2 User Prediction Model. This figure shows the predicted versus actual throughput results and an ideal fit line

6.4.5 Analysis of the Two User Throughput Prediction Model

The results of the fitting technique for this model show the results to be fairly accurate. This combined with the simple, intuitive nature of the prediction model make it a valuable contribution for wireless LAN installation design and deployment. However, there are two other very important parameters that need to be considered that were neglected in the two user model (Equation 6.4) due to a lack of available data:

1. The first of these is the saturation of the available bandwidth. For all measurements taken and used to derive these models, there was no evidence that the clients were being limited by the users utilizing all of the available bandwidth at a single access

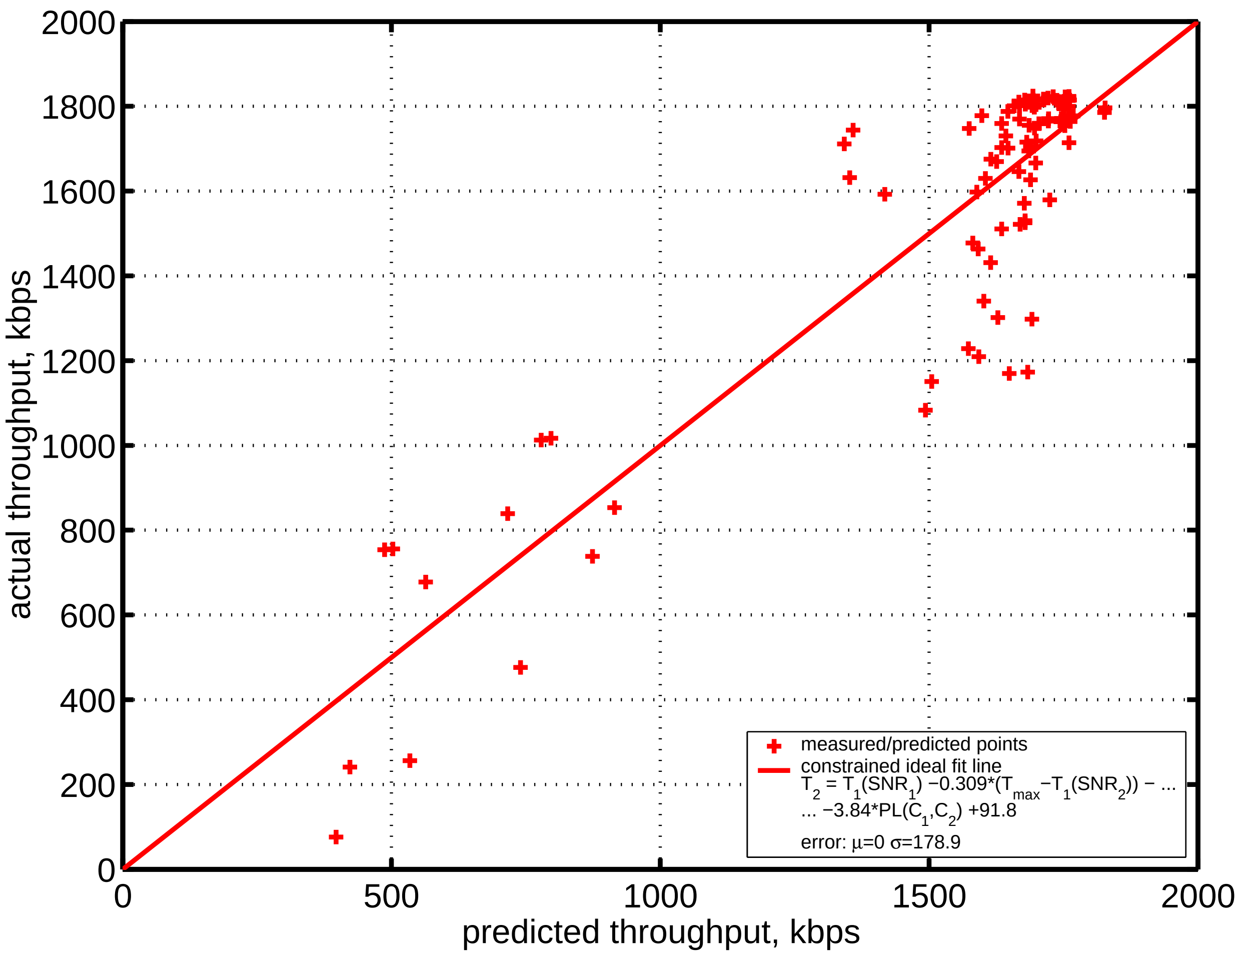

Figure 6.11: WaveLAN Scenario B, 2 User Constrained Prediction Model. This figure shows the predicted versus actual throughput results and an ideal fit line

point. The $T_{loss-compete}$ parameter could be used with some modification to model the saturation that occurs when all available bandwidth is utilized. The typical saturation point of wireless LANs has been reported to occur at approximately 2 or 3 users operating at maximum usage [1]. Therefore it was not possible to model the saturation that occurs using the measured data presented in this thesis.

2. The second unconsidered difficulty is the user access patterns. All measurements taken for this thesis had users attempting to access the wireless network as often as possible. Users were always trying to transmit or receive in a continuous fashion, thus the probability of two users attempting to transmit at the same time was extremely high. However, it would extremely useful if an Erlang or similar traf-

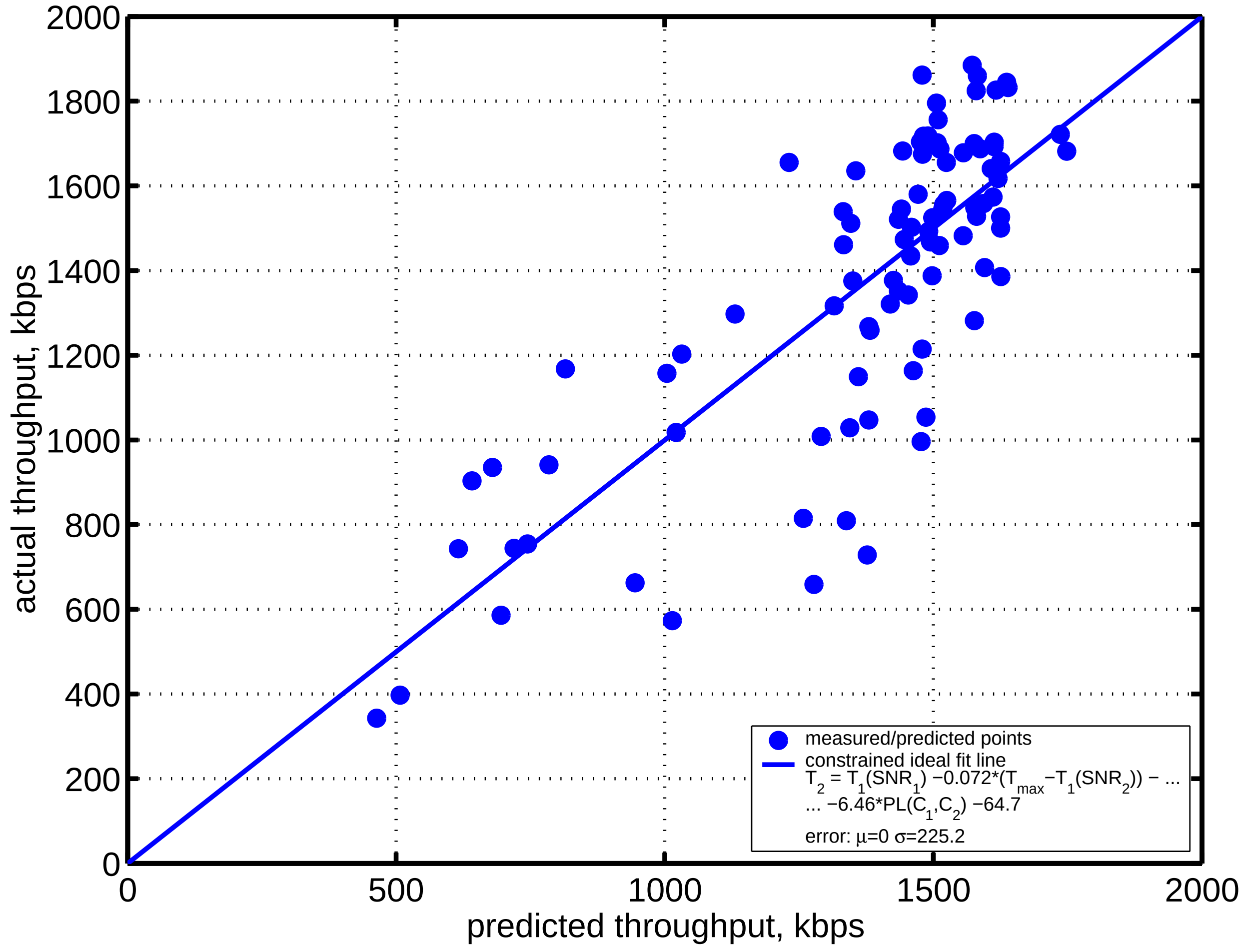

Figure 6.12: 3Com Scenario B, 2 User Constrained Prediction Model. This figure shows the predicted versus actual throughput results and an ideal fit line

fic model were applied to allow wireless LANs to calculate an actual user capacity based on the shared nature of a network in which users are not attempting to transmit continuously. This sort of analysis is beyond the scope of this thesis.

6.5 Prediction Models for Scenario C – One User with Interference

Scenario C builds directly from the exponential Scenario A model given in Equation 6.3. The presence of an interfering client tends to lower the throughput of the desired client. As a result, the model developed for Scenario C was a modified version of the exponential model developed in Scenario A. The exponential model was chosen over the linear model because the exponential has fewer parameters and is more intuitive to use. The model developed below introduces a $T_{loss-int}$ factor similar to the T_{loss} factors introduced in the Scenario B model (Equation 6.4). This loss factor, $T_{loss-int}$ models the additional losses in throughput due to the presence of interference.

6.5.1 Quantifying the Interference Power

In order to predict the throughput performance of a wireless LAN in the presence of an interfering wireless LAN system, the interference must be quantified. This is more challenging than it might first appear. The difficulty is that all interference measurements were made with another wireless LAN system in use as the interferer. As a result, each measurement had two sources of interference and two recipients of interference. This is shown diagrammatically in Figure 5.26, where the intentional transmissions are represented by the dashed lines. The interfering client and interfering access point are attempting to communicate with one another while the desired client and desired access point attempt to do the same thing. Additionally, there is intentional interference between the interfering system and the desired system. In order to accurately predict the throughput experienced between the desired client and the desired access point, four signal strengths must be considered. That is, the interference power at both the desired client and the desired access point from both the interfering client and the interfering server need to be taken into account. To do this, a combination of measurements and predictions, as outlined in Section 5.7.2 were performed to find four RSSI values for each throughput measurement. Section 5.7.2 also described how to calculate a net SIR using Equation 5.3. Thus, using the equation, each throughput measurement has an associated net SIR value associated with it. These net SIR values are used in the throughput prediction model for Scenarios C and D.

6.5.2 Formulation of the Scenario C Model

A $T_{loss-int}$ parameter is introduced in this section for the Scenario C, single client plus interference throughput prediction model. $T_{loss-int}$ is intended to model the additional loss in throughput caused by the interference. The main formula for the model is shown in Equation 6.8. In Equation 6.8 a user's throughput is predicted from the theoretical SNR and the measured SIR values for the client's location. The SNR value is the signal to noise measurement made when the interfering client was not operating. Referring to Figure 5.26, the SIR value is actually based on the net effect of all the interference power from an interfering client and an interfering access point on the desired client and desired access points. A detailed treatment of how to compute the SIR parameter is presented in Sections 6.5.1 and 5.7.2. Also in Equation 6.8, $T_1(SNR)$ is the exponential equation from scenario A and is defined in Equation 6.3. $T_{loss-int}$ is a new parameter and is the loss due to the interference from another operating wireless LAN system.

$$T_i(SNR, SIR) = T_1(SNR) - T_{loss-int}(SIR) \tag{6.8}$$

The basis for the $T_{loss-int}$ parameter of Equation 6.8 os that as the SIR at a user location decreases (i.e. as the interference power from another wireless LAN system increases) the throughput of the client will decrease. Thus, a linear dependence on the SIR is used to model $T_{loss-int}$. This formulation is given in Equation 6.9.

$$T_{loss-int}(SIR) = A_i * SIR + B_i \tag{6.9}$$

Equation 6.9 is a linear fit to the measured data. This fit is accomplished using the parameters A_i and B_i. A_i is a slope in units of kbps/dB and B_i is an intercept point in units of kbps. Note that A_i is expected to be a negative value because the loss in throughput is inversely proportional to the user's SIR.

6.5.3 Interference Model Parameters

For Equations 6.8 and 6.9, SIR and SNR are dB values, T_{loss_int}, T_1 and T_i are in kbps. The parameters for $T_1(SNR)$ are those found to the optimum values from the data in

Scenario A. All parameters vary depending on the wireless LAN vendor, and may vary for different network applications.

For Equations 6.8 and 6.9, it is important to note that SIR can be negative but still result in a non-negative throughput, T_i. This is reasonable because SIR is not based on the power which is passed through the wireless LAN despreader and bandwidth filter. Recall that the desired system was always used on channel 1, but the interference client was always operated on channel 3. Thus, nearly half of the power transmitted by the interfering system is not in the same frequency band as the desired system. In addition, the spreading codes used by the IEEE 802.11 wireless LANs are designed to have low cross-correlations. As a result, a wireless LAN could potentially operate despite a negative SIR value. The notion of a wireless LAN being able to operate despite the presence of a stronger interfering signal than the desired signal (i.e. a negative SIR) is somewhat counter-intuitive. However, this is because an easily predicted, shifted value for the SIR is intentionally used. By using an easily predicted model for SIR , rather than an SIR corrected (i.e. shifted) for cross-correlations, channel bandwidth overlaps, and other factors, it is easier to directly implement the resulting throughput prediction model in a software tool, such as SitePlanner, which can predict the received signal strength at a given location, but not the complexities involved in the cross-correlation of different IEEE 802.11 spreading codes.

6.5.4 Calculating the Interference Model Parameters.

The same optimization technique was used with Equations 6.8 and 6.9 of this model, as was used for the prior two models. The MATLAB *nlinfit* function was used to compute ideal values for the model parameters A_i and B_i. The values for T_{max}, α and SNR_0 remain unchanged from the optimal values found from the data given in Scenario A. Note that in this case, a simple MMSE fit [20] could have been used. However, the *nlinfit* is equally effective and in this case is guaranteed to find a best fit, rather than a local minima. The resulting values for A_i and B_i are given in Table 6.4. Additionally, the mean and standard deviations of the resulting error in the equations are given in Table 6.2. Figures 6.13 and 6.14 give an indication for how well the throughput prediction model matches the measured data. These figures show a best fit line running across the diagonal and allow the comparison of the predicted and measured data points.

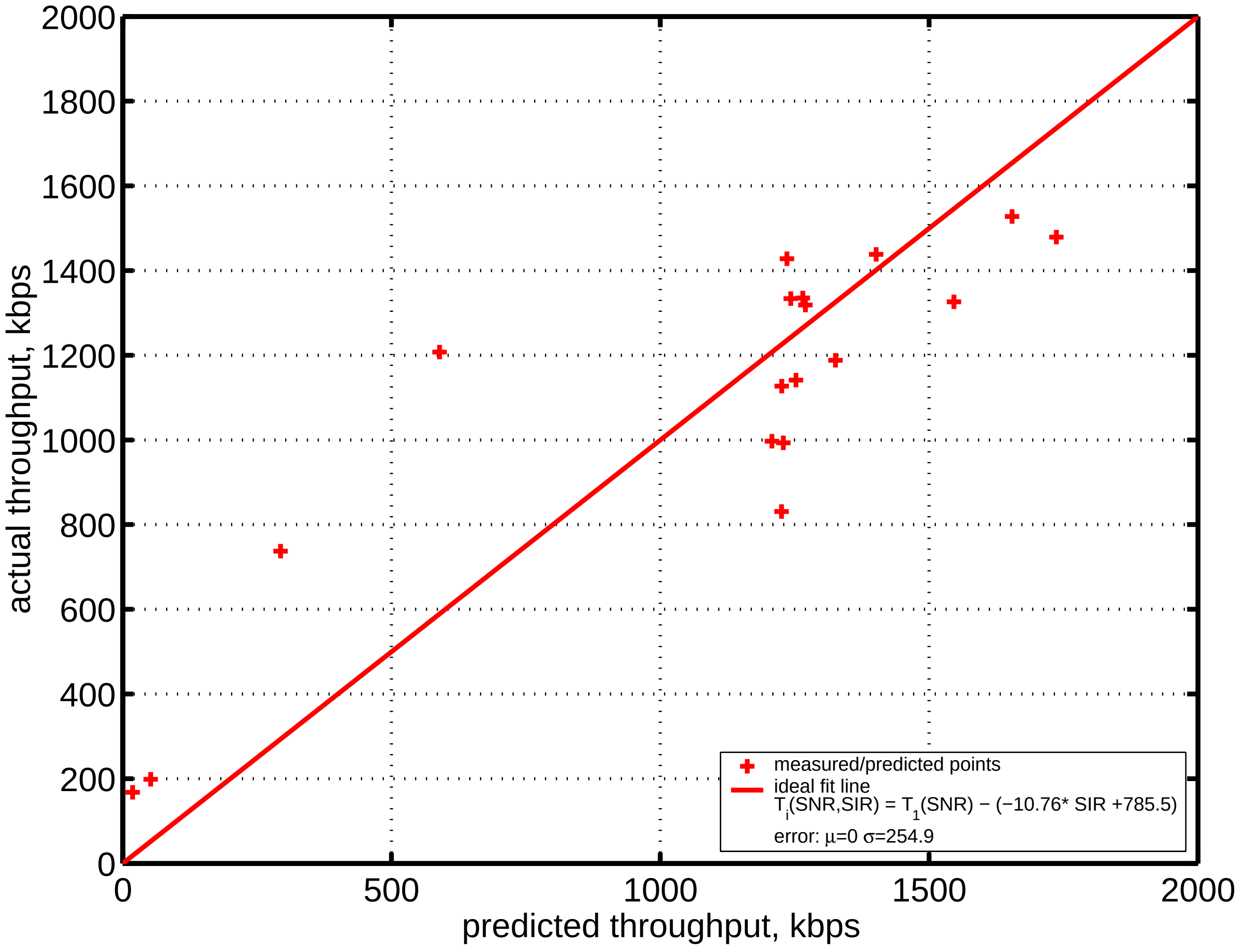

Figure 6.13: WaveLAN Scenario C, One User with Interference, throughput prediction Model. This figure shows the predicted versus actual throughput results and an ideal fit line

6.5.5 Constraining the One User with Interference Model

As was done for the Scenario B models, the Scenario C throughput prediction model uses Equations 6.8 and 6.9 and was constrained to force the case where zero SNR, and infinite SIR results in zero throughput. This is conceptually the same as forcing SNR_0 in Equation 6.3, which defines $T_1(SNR)$, to be zero. This simplifies the single user with interference, throughput prediction model to be fully specified by 4 parameters. To constrain the One User with Interference model of Equation 6.8, $T_1(SNR)$ was configured to use the constrained exponential version of the One User, No Interference model. Then, the process of finding a best fit for A_i and B_i was repeated to find a constrained versions of the two

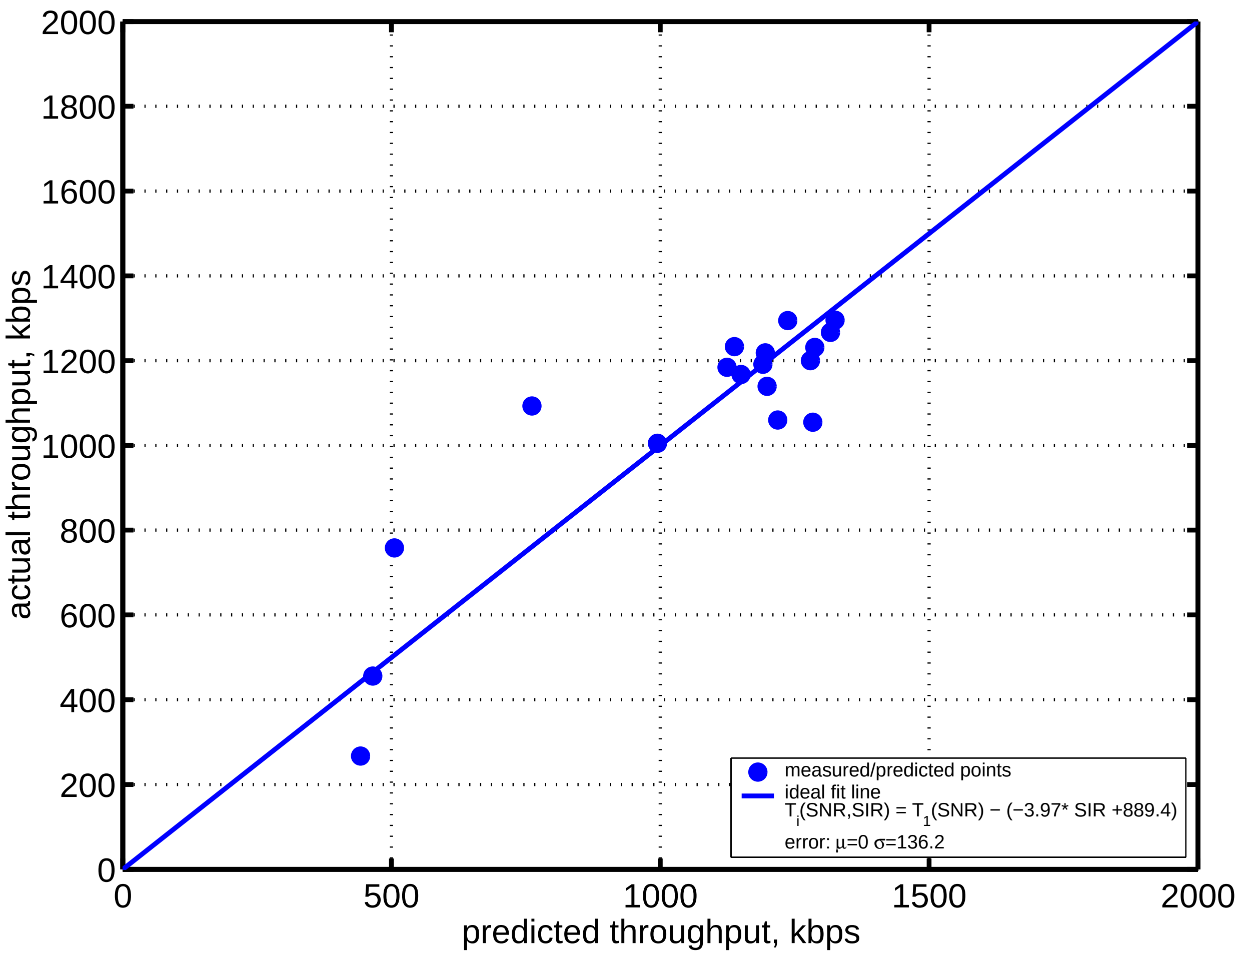

Figure 6.14: 3Com Scenario C, One User with Interference, throughput prediction model. This figure shows the predicted versus actual throughput results and an ideal fit line

parameters. These new parameters are presented in Table 6.4 and in new plots comparing the measured and predicted versions of throughput. These plots are shown in Figures 6.15 and 6.16.

6.5.6 Analysis of the One User with Interference Model

The single user with interference model presented in Equations 6.8 and 6.9 benefits from the fact that it is extremely intuitive and simple to compute. The single user interference model needs only two new parameters (i.e. A_i and B_i in Equation 6.8) and a single new variable (i.e. the SIR). However, despite only adding a single new variable compared to

Table 6.4: The Calculated, Ideal Values and Error Statistics for the Scenario C, Throughput Prediction Model, given in Equations 6.8 and 6.9 for One User with WLAN Interference.

Manufacture/Model	A_i	B_i	μ	σ
WaveLAN	-10.76 kbps/dB	785.5 kbps	0 kbps	254.9 kbps
WaveLAN, Constrained	-11.2 kbps/dB	788.9 kbps	0 kbps	258.3 kbps
3Com	-3.97 kbps/dB	889.4 kbps	0 kbps	136.2 kbps
3Com, Constrained	-4.32 kbps/dB	905.1 kbps	0 kbps	135.6 kbps

the single user model from Scenario A, it still provides a fairly accurate fit to the measurement data. The accuracy could perhaps been better, but not significantly so, even if several parameters were added to Equation 6.8 predict the throughput experienced by a user. Thus, the simple, powerful model of throughput presented in Equations 6.8 and 6.9 is extremely valuable for wireless LAN design and installations.

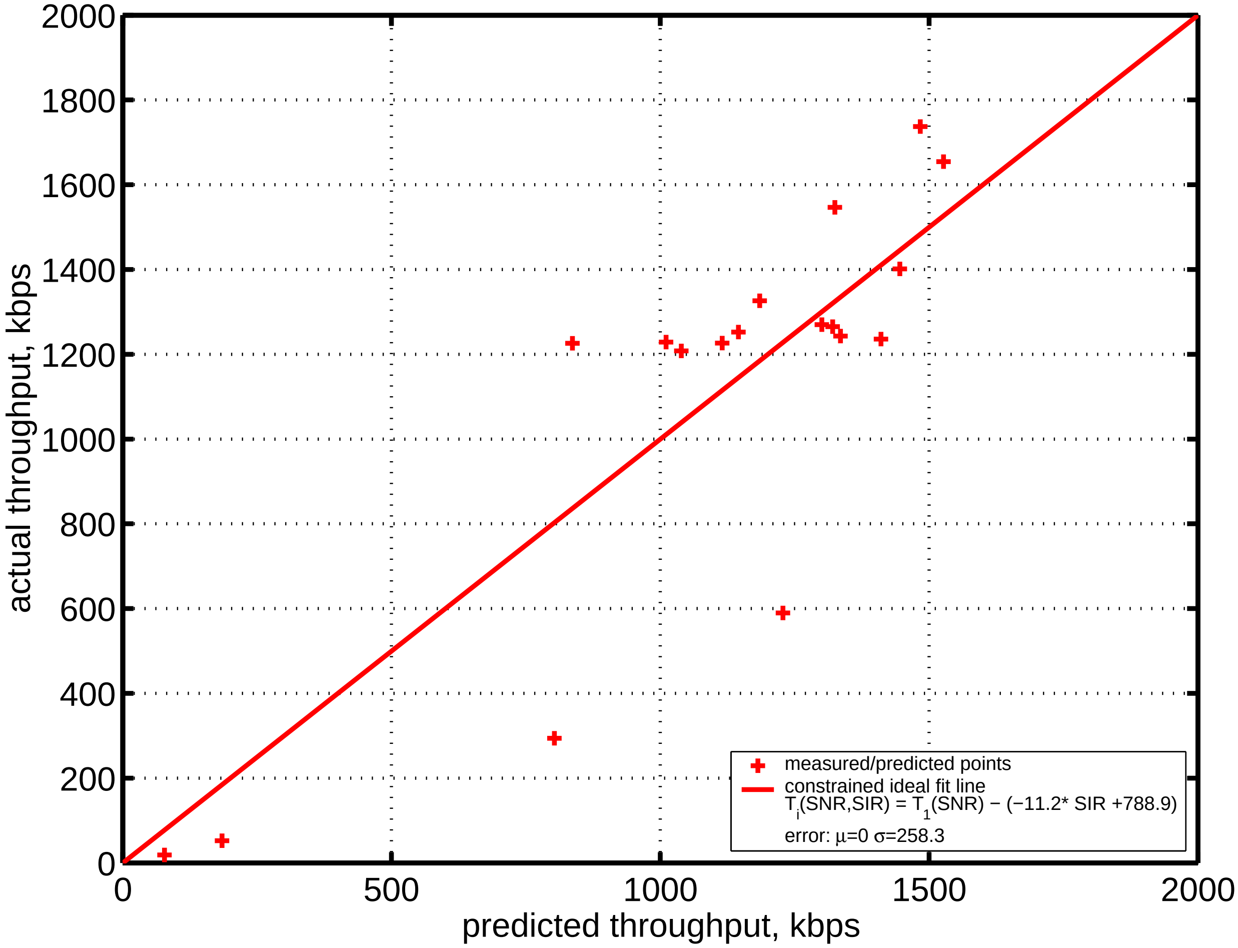

Figure 6.15: Constrained WaveLAN Scenario C, One User with Interference, throughput prediction Model. This figure shows the predicted versus actual throughput results and an ideal fit line

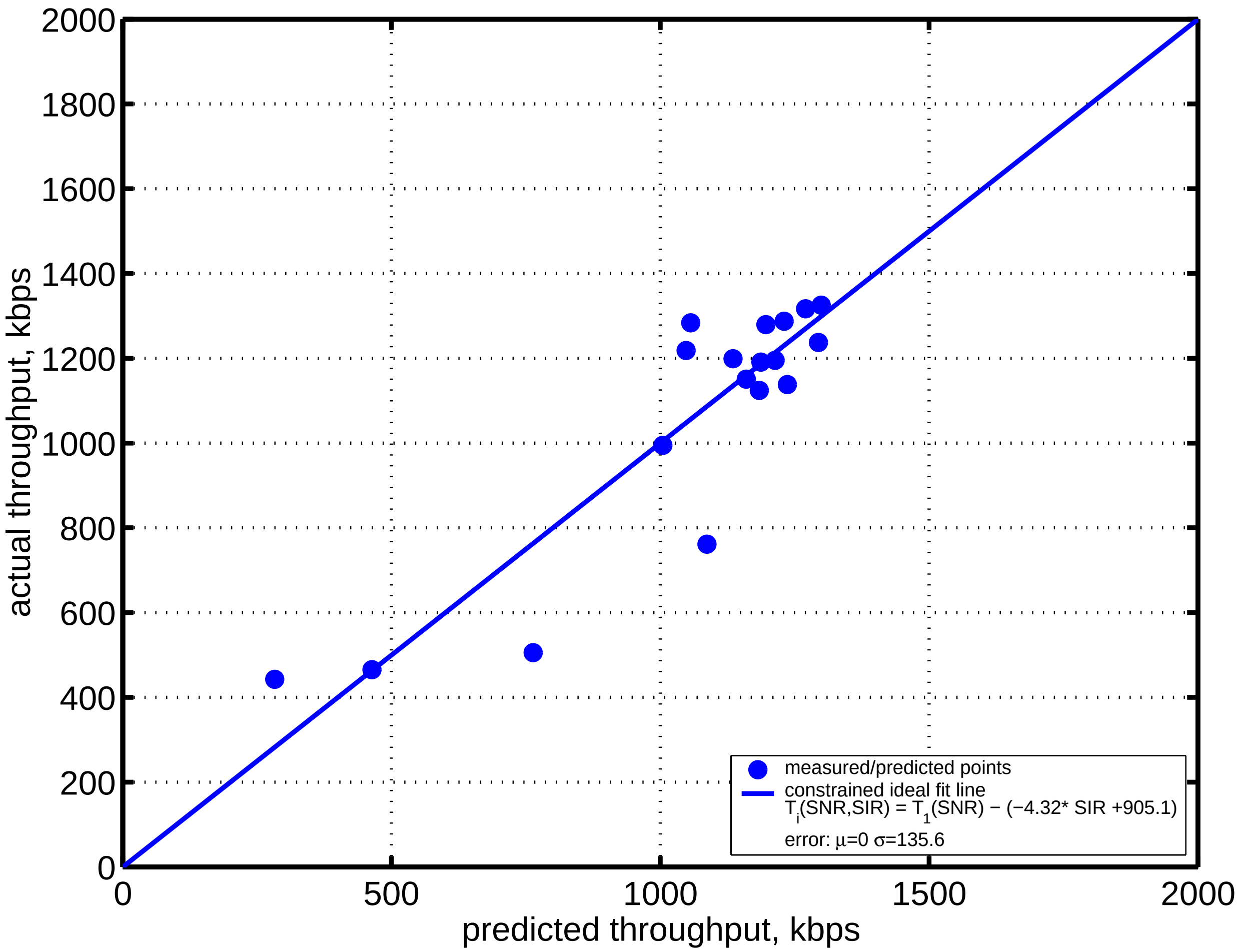

Figure 6.16: Constrained 3Com Scenario C, One User with Interference, throughput prediction model. This figure shows the predicted versus actual throughput results and an ideal fit line

6.6 Prediction Models for Scenario D – Two Users with Interference

Scenario D has been created as a synthesis of Scenarios B and C. By combining the models for throughput from Scenarios B and C, (given respectively in Equations 6.4 and 6.8) into a single model, a validation of the all models is possible. For this final Scenario, no non-linear curve fitting was used. Instead, the models created form the work done in Scenarios B and C were combined into a single model presented in Equation 6.10 in which both the effect of interference and multiuser competition is taken into account as shown below in Equations 6.14 and 6.11. This section presents the work performed in comparing the non-optimized model with the measurements made for this Scenario D.

6.6.1 Creating the Two Client with Interference Model

In order to create the throughput prediction model for Scenario D, the Scenario B and Scenario C models were combined into a single model with a new offset value, T'_{offset}, which aligned the predicted and measured data. The final equation models throughput for two clients with interference and is presented in Equation 6.10.

$$T = T_1(SNR_1) - T'_{loss-compete} - T'_{loss-hidden} - T'_{loss-int} - T'_{offset} \tag{6.10}$$

The Scenario D Model of Equation 6.10 is based on $T_1(SNR_1)$, which is the Scenario A Model given in Equation 6.3. Since there are two clients in this Scenario, SNR_1 is the signal to noise ratio of the considered client, as opposed to SNR_2 which is the signal to noise ratio of the client sharing the access point with the considered client. Note that both clients are desired, not inferring clients.

Equation 6.10 is made up of all the T_{loss} parameters as the Scenario B and Scenario C equations combined. However, each of these factors have been slightly modified to keep the equations analogous to their original intent. For example, the $T_{loss-equation}$ has been modified, as can be seen in Equation 6.11, to use the same optimized parameter $A_{compete}$ as was found in Scenario B. $A_{compete}$ is still multiplied by the theoretical throughput of the other user. However, the other user's theoretical throughput is now calculated using the

Scenario C equation which takes into account the effect of the interference on the other client's theoretical throughput.

$$T'_{loss-compete} = A_{compete} * (T'_{max} - [T_1(SNR_2) - (A_i * SIR_2 + B_i)]) \quad (6.11)$$

T_{max} is also adjusted to be equal to the maximum theoretical throughput that will be experience in the interference environment, as is shown in Equation 6.12.

$$T'_{max} = T_{max} - (A_i * max(SIR) + B_i); \quad (6.12)$$

No modifications were made to the hidden terminal adjustment equation, Equation 6.13 nor the interference adjustment equation, Equation 6.14, which are both presented again here in Equations 6.13 and 6.14.

$$T'_{loss-hidden} = A_{hidden} * PL(C_1, C_2) \quad (6.13)$$

$$T'_{loss-int} = (A_i * SIR_1 + B_i) \quad (6.14)$$

Lastly, T_0, the throughput intercept value was adjusted to align the measured and predicted results. The value of T_0 as well as all other variable parameters are shown in Table 6.5. Note that the values of these parameters are the results of the optimization procedures carried out for Scenarios B and C. Thus, the values in Table 6.5 are identical to those used in Scenarios B and C as given in Tables 6.3 and 6.4 respectively.

6.6.2 Analysis of the Two User with Interference Model

The results of the modifications to the Scenario B and Scenario C equations are illustrated in Figures 6.17, 6.18, 6.19, and 6.20. As can be seen from the graph, there is a quite good agreement in the predictions versus the measured throughput values. The mean error and error standard deviations of the resulting measurements are shown in Table 6.2 and confirm these results.

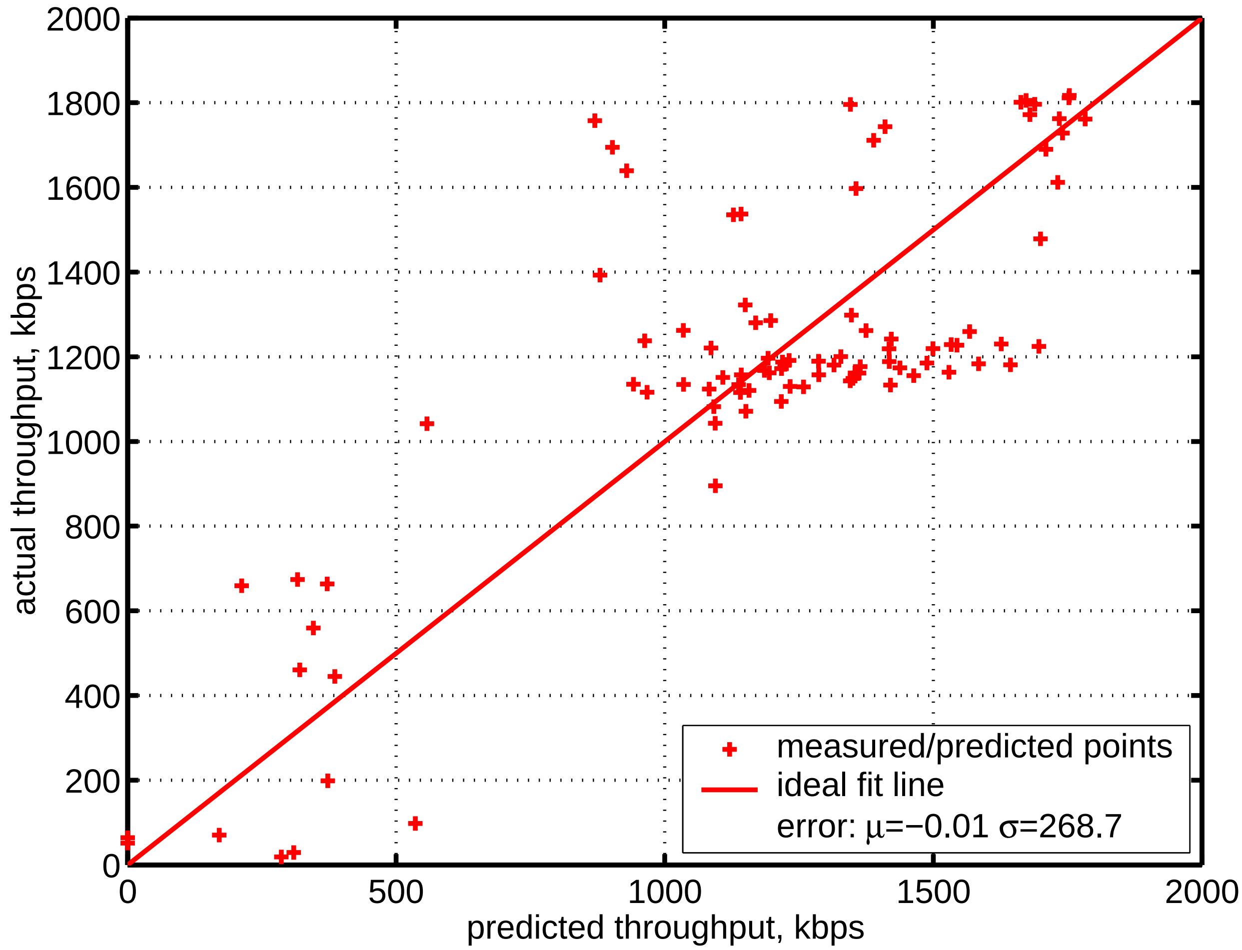

Figure 6.17: WaveLAN Scenario D, 2 User with Interference, throughput prediction Model. This figure shows the predicted versus actual throughput results and an ideal fit line

Table 6.5: Table of Parameter Values Used for the Scenario D Prediction Model for Equation 6.10

Parameter	WaveLAN Value	3Com Value
$A_{compete}$	0.359	0.034
A_{hidden}	3.48	6.44
A_i	-10.76	-3.97
B_i	785.5	889.4
$T_0{}'$	477.34	165.37

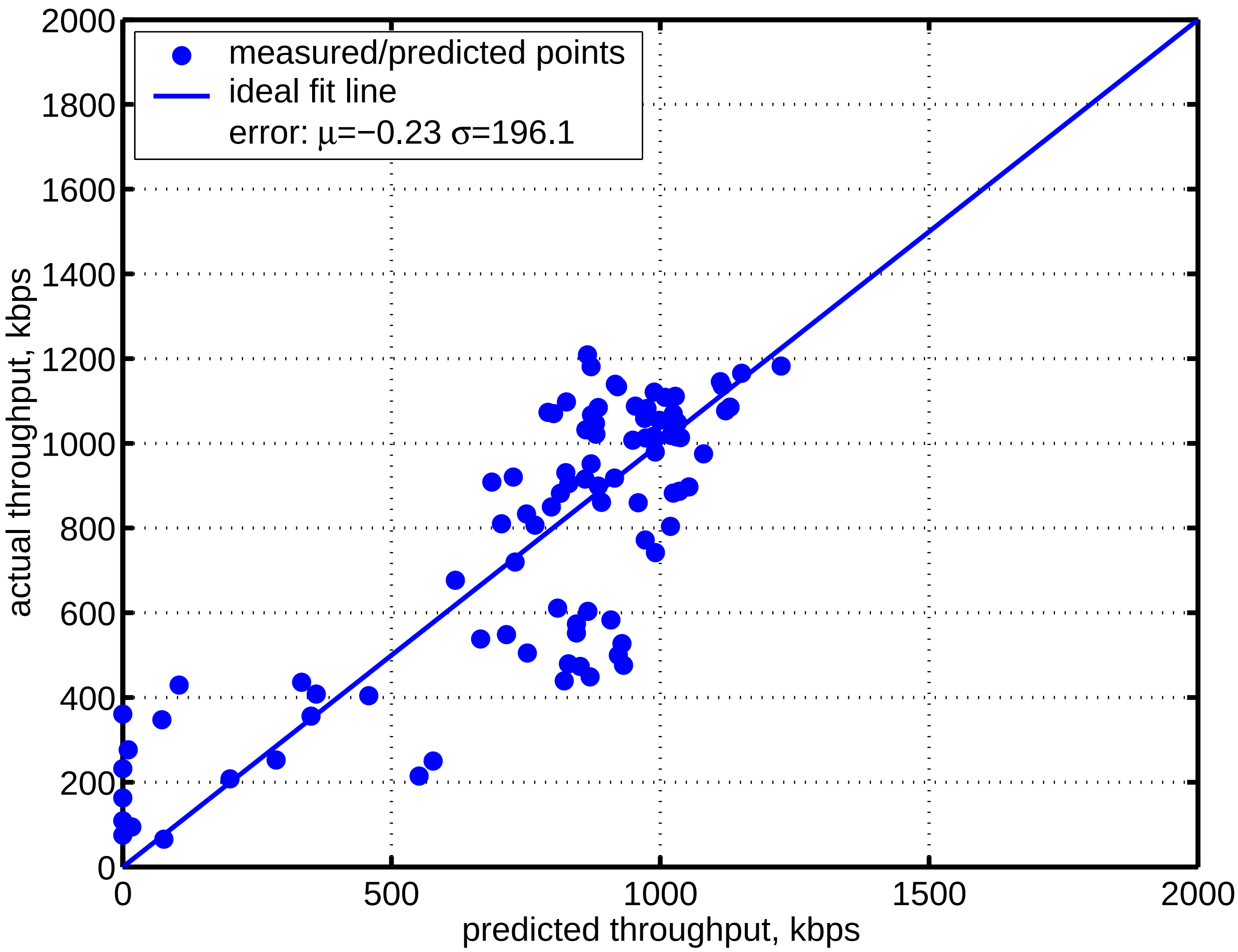

Figure 6.18: 3Com Scenario D, 2 User with Interference, throughput prediction model. This figure shows the predicted versus actual throughput results and an ideal fit line

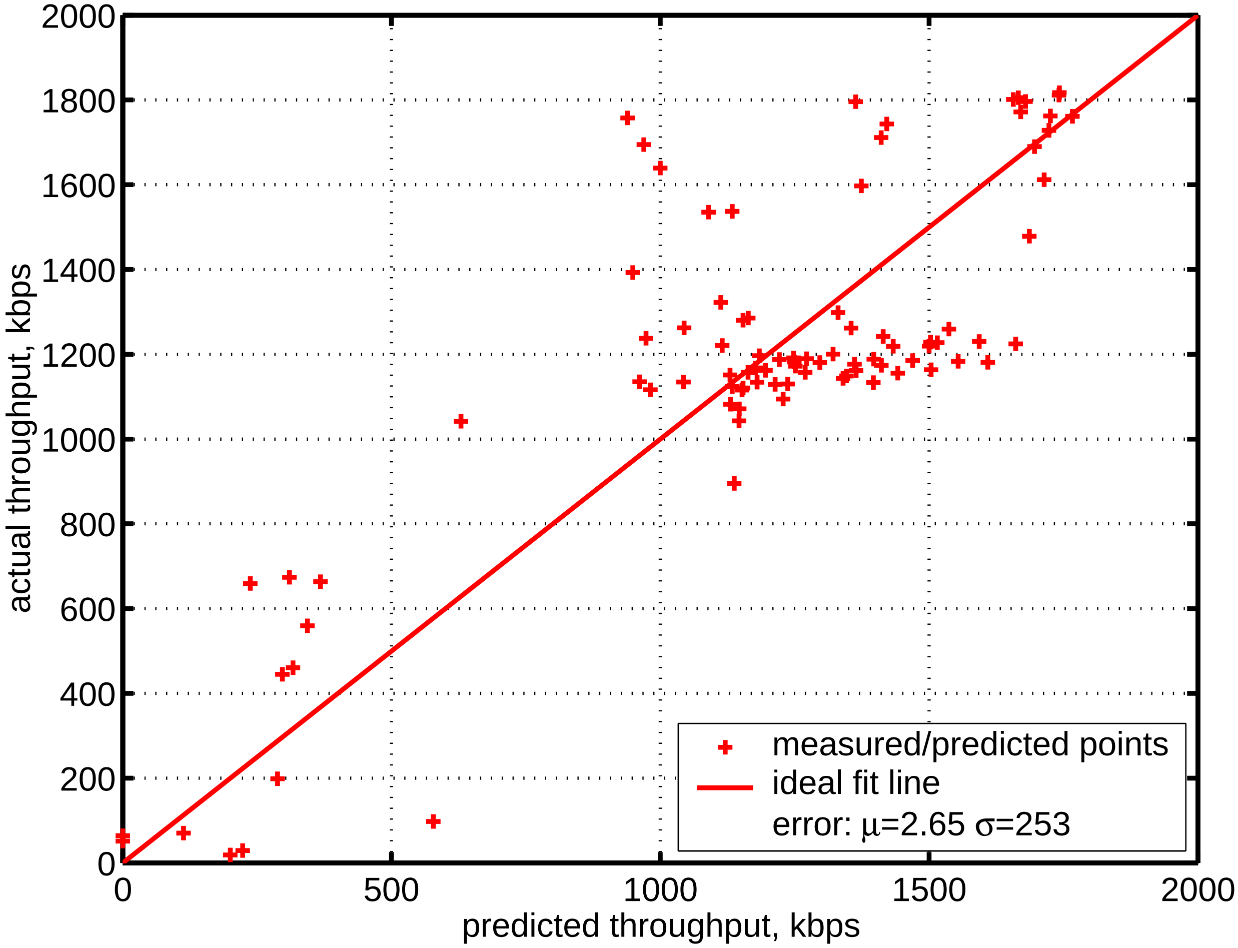

Figure 6.19: WaveLAN Scenario D, Constrained 2 User with Interference, throughput prediction Model. This figure shows the predicted versus actual throughput results and an ideal fit line

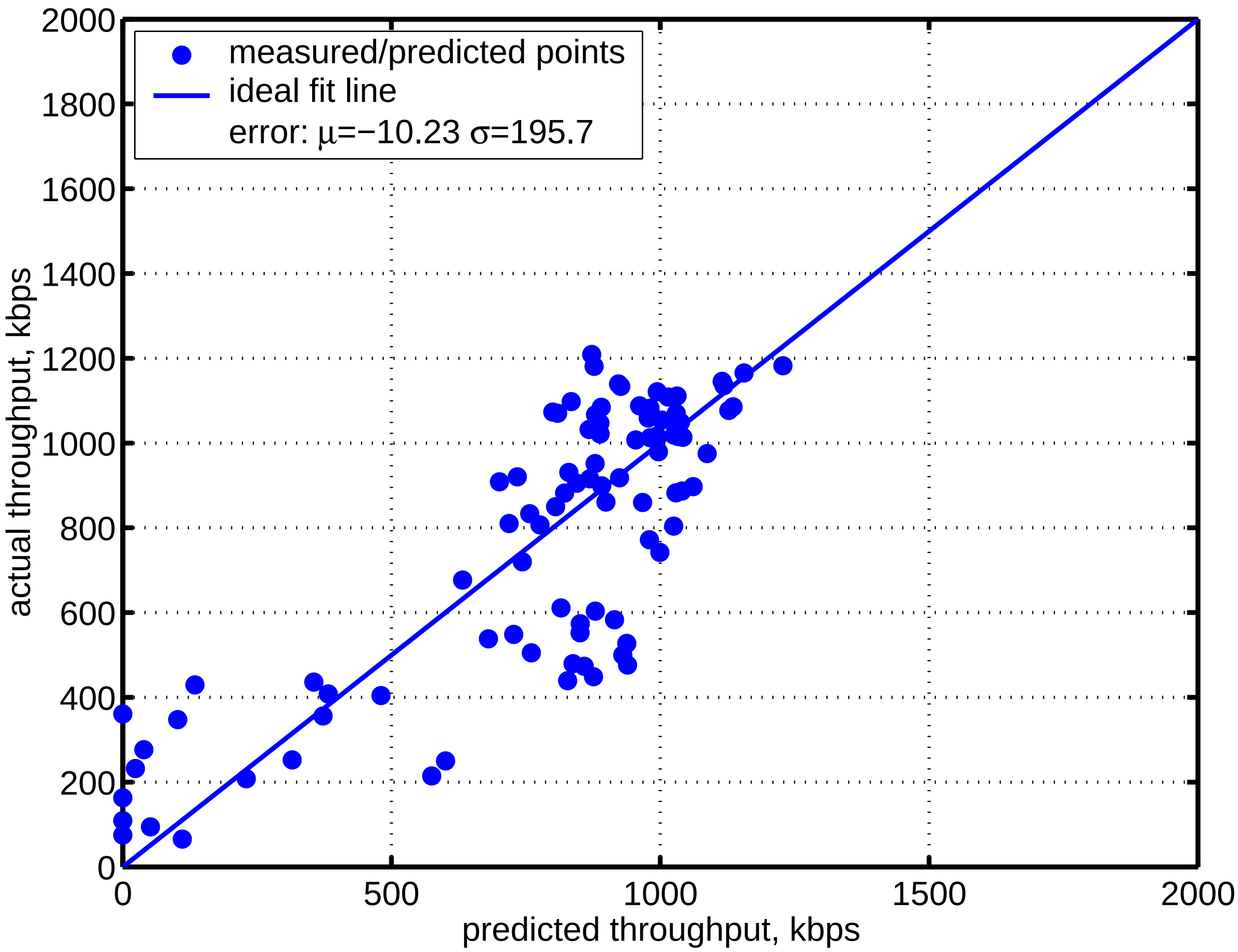

Figure 6.20: 3Com Scenario D, Constrained 2 User with Interference, throughput prediction model. This figure shows the predicted versus actual throughput results and an ideal fit line

6.7 Procedure for Optimizing Models in a Computer Simulation Tool

In order to implement the above models in a computer simulation tool capable of making site-specific measurements in a full three dimensional model of an environment, the following steps would be followed:

1. Model the environment in which the wireless LAN will operate. Be sure to include the locations of access points.
2. Measure the throughput of the wireless LAN at a wide range of locations and with several different numbers of users. Measurements should be taken for several directions at each location.
3. Either measure or predict the RSSI at each measurement location. Ideally this step would be performed simultaneously with the previous step.
4. Perform spatial averaging of the all measurements
5. If RSSI measurements were made in step 3, optimize a throughput prediction model. Equation 5.1 can be used for this step.
6. Use the single client model and Equation 6.3 to calculate ideal values for T_{max}, α and SNR_0.
7. In step 2, if throughput measurements were made for multiple users and/or in the presence of other DSSS wireless LANs, optimize the parameters of Equations 6.8 and/or 6.4 respectively.
8. Use the models with the calculated parameters from steps 6 and 7 to predict the performance of an access point placed in a different location.
9. Repeat the previous step until the desired performance is found for a certain access point location.
10. Repeat steps 8 through 9 as often as needed for multiple access points.
11. Install the access points in the final locations from step 9.

Note that some flexibility in the above steps could easily be followed. However, the order given is intended for a more logical description.

6.8 Summary of Prediction Models

This chapter has presented several, easily simulated prediction models. Reviewing Table 6.2, the mean and standard deviations are all at levels that are quite reasonable for prediction purposes. As a result, it can be safely stated that the models presented here can provide a major help to non-wireless network engineers who do not comprehend wireless specifications, but do understand network statistics.

Chapter 7

Predicting Throughput from a Building Model

7.1 Motivation for Predicting Throughput from a Building Model

As a final verification of the entire concept of this thesis, throughput is predicted directly from a model of the wireless environment. This procedure reflects how the models in this thesis are intended to be used. Additionally, carrying out a throughput prediction from the SitePlanner model of the environment will provide some verification of the measurement, modeling and predictions carried out throughout this thesis.

7.2 Prediction Technique

The prediction of throughput at a certain location in a modeled building environment is carried out in two steps. First, the signal strength at a given location is predicted. Next, the appropriate Scenario Throughput Prediction model is selected. The predicted signal strength is then used as an input to the throughput prediction model. Note that it should be expected that the mean and standard deviations of error should increase somewhat using this technique. This is because errors in throughput prediction are compounded with the errors in the RSSI prediction model. However, the error statistics should be similar for this technique as for the error statistics for the throughput prediction models based on the measured signal strength.

7.2.1 Signal Strength Prediction

Equation 5.1 is used in SitePlanner to predict the signal strength at each of the original measurement locations. These locations and the SitePlanner model of the environment are shown in Figure 5.25. The model parameters shown in Table 5.1 were used for these signal strength positions because they were found to be optimal values for predicting the path loss of wireless LAN signals in Section 5.7.1. This was done using Equation 5.1, but other prediction models were available for use in SitePlanner [26]. The signal strength has been predicted for each average location using omnidirectional antennas. Thus, the predicted signal strength values will be used with the locationally averaged throughput data.

7.2.2 Throughput Prediction from Predicted Signal Strength

The next step to predicting the throughput is to take the predicted signal strength at each location, and use it as an input to the throughput models for each of the four scenarios. Then, by comparing this predicted throughput to the measured throughput, error statistics can be calculated using the optimized models given in equations 6.3, 6.4, 6.8 and 6.10. Figures 7.1, 7.2, 7.3, 7.4, 7.5, 7.6, 7.7, 7.8 illustrate the results of this prediction process. Additionally, Table 7.1 presents a comparison of using the throughput prediction models with either a model of a building environment or using RSSI measurements.

7.3 Successful Prediction of Throughput

From Table 7.1 and the figures presented in this chapter, it is clear that the use of empirical models to predict site-specific, location based throughput is definitely feasible. As the results indicate, a non-technical networking professional could quite reasonably use the prediction models developed in this thesis to design wireless LAN networks.

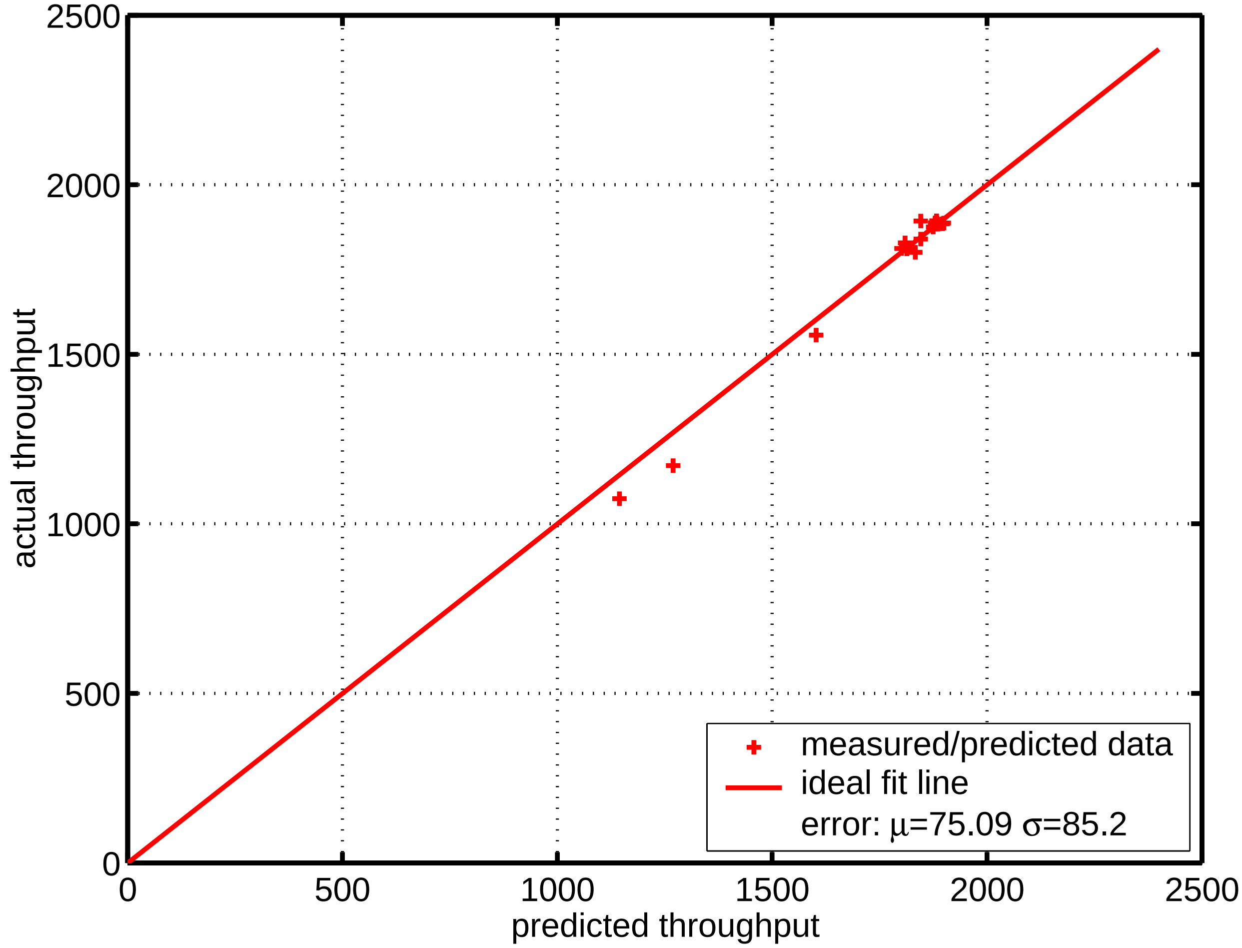

Figure 7.1: Comparison of the fit for Scenario A when throughput is predicted directly from a model of the building environment for the WaveLAN network.

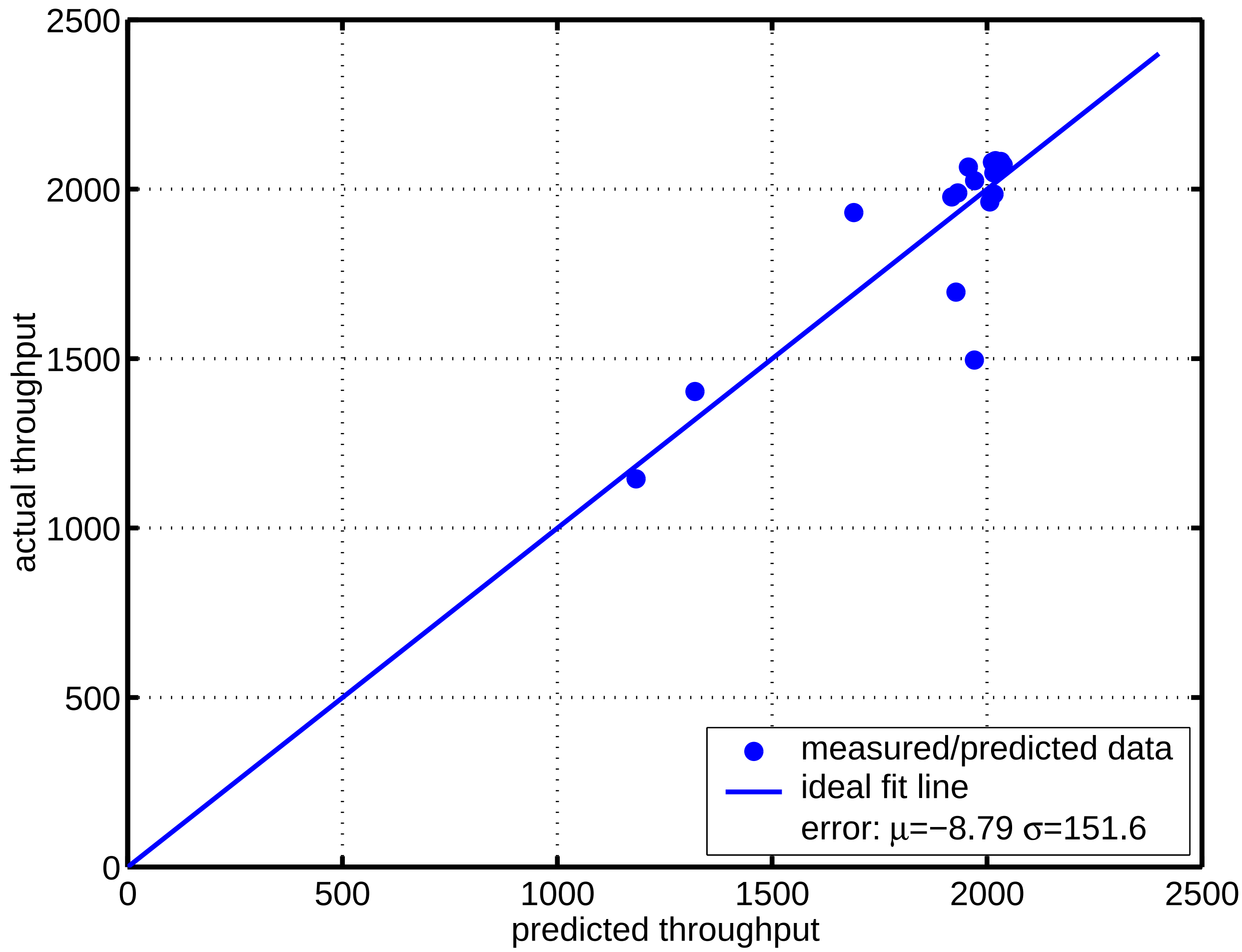

Figure 7.2: Comparison of the fit for Scenario A when throughput is predicted directly from a model of the building environment for the 3Com network.

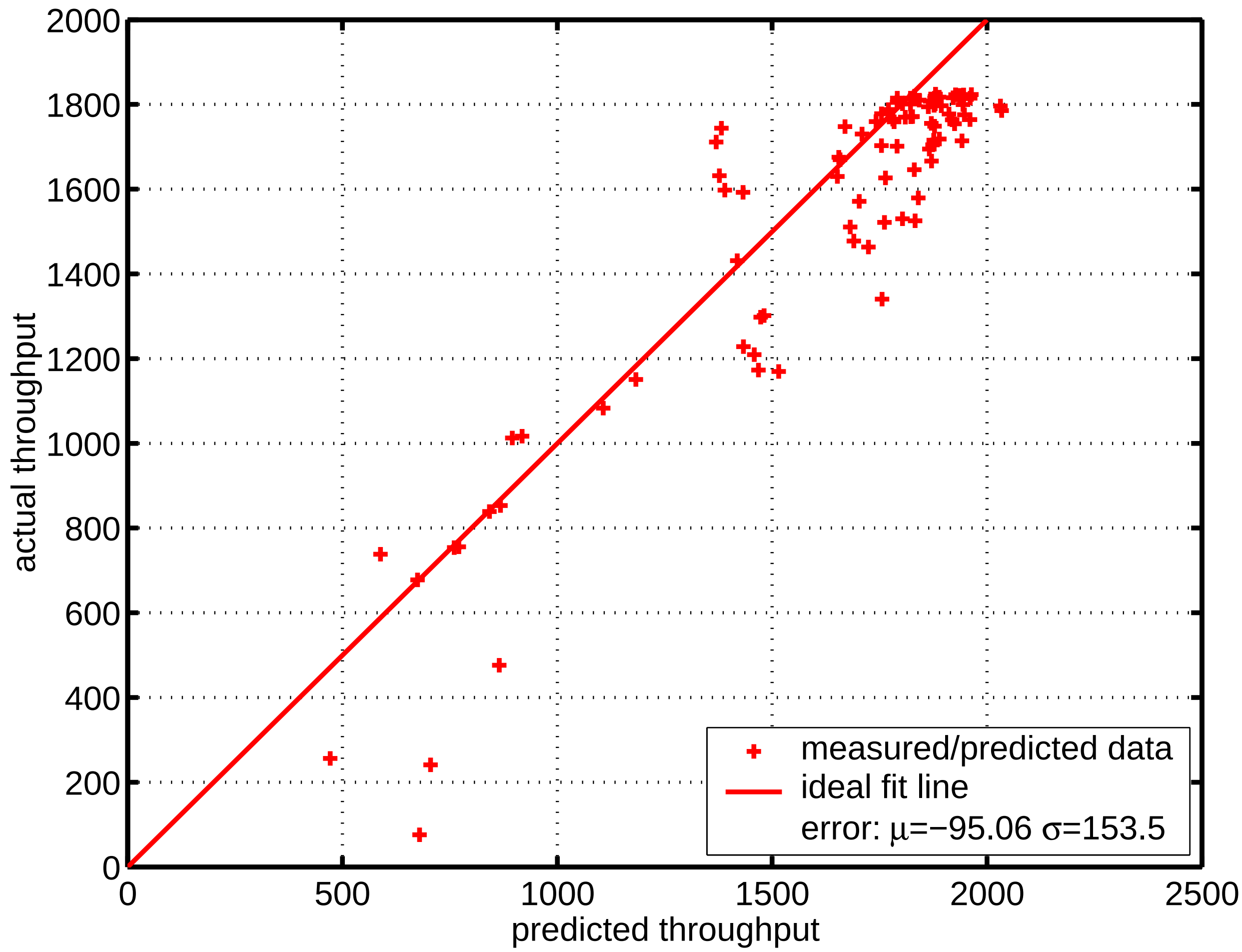

Figure 7.3: Comparison of the fit for Scenario B when throughput is predicted directly from a model of the building environment for the WaveLAN network.

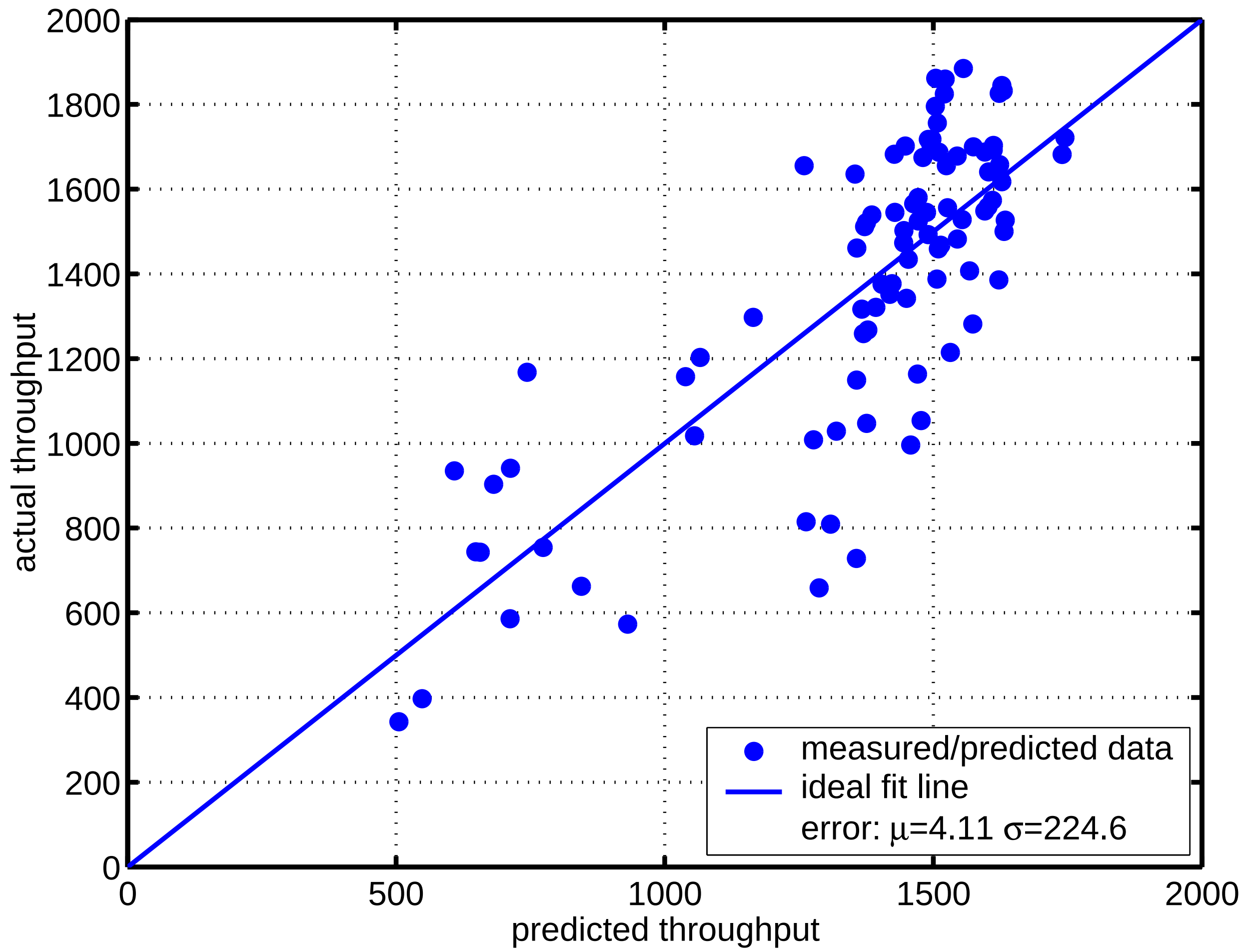

Figure 7.4: Comparison of the fit for Scenario B when throughput is predicted directly from a model of the building environment for the 3Com network.

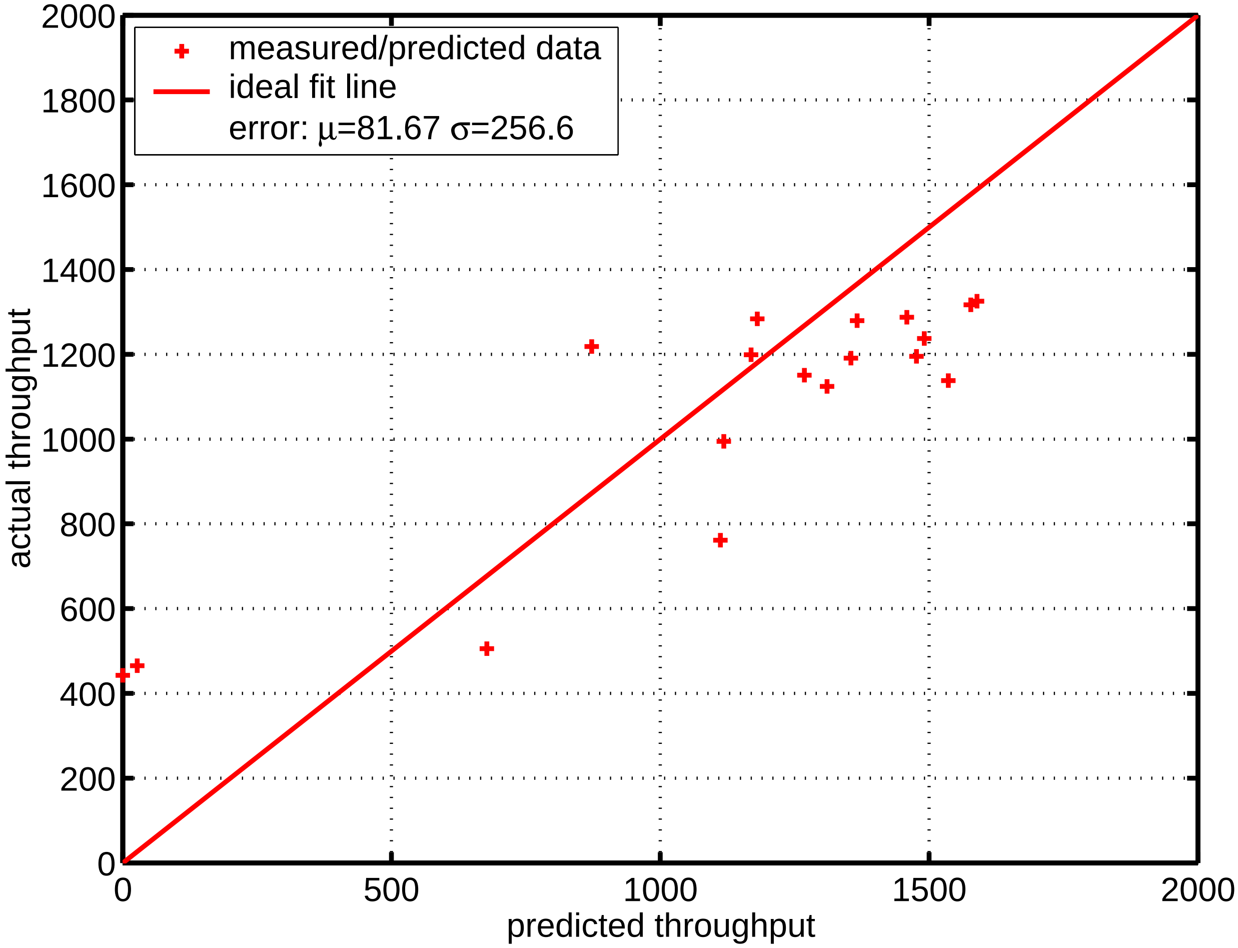

Figure 7.5: Comparison of the fit for Scenario C when throughput is predicted directly from a model of the building environment for the WaveLAN network.

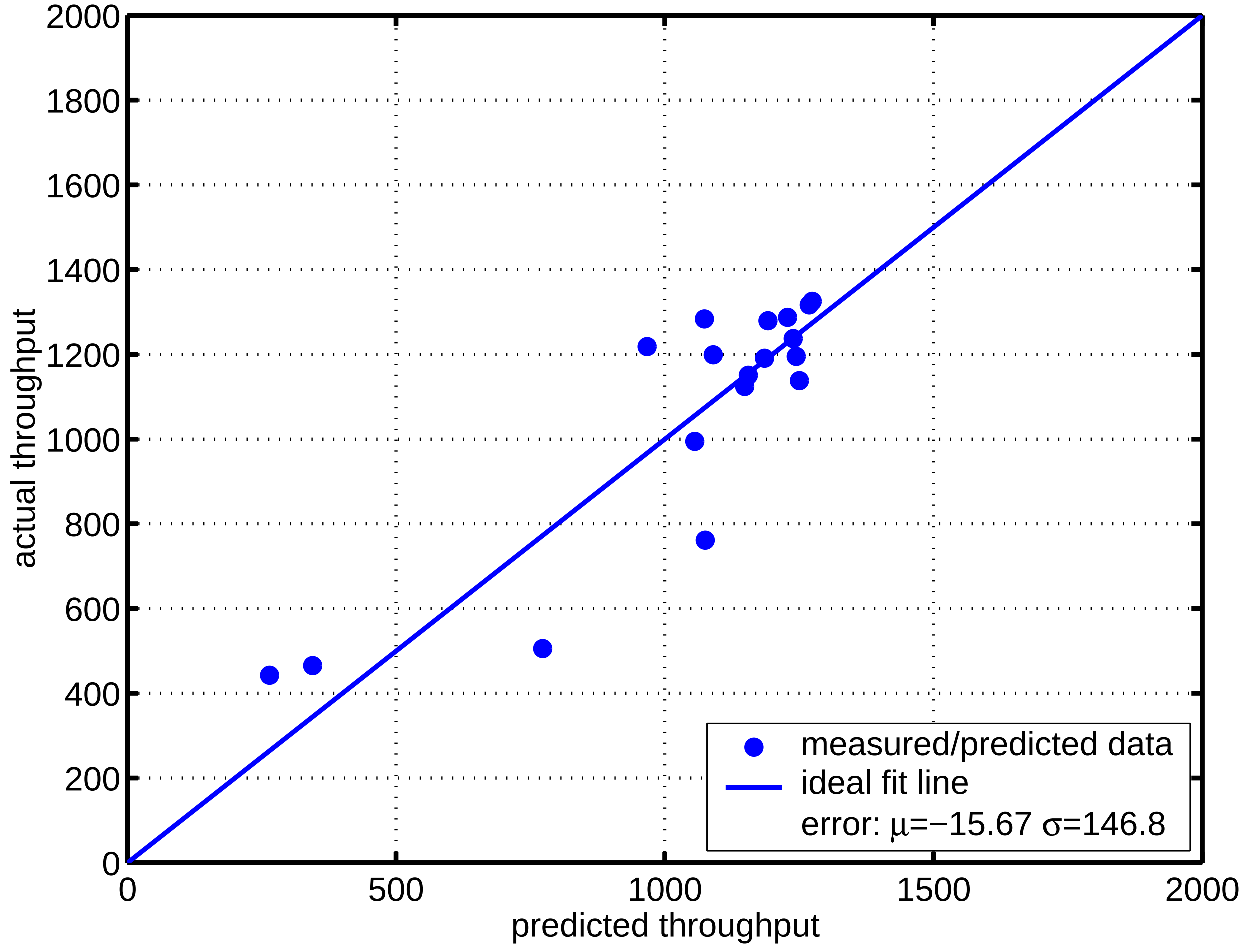

Figure 7.6: Comparison of the fit for Scenario C when throughput is predicted directly from a model of the building environment for the 3Com network.

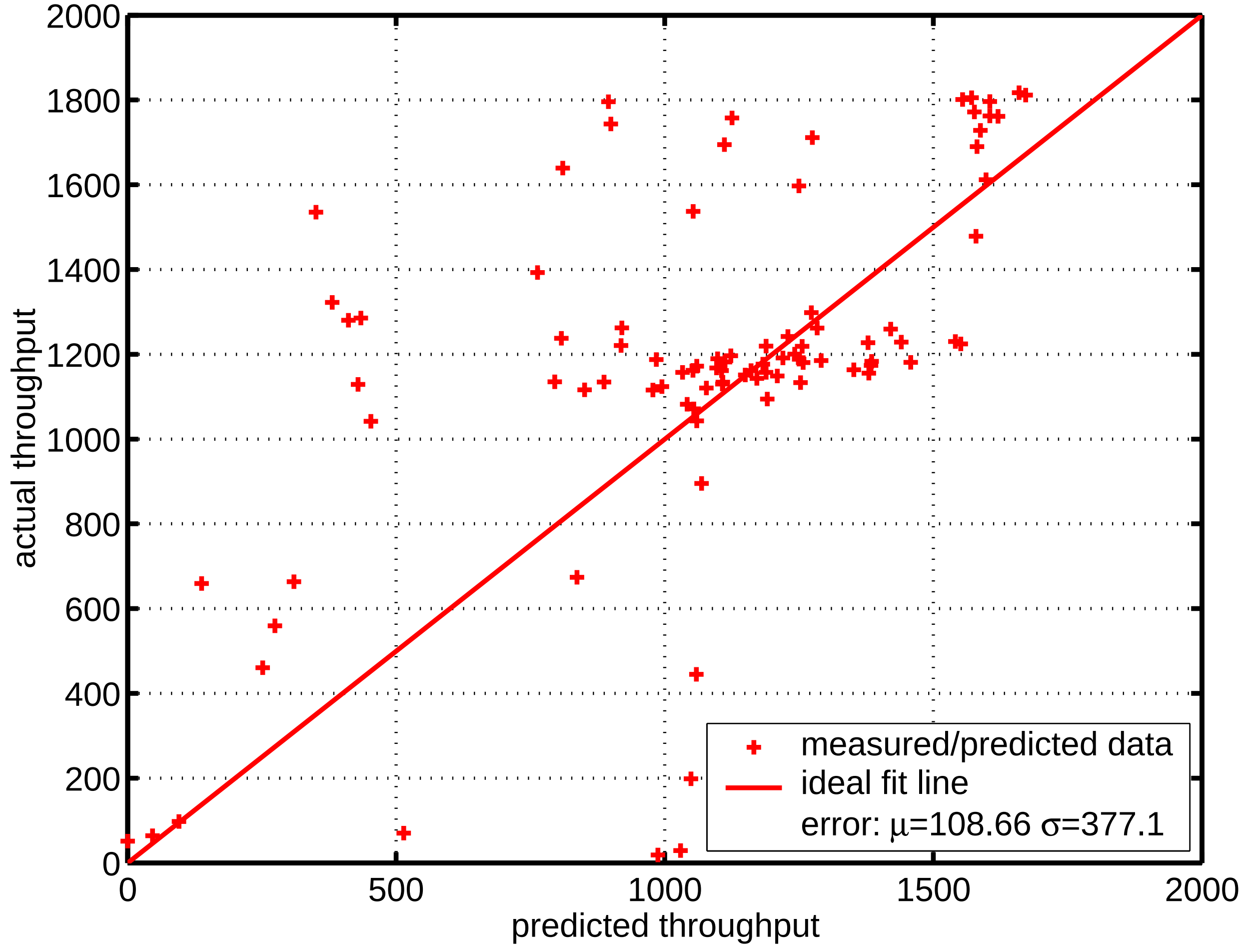

Figure 7.7: Comparison of the fit for Scenario D when throughput is predicted directly from a model of the building environment for the WaveLAN network.

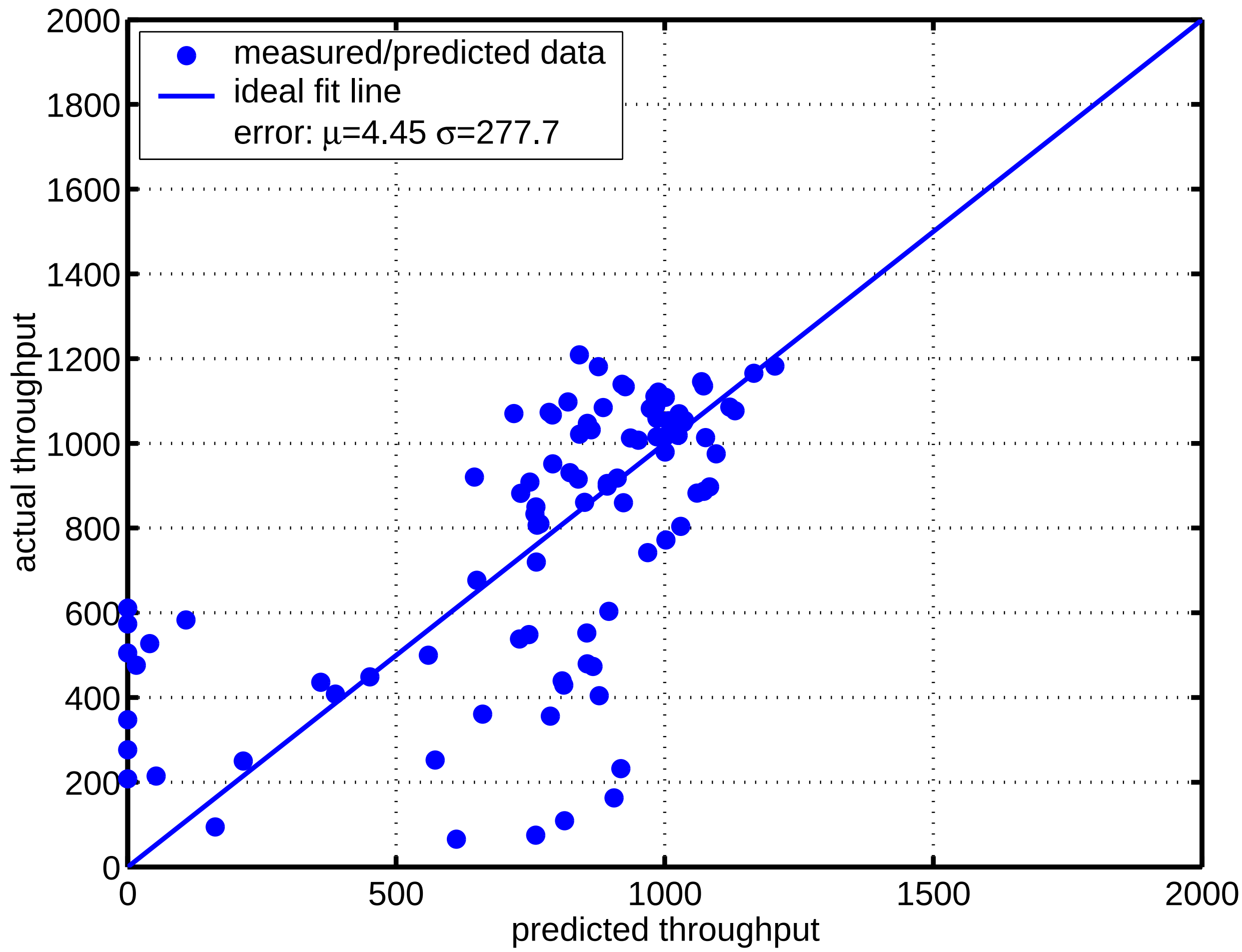

Figure 7.8: Comparison of the fit for Scenario D when throughput is predicted directly from a model of the building environment for the 3Com network.

Table 7.1: Comparison of Throughput Model Error Statistics for Starting from a Building Model or from RSSI Measurements

Model	Using Building Environment		Using RSSI Measurements	
	Mean Error	Error Standard Deviation	Mean Error	Error Standard Deviation
WaveLAN Scenario A	75.09 kbps	85.2 kbps	0 kbps	34.6 kbps
3Com Scenario A	-8.79 kbps	151.6 kbps	0 kbps	147.3 kbps
WaveLAN Scenario B	81.67 kbps	256.6 kbps	0 kbps	185.2 kbps
3Com Scenario B	-15.67 kbps	146.8 kbps	0 kbps	225.2 kbps
WaveLAN Scenario C	95.06 kbps	153.5 kbps	0 kbps	254.9 kbps
3Com Scenario C	4.11 kbps	224.6 kbps	0 kbps	136.2 kbps
WaveLAN Scenario D	108.66 kbps	377.1 kbps	0 kbps	268.7 kbps
3Com Scenario D	4.45 kbps	277.7 kbps	0.2 kbps	196.1 kbps

Chapter 8

Conclusions

8.1 Summary of Findings

This thesis has demonstrated that the network performance of a wireless LAN will vary with the site-specific location of a user, the number of users sharing the wireless LAN and the strength of interfering wireless LAN signals. How the performance of a wireless LAN varies is of critical importance to the professional designing a wireless LAN network. Simulation tools exist today which will allow anyone to predict the signal strength of a wireless LAN access point at any location, but this information is useless to most wireless LAN installers. This is because wireless LAN installers are typically networking professionals with little or no knowledge in regards to wireless packet networks. For this reason, this thesis has developed throughput prediction models for IEEE 802.11b wireless LANs. These models are designed to be intuitive and easy to implement in a simulation product, but still be accurate. This thesis has achieved this goal, as is outlined below.

8.1.1 Wireless LAN Measurement Software

This thesis has presented new wireless LAN measurement software products LANFielder and SiteSpy which allow consistent, non-optimistic measurements of the performance of wireless LAN networks. In addition, this patent pending software utilizes a three dimensional model of the measurement environment which allows a user to actually record the throughput in the modeled location of the precise physical location. The ability of LAN-Fielder to store throughput measurements site-specifically is key to being able to predict and adequately measure the performance of wireless LANs. The software is capable of measuring RSSI, throughput, packet error rates, and packet latencies on a location basis. As a result, the wireless LAN software developed as part of the thesis is capable of recording wireless LAN performance data in a site-specific manner. This is an important contribution to the wireless industry because it allows easy measurement of wireless LAN installations.

8.1.2 Wireless LAN Measurement Campaign

An extensive wireless LAN measurement campaign has been conducted as part of this thesis. The measurements have been extensively presented to provide installers with a concept of how a wireless LAN will perform in a variety of environments and usage scenarios. The measurements were conducted using the software presented in this thesis to allow validation of the software. The wireless data was measured in a variety of different environments and for two different IEEE 802.11b test-bed systems. Additionally, all measurements were performed for a variety of different typical usage scenarios. These scenarios were comprised of one and two user situations each with and without the presence of an interfering IEEE 802.11 DSSS system.

The measurements presented show definite trends throughout the different scenarios. The thesis compared the throughput with the signal to noise ratio at each measurement location. In general this comparison revealed a constant throughput above a certain signal to noise ratio (SNR). Below this certain, critical SNR value, the measured throughput shows a steady decrease. Additionally, the measurement data has shown evidence in competition between multiple users causing a decrease in the measured throughput of wireless LAN users. The presence of interference also tended to cause the wireless LAN systems to drop to lower data transmission rates. These observations were used to to develop empirically based throughput prediction models.

8.1.3 Throughput Prediction Models

The throughput prediction models presented in this thesis have proven to be both effective at predicting the throughput based on just a model of a building environment or based upon RSSI measurements using wireless LAN card hardware. In the development of these models, linear, piece-wise fits to the data was first considered. These models worked quite well, but were considered to be too complicated. Thus, exponential fits to data were instead used. These models proved to be extremely close in terms of accuracy, but were simpler to specify.

To develop the throughput prediction models, the basic exponential model was first fit to the single user measurements. The resulting model was then used as the basis for further models for the two user, one user with interference and two user with interference models. To do this throughput loss factors were introduced. These loss factors were based on the expected causes for a loss in throughput: competition for shared resources, hidden terminal difficulties, and the presence of strong interference resulting in transmission delays. Using these loss factors, a model for each of the measurement scenarios was created. These models were then fit to the measured data.

The accuracy of these models has been shown to be as good as zero mean error and an error standard deviation of 150 kbps, or as poor as a mean error of 26 kbps and a error standard deviation of 270 kbps, depending on the usage scenario. The most models, however, had zero mean error and error standard deviation of less than 200 kbps. These results imply the developed models fit the measurements quite well. Since the models are also easily implemented in a computer-aided simulation environment, the models are an important contribution to wireless LAN system design.

8.2 Further Research

This work has opened up several areas of possible new research. For instance, the measurements and models presented in this thesis only cover some of the most basic usage patterns of wireless LANs. Additional measurements of wireless LANs in operation using the PCF modes, or having RTS/CTS or even WEP functionality enabled would be quite interesting to compare to the results presented in this thesis. Additionally, it would be useful to explore scenarios in which FHSS or high rate (i.e. IEEE 802.11b) DSSS wireless LANs or even Bluetooth or HomeRF modems were used as interference sources. Lastly, it would be interesting to measure scenarios in which a large number of users were all accessing the wireless LAN using different server computers and different usage levels.

It should be noted that the data collected seems to indicate that the developed throughput prediction models seem to fair better for the WaveLAN equipment in the single client cases but better for the 3Com equipment in the interference free cases. This indicates that perhaps different models for different manufacturers would be potential area of research. This was however avoided in this thesis because a single, intuitive model for throughput prediction was desired. Regardless, the models presented in this thesis provide a useful groundwork for more advanced wireless LAN modeling.

Appendix A

An Overview of the IEEE 802.11 and IEEE 802.11b Standards

A.1 Organization of the Standard

The IEEE 802.11b standard is broken into two main layers: the MAC or Media Access Control layer and the PHY or Physical Layer. These two layers allow a functional separation of the standard and, more importantly allows a single data protocol to be used with several different RF transmission techniques. Since the goal of this thesis is to predict the performance of IEEE 802.11b wireless LAN products, this chapter will present an overview of the DSSS function of the PHY layer and a basic description of the MAC layer.

A.2 Physical Layer

The PHY layer of the 802.11 standard defines the different RF transmission techniques. There are three basic transmission techniques: Frequency Hopping Spread Spectrum or FHSS, Direct Sequence Spread Spectrum or DSSS, and Diffuse Infrared. The relationship of these three standards is shown in figure A.1. Note that the diffuse infrared PHY access technique has received little attention and will be neglected in this thesis as it is not relevant to the research which has been conducted.

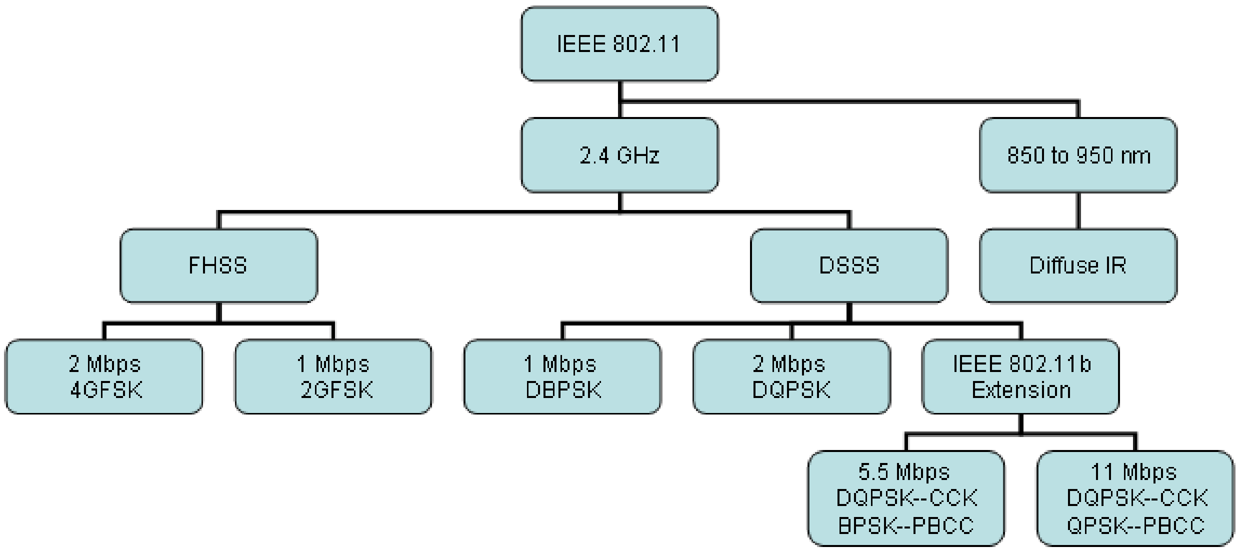

Figure A.1: A diagrammatic overview of the IEEE 802.11 standard as defined in [30] and [31].

A.2.1 DSSS Frequency Band and Channels

The two remaining PHY access techniques operate in in channels spread between 2.4 and 2.497 GHz. However, the FHSS technique has also been substantially less popular than the DSSS technique. This is mainly due to the higher bandwidth available to the DSSS implementation and the fact that the DSSS function lends itself better to interoperability between different implementations. For this reason, this thesis focuses on the more popular technique of DSSS used under the IEEE 802.11b standard. The specific channels available vary by country and the regulation agencies which controls the spectrum allocation. An description of the spectrum available in different countries is shown in Table A.1. The DSSS channels defined by the IEEE 802.11 standard for different countries are also shown in Figure A.2.

A.2.2 DSSS Advantages

The DSSS technique has two major advantages. It provides a spreading gain against narrowband interference signals and it spreads the transmitted signal across a wide range so the transmission resembles noise to a narrowband receiver. These two characteristics

Table A.1: World Wide Spectrum Allocation for IEEE 802.11 and 802.11b use as defined in [30] and [29].

Country	Regulatory Agency	Frequency Range Available	DSSS Channels Available	FHSS Channels Available
United States	FCC	2.4 to 2.4835 GHz	1 through 11	2 through 80
Canada	IC	2.4 to 2.4835 GHz	1 through 11	2 through 80
Japan	MKK	2.4 to 2.497 GHz	1 through 14	2 through 95
France		2.4465 to 2.4835 GHz	10 through 13	48 through 82
Spain		2.445 to 2.475 GHz	10 and 11	47 through 73
Remainder of Europe	ETSI	2.4 to 2.4835 GHz	1 through 13	2 through 80

are why DSSS was originally used by the military because it is difficult to jam and difficult to detect by narrowband radios. These two characteristics also make the DSSS technique ideal for coexisting with other narrowband users.

A.2.3 Data Bandwidths and Modulation Techniques

DSSS is currently a very popular transmission technique because it has the highest data bandwidth available. The DSSS transmission technique is defined in both the IEEE 802.11 and IEEE 802.11b standard. The 802.11b standard was introduced after the 802.11 standard to define 2 different modulation techniques in addition to the two originally defined in the original 802.11 standard. The original 802.11 standard originally defined two data bandwidths (e.g. bit rates): 2 Mbps and 1 Mbps. These bandwidths use a Barker sequence described below. However, the data is modulated using DQPSK and DBPSK for the 2 Mbps and 1 Mbps bandwidths respectively. The IEEE 802.11b standard adds 5.5 Mbps and 11 Mbps data bandwidths to the original 2 Mbps and 1Mbps transmission techniques.

Switching Between Data Bandwidths

A compliant IEEE 802.11b product can use any of the four transmission bandwidths. Ordinarily a wireless LAN card will operate at the highest possible bandwidth. However, as a user moves further away from an access point or if an interference source is present, the highest bandwidth may not provide reliable transmission of data. To combat this, the wireless LAN card will drop to a lower data transmission bandwidth. The lower rates are

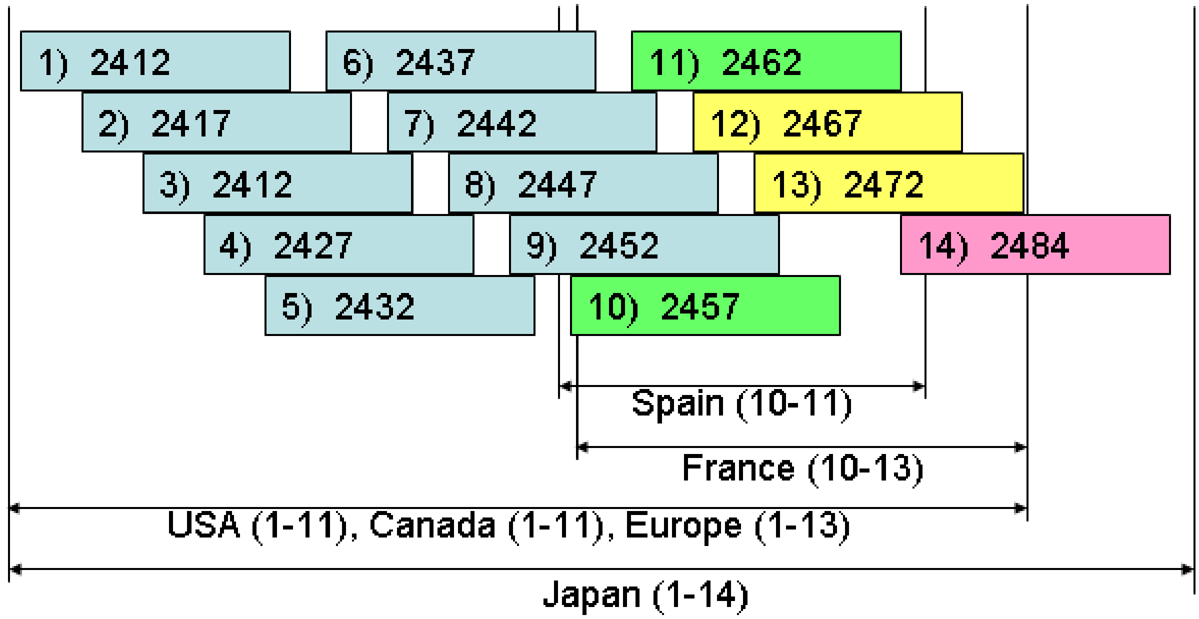

Figure A.2: Illustration to scale of IEEE 802.11 DSSS Channels as defined in [30]. The channels are labeled by channel number and center frequency, in MHz. Note that channels 12 and 13 are not used in the USA or Canada even though they lie within the allocated frequency bands. Note also that channel 10 is defined for use in France even though it slightly exceeds the valid frequency band available in France. Lastly, also note that channel 14 is slightly offset from the other channels which all have 5 MHz spacings. All Channels have 22 MHz Bandwidths.

more tolerant to noise and thus can be more reliable than the faster transmission rates. The IEEE 802.11 standard does not define what criteria to use to decide which data transmission rate to use. The standard only requires that all compliant products support all data bandwidths for compatibility purposes. The next sections explain how the different data bandwidths are achieved.

A.2.4 1 and 2 Mbps Transmission Rates

Direct Sequence Spread Spectrum uses a PN spreading code to spread transmitted data over a wide bandwidth. This can be thought of as XORing a stream of data bits with a specific PN sequence. In the IEEE 802.11 standard, a single PN code is used by every user in the network. (See Section A.3.3 for information about multiple access techniques). This

PN code is the 11 bit barker sequence: +1 -1 +1 +1 -1 +1 +1 +1 -1 -1 -1. The technique of XORing data with the Barker sequence is shown in Figure A.3.

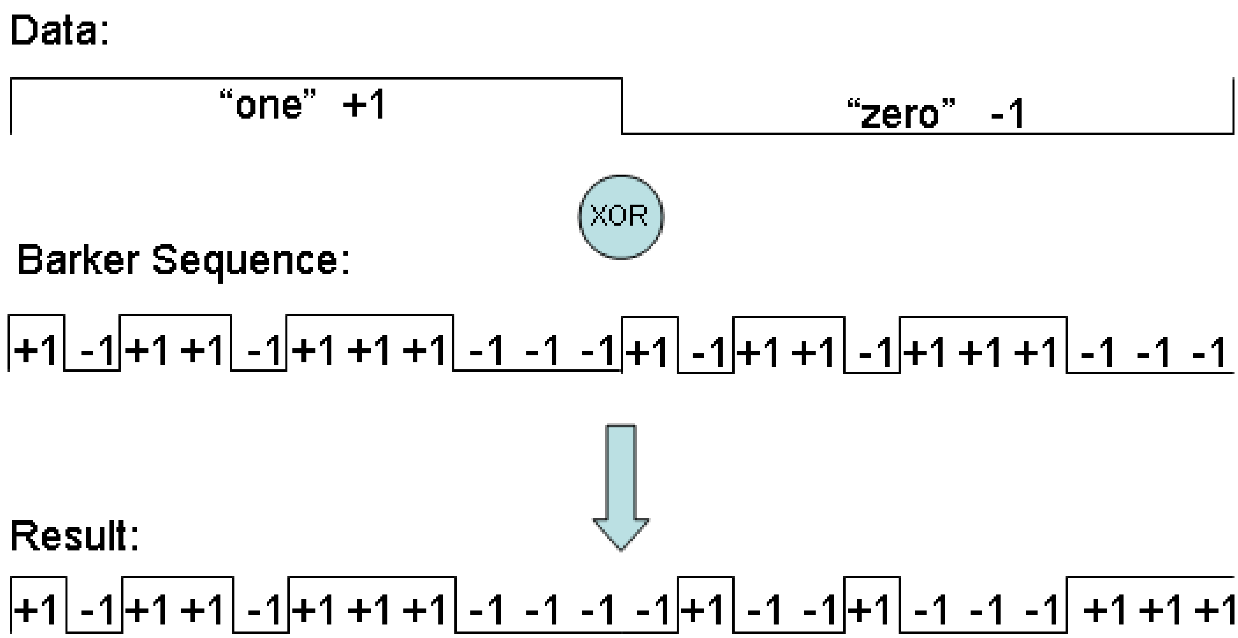

Figure A.3: Illustration of the process of spreading data using a Barker Sequence. Adapted from [36].

The figure shows how a "one" or a "zero" is transmitted as 11 bits of data represented by the original Barker sequence or the inverse of the Barker sequence. By combining data with a bandwidth at baseband of 1 MHz, spreading it and up-converting it to the desired 2.4 GHz channel results in an RF channel bandwidth of 22 MHz, as is shown in Figure A.4. The data is transmitted using either BPSK or QPSK encoding to provide the 1 Mbps and 2 Mbps transmission rates defined in the IEEE 802.11 standard.

A.2.5 5.5 Mbps and 11 Mbps Transmission Rates

The IEEE 802.11b standard defines two additional data bandwidths of 5.5 and 11 Mbps respectively. For these two transmission rates, the data is no longer spread using the Barker Sequence defined above. In order to increase the data rate to 5.5 and 11 Mbps, the IEEE 802.11b standard uses Complimentary Code Keying (CCK) or, optionally, Packet Binary Convolutional Coding (PBCC). Both techniques are discussed in the next two sections.

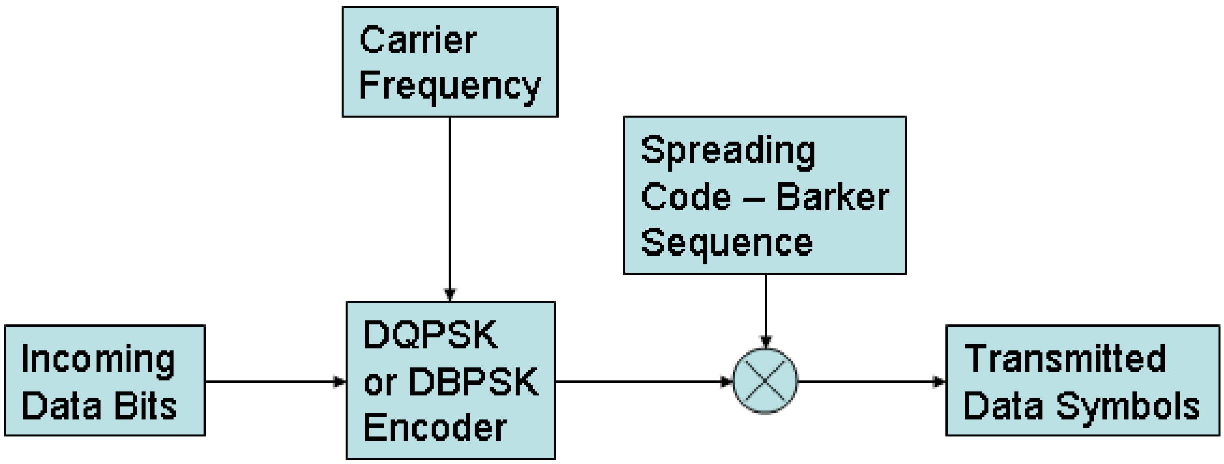

Figure A.4: Illustration of the basic IEEE 802.11 Spreading technique. Created using information from [30] and [31].

A.2.6 CCK Encoding

The IEEE 802.11b standard requires modems to support a technique known as Complementary Code Keying to simultaneous spread data across a 22 MHz channel while transmitting more data bits per 11 spread bits than the 1 or 2 bits transmitted in the plain IEEE 802.11 standard.

CCK Based 11 Mbps Transmission Technique

The 11 Mbps version of CCK encoding works using an 8 chip spreading sequence. This can be done because the 8 bit sequence still runs at a rate of 11 Megachips per second, which results in a spreading factor of 11. CCK encoding, however, does not use a static PN code. It calculates a different spreading code based on the incoming data. This is done by breaking the incoming bits into symbols of 8 bits in duration. An 8 chip spreading code is found from the 8 data bits. Each chip is then encoded using the same DQPSK constellation and then transmitted.

The 11 Mbps CCK technique of calculating the CCK spreading code can be broken down into two steps. First, 8 data bits are split into pairs called dibits. The four dibits are first used to calculate four phase angles, φ_1 through φ_4. The 2nd through 4th dibits are converted to a phase angle, φ using the mapping shown in Table A.2. The first dibit is converted to φ_1 using Equation A.1. That is, $\varphi_1(i)$ found as the value of $\varphi_1(i-1)$ for the previous symbol plus an offset angle found using Table A.2 plus 180 degrees if this is an odd symbol, 0 if it is an even symbol. In Equation A.1, i is the current symbol. At time zero, i starts at zero. This allows the $\pi * mod(i,2)$ portion of the formula to provide a 180 degree (π radian) shift to every odd numbered symbol. The $offset(1^{st} \quad dibit$ portion of the equation is simply a table look-up procedure based on the first dibit, using Table A.2. Correct use of this table and Equation A.1 results in the first of four phases, φ_1.

Table A.2: Mapping of dibits to Angles for CCK Modulation. Created using information from [31].

dibit	φ value or offset angle, radians
00	0
01	$\pi/2$
10	π
11	$3\pi/2$

$$\varphi_1(i) = \varphi_1(i-1) + offset(1^{st} \quad dibit) + \pi * (mod(i,2)) \tag{A.1}$$

In the second step the four phases are used to calculate 8 complex chips using the mapping in Table A.3. The 8 complex chips are then mapped to the same QPSK constellation and transmitted. Note that the QPSK constellation is actually a DQPSK constellation. The value of φ_1 offsets all of the chips by the same angle, which is an offset to the QPSK constellation used for the previous set of 8 complex chips.

CCK Based 5.5 Mbps Transmission Technique

The 5.5 Mbps version of CCK modulation is carried out in the same manner that the 11 Mbps version is carried out, except that only four bits are encoded per symbol instead of eight. Since, there are only four bits per symbol, the φ values are calculated differently. In this case the first dibit is still used to encode φ_1 in exactly the same manner as before using Equation A.1. The remaining φ values are calculated based on the 3rd and 4th bits. The formulas for these calculations are given in Equations A.2, A.3 and A.4. Once the φ values have been calculated, the chips are calculated and mapped to the QPSK

Table A.3: Mapping of Angles into Complex Chips for CCK Modulation. Created using information from [31].

Chip Number	Formula
0	$e^{j(\varphi_1+\varphi_2+\varphi_3+\varphi 4)}$
1	$e^{j(\varphi_1+\varphi_3+\varphi 4)}$
2	$e^{j(\varphi_1+\varphi_2+\varphi 4)}$
3	$-e^{j(\varphi_1+\varphi 4)}$
4	$e^{j(\varphi_1+\varphi_2+\varphi_3)}$
5	$e^{j(\varphi_1+\varphi_3)}$
6	$-e^{j(\varphi_1+\varphi_2)}$
7	$e^{j\varphi_1}$

constellation exactly as was done for the 11 Mbps version of CCK.

$$\varphi_2 = (3^{rd} \quad bit) * \pi + \pi/2 \tag{A.2}$$

$$\varphi_3 = 0 \tag{A.3}$$

$$\varphi_4 = (4^{th} \quad bit) * \pi \tag{A.4}$$

The above process can be thought of in a slightly different manner. Instead of calculating the chips as above, the chips can be thought of as being calculated using the equations shown in Table A.3 modified to not include φ_1. That is, each chip would be a function of just φ_2, φ_3, and φ_4. Then, instead of encoding each chip using QPSK, the entire symbol is encoded using a DQPSK encoding based on the value of φ_1 as a phase difference from the previously transmitted symbol. Then the symbol is spread using the CCK chips calculated using the modified equations. This alternative way of considering the process is show in block diagram form in Figure A.5.

CCK Encoding Example

To illustrate the CCK spreading process, the sequence 00011011 as an odd symbol in which the phase of the previous symbol was pi/2 will be encoded. The resulting dibits are 00, 01, 10 and 11. The phases would first be calculated to be $\varphi_1 = \pi/2$, $\varphi_2 = \pi/2$, $\varphi_3 = \pi$, $\varphi_4 = 3\pi/2$. Next, the 8 complex chips would be calculated to be $e^{3j\pi/2}$, $e^{j\pi}$, $e^{j\pi/2}$, $-e^0 = e^{j\pi}$,

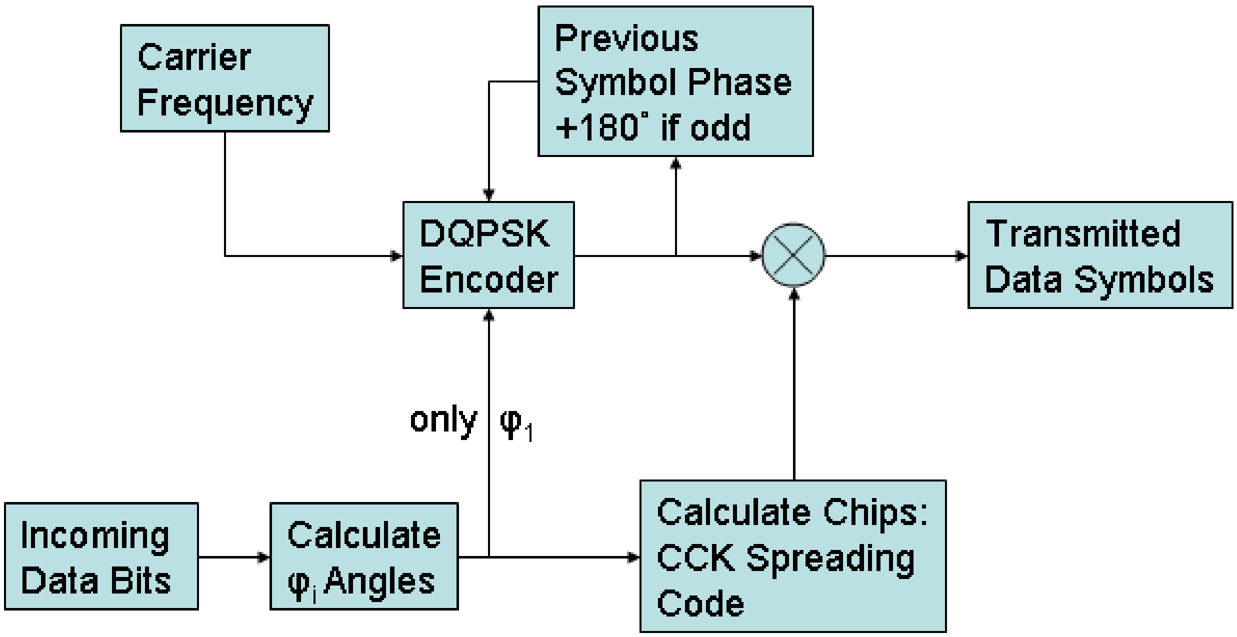

Figure A.5: Conceptual block diagram of the CCK encoding process carried out by the transmitter. Created using information from [31].

e^0, $e^{3j\pi/2}$, $-e^{j\pi} = e^0$, and $e^{j\pi/2}$. Each chip would be encoded on the QPSK constellation based on the phase of each of the chips.

A.2.7 PBCC Encoding for 11 and 5.5 Mbps Transmission Rates

The IEEE 802.11b standard also supports an optional technique known as Packet Binary Convolutional Coding (PBCC) to achieve the 5.5 and 11 Mbps data rates. The PBCC encoding technique uses a standard 1/2 rate, 64 state, rate $\int$ code. The PBCC technique feeds the data bits into the 1/2 rate encoder. The 1/2 rate encoder, by definition, generates 2 output bits for each input data bit. The output of the encoder is mapped to a QPSK constellation for the 11 Mbps data rate and to a BPSK constellation for the 5.5 Mbps data rate. To provide some pseudo-randomness to this technique a pseudo-random cover code is used to vary the QPSK or BPSK constellation used. A block diagram of this technique is shown in Figure A.6

A 256 bit cover sequence is used to vary the QPSK or BPSK constellation used. The cover sequence is generated by taking the 16 bit sequence: 0011 0011 1000 1011 and rotating it 3

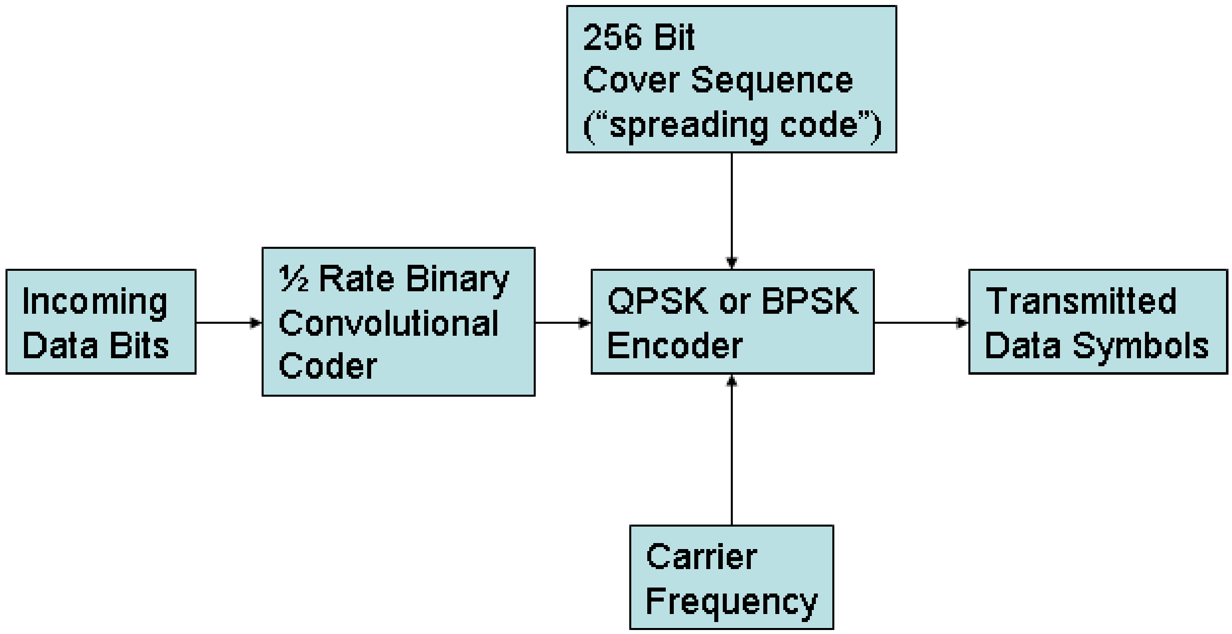

Figure A.6: A block diagram of the PBCC encoding technique. Created using information from [31].

bits to the left 15 times to generate the 256 bit sequence made up of 16 sequences of 16 bits each. Thus the 17th through 32nd bits of the full sequence are 1001 1100 0101 1001. The 33rd through 48th bits are 1110 0010 1100 1100. This 256 bit sequence is used repeatedly to vary the constellation used to transmit each chip. Specifically, if the cover sequence is a 0, one constellation is used. If the cover sequence is a 1, then the constellation is rotated by +pi/2 and used. After the end of the 256 bit sequence, the sequence is repeated.

Note that no ”spreading” takes place in this technique. Rather, the data is encoded directly at the desired data rate (either 11 Mbps or 5.5 Mbps). However, the use of the cover sequence will act to randomly distribute the data transmission across the full 22 MHz channel bandwidth and the ”spreading” can be thought of as occurring inside the QPSK or BPSK encoder.

A.2.8 PHY Layer Packet Format

The IEEE 802.11b standard defines two different packet structures that are used in the DSSS standard. There is a short and a long packet format as shown in figure A.7. The

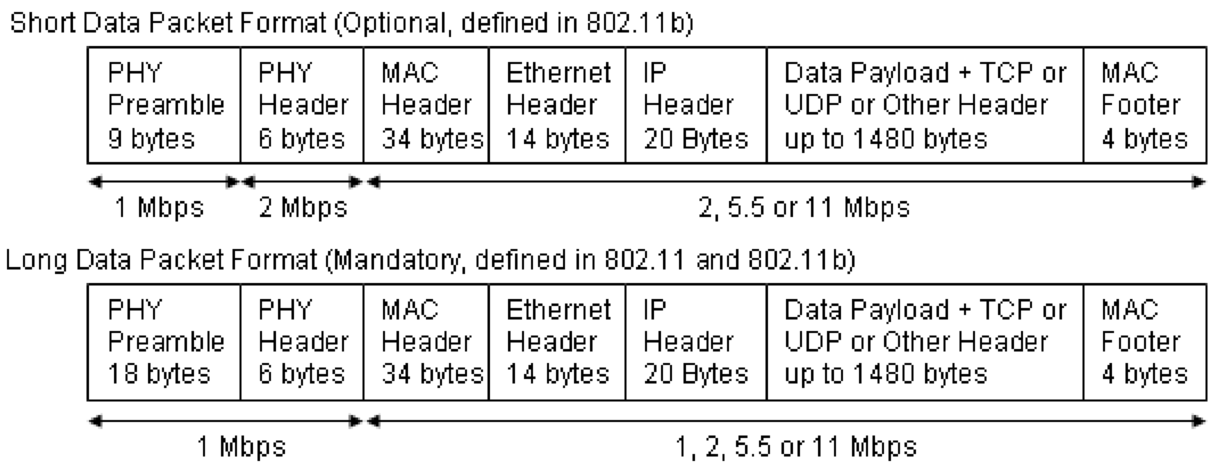

Figure A.7: Basic structure of the IEEE 802.11b packet as it's transmitted on the physical layer is shown here. Adapted from [31].

short packet format is intended to reduce the overhead of transmissions while the long packet format is to maintain compatibility with IEEE 802.11 networks. The PHY preamble is used to allow the receiver to get synchronized to the transmitter. The PHY header is the overhead needed by the PHY layer. The remainder of the packet contains the data passed to the PHY layer by the MAC layer as is shown in A.8. Note that different parts of the packet are transmitted at different transmission rates.

A.3 MAC Layer

A.3.1 Basic Network Layout

The IEEE 802.11 standard defines two types of networks: Adhoc and Infrastructure. Adhoc networks are self-configuring networks between mobile and portable wireless clients. Infrastructure networks use fixed, interconnected access points to provide connectivity to mobile and portable wireless clients. This thesis will focus on infrastructure networks. Since these networks use fixed location access points, it is import to carefully select the locations of these access points. This is the main motivation for this thesis. Infrastructure based wireless networks need to provide some level of service, either in terms of coverage area or in network performance, or in both. However, in order to carefully place the access points which make up a wireless LAN infrastructure, design rules are needed. Prediction models which can be used to design wireless LAN infrastructure networks are presented later in this thesis.

A.3.2 MAC Layer Packet Structure

The basic format of packets passed to the PHY layer from the MAC layer is shown in figure A.8. Note that this is the basic format for all packets sent by the MAC layer. Some actual packets do not actually contain all of the fields. However, all fields are present in all data packets. Up to four addresses are needed because it is sometimes necessary to identify the address of the access point used by the transmitter or receiver. Thus, if two wireless LAN users are sending packets to one another but each is using a different access point, the 802.11 MAC address of both access points and both clients will be present in the four address fields.

Frame Control 2 bytes	Duration and ID 2 bytes	Address 1 6 Bytes	Address 2 6 Bytes	Address 3 6 Bytes	Sequence Control 2 Bytes	Address 4 6 Bytes	Frame Body 0 to 2312 Bytes	Frame Check Sequence 4 Bytes

Figure A.8: The structure of packet created at the MAC Layer. Adapted from [30].

A.3.3 Multiple Access, DCF and CSMA/CA

Regardless of the physical layer used, all IEEE 802.11 wireless LAN clients use the same channel to transmit on. This means the standard needs to define a way in which clients know when they can transmit and when they can not. This is handled using several multiple access mechanisms. The most basic of these is the Carrier Sense Multiple Access with Carrier Avoidance (CSMA/CA) mechanism. This mechanism is defined as part of the Distributed Coordination Function (DCF) of the IEEE 802.11b standard. The DCF is the mandatory method by which clients work together and differ access to the medium so that the all users can use the same wireless channel.

CSMA/CA is based on the multiple access technique used in wired Ethernet connections, Carrier Sense Multiple Access with Collision Detection, CSMA/CD. In both types

Table A.4: Different Interframe Spacings (IFS) as defined in [30].

Abbreviation	Meaning
SIFS	Short Interframe Spacing
PIFS	Point Coordination Function (PCF) Interframe Spacing
DIFS	Distributed Coordination Function (DCF) Interframe Spacing
EIFS	Extended Interframe Spacing

of CSMA users first sense the transmission medium to see if anyone is transmitting just before transmitting a packet of data. This only partially avoids the possibility of packets being transmitted by two users at the same time. When two or more packets are transmitted simultaneously, or overlapping in time, a "collision" is said to have taken place. In wired Ethernet connections a user is able to detect when a collision has taken place because a network card is setup to be able to transmit and receive on different physical wires that make up the actual Ethernet cable. This is not possible in wireless Ethernet because when a wireless LAN card is transmitting it can not listen to detect if packets collide.

To partially cope with the inability to detect a collision, the IEEE 802.11 standard attempts to avoid collisions using carefully designed waiting periods that allow multiple users to defer access to the shared wireless channel to one another. That is, IEEE 802.11 clients will always ensure a channel has been idle for a certain period of time before transmitting. The process of deciding how long to wait as governed by the basic DCF is illustrated via a flowchart in Figure A.9.

The following presents a basic overview of how the DCF progresses.

1. In Figure A.9, flow begins at the top left portion of the diagram. When a data packet is ready for transmission, a client will first sense the medium. If it is idle and remains so for a period of time know as the Interframe Spacing (IFS) period, the packet can be immediately transmitted. A standard, unicast packet (called a directed MSDU or MAC Service Data Unit) needs to be acknowledged by the receiver by a short ACK packet. If this transmission fails by the ACK not being received, the client enters the same defer state as if the medium was initially detected to be busy during the IFS period.

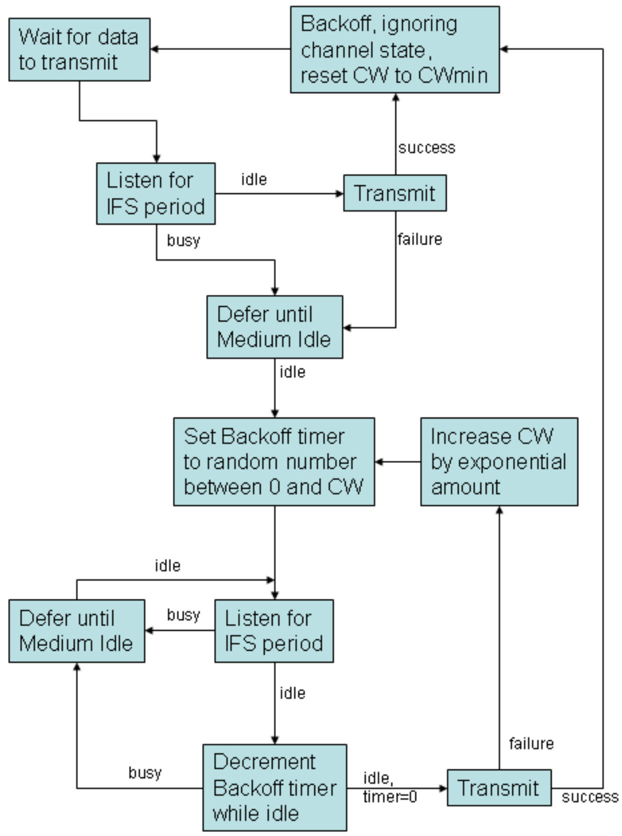

Figure A.9: Flow chart of the process of sending data packets under the CSMA/CA based Distributed Coordination Function (DCF). Created using information in [30].

Note that their are four different types of IFS frame spacings. These are shown in Table A.4. The shorter IFS frame spacings are used for higher priority transmissions and to ensure certain events like acknowledgments (ACKs) will occur before another packet transmission.

2. If the medium was not idle or the transmission fails, the client must defer until the medium is free. This is done using a special timer known as the Network Allocation Vector (NAV). To set this timer the client reads a field in the header of the packet currently being transmitted that tells the client how long the current user will continue to use the medium, through the current transmission, or through the current transmission and immediate transmissions after the current transmission. In this way the client does not need to continue to sense the state of the channel until the NAV timer has expired.

3. After the NAV timer has expired or the client has sensed that the channel is no longer busy, the client will calculate a backoff interval. This backoff interval is a uniform random number. This number is chosen from the interval between 0 and the value of the Contention Window (CW) inclusive. The contention window is initially set to be equal to CW_{min} which is a value defined by the PHY layer.

4. After calculating the value of the backoff interval to use, the senses the channel for an IFS period. If the channel is idle at the end of this period the client will set a backoff timer equal to the value of the backoff period calculated previously. This timer is periodically decremented while the channel continues to stay idle. If the channel becomes busy either during the IFS period or during while the backoff counter is being decremented but before it reaches zero, the client goes into a defer state *without* changing the value of the backoff timer.

5. If the client goes back into a defer state it goes through the same process as before in which the client waits for the medium to become idle by first sensing the medium and then by setting the NAV timer. When the medium returns to an idle state, the client must wait for an additional IFS period before continuing to decrement the backoff counter.

   Note that the backoff counter does not get reset to the initially calculated backoff interval each time the client goes into the Defer state. Rather, the backoff counter is decremented whenever the medium has been idle for at least an IFS period.

6. When the medium has been idle for an IFS period and the backoff counter reaches zero, the client will transmit it's data.

7. If the client discovers that the transmission has failed then the client must exponentially increase the value of CW using equation A.5. As a result, if the medium is very busy, exponential increases in the maximum backoff delay will occur and the probability of packet collisions will decrease. After increasing CW, the client generates a new value for the backoff interval and re-senses the state of the channel.

8. After either of the two transmit states have been completed successfully (by having been properly acknowledged by the receiver using an ACK packet), several things happen. First, the value of CW is reset to CW_{min} after successful transmission occurs. Second, the client goes through a mandatory backoff interval in which the

state of the medium is ignored. The client then goes back to the initial state in which the client waits for data to be ready for transmission.

$$CW_{new} = min(2 * (CW_{old} + 1) - 1, CW_{max}) \tag{A.5}$$

A.3.4 RTS/CTS and the Hidden Terminal Problem

The DCF implementation of IEEE 802.11 does not handle a problem referred to as the hidden terminal problem. This problem occurs when a mutual receiver is in range of two transmitters which are not in range of one another. In this case attempting to detect if the medium is free does not necessarily work because the two transmitters can not detect one another's transmissions. Thus the packets from the two transmitters will collide at the common receiver. To combat this problem, IEEE 802.11 adds an optional RTS/CTS mechanism. In this technique instead of transmitting a data packet after waiting for a free medium, a client will transmit a short Ready To Send (RTS) packet to request the use of the medium. If this succeeds, the receiver will quickly (after a SIFS period) reply with a short Clear To Send (CTS). After the successful exchange of an RTS/CTS pair the actual transmission takes place. This method allows hidden terminals to hear either a CTS or an RTS packet and know to differ access using the NAV functionality described previously. It also means that if packets do collide only a short RTS or CTS packet is lost rather than a long data packet. It is important to note though that this functional is optional to include and is enabled in one of three modes: always on, always off or on for packet sizes above a certain threshold.

A.3.5 Additional Optional Provisions of the MAC Layer

The Point Coordination Function

Another optional protocol that is part of the IEEE 802.11 standard is the Point Coordination Function (PCF). This function allows time critical or delay sensitive packets to be given priority over regular data transmissions. The PCF uses a polling procedure to setup a contention free period which takes priority over the DCF procedure. During the PCF established contention free period, a single host poles clients and allows them to transmit. In this way delay sensitive packets such as voice or video can be given priority over other data.

Wired Equivalent Privacy

The Wired Equivalent Protocol (WEP) is intended to provide a simple layer of protection for wireless network connections. By their very nature, wireless networks are easy to connect to and be eavesdropped on. WEP is a single shared key system in which a basic 40 bit encryption is applied to packet transmissions on the network. Without knowledge of the WEP key, packets can not easily be decoded by an unauthorized user. The use of WEP though tends to slow down transmissions and increases the overhead of packet transmissions, thereby lowering the bandwidth available. In addition, several security flaws have been found in the technique [3].

Bibliography

[1] Bianchi, G.;Fratta, L.; Oliveri, M. “Performance Evaluation and Enhancement of the CSMA/CA MAC Protocol For 802.11,” *Seventh IEEE International Symposium on Personal, Indoor and Mobile Radio Communications, Proceedings, (PIMRC ’96)*, Vol. 2, p. 392-396, 1996.

[2] Bing, B. “Measured Performance of the IEEE 802.11 Wireless LAN,” *Local Computer Networks, 1999. Conference on. (LCN ’99).* p. 34-42. 1999.

[3] Borisov, N.; Goldberg, I; Wagner, D. “(In)Security of the WEP algorithm.” http://www.isaac.cs.berkeley.edu/isaac/wep-faq.html Viewed: July 4, 2001

[4] Chen, S-A; Lee, Y-H; Yen, R.Y.; Zheng, Y-J; Ko, C-H; Sheu, S-T; Chen, M-H. “Optimal prediction tool for wireless LAN using genetic algorithm and neural network concept.” *Communications, 1999. APCC/OECC ’99. Fifth Asia-Pacific Conference on ... and Fourth Optoelectronics and Communications Conference*, Vol. 1, p. 786 -789, 1999.

[5] Demir, T.; Komar, C.; Ersoy, C. “Measured Performance of an IEEE 802.11 Wireless LAN,” *Proceedings of the Fifteenth International Symposium on Computer and Information Sciences*, Istanbul, Turkey. p. 246-254, Oct 2000.

[6] Duchamp, D.,; Reynolds, N. F., “Measured Performance of a Wireless LAN,” *Local Computer Networks, 1992. Proceedings., 17th Conference on.* p. 494-499, 1992.

[7] Durgin, G.; Rappaport, T. S.; Xu, H., ”Measurements and Models for Radio Path Loss and Penetration in and Around Homes and Trees at 5.85 GHz,” IEEE Transactions on Communications, Vol. 46, No. 11, November 1998, pp. 1484-1496.

[8] Feigin, J.; Pahlavan, K. “Measurement Of Characteristics Of Voice Over IP In A Wireless LAN Environment,” IEEE International Workshop on Mobile Multimedia Communications, 1999. (MoMuC ’99), p236-240.

[9] Hope, M.; Linge, N., “Determining the Propagation Range of IEEE 802.11 Radio LAN’s for Outdoor Applications,” *Local Computer Networks, 1999, Conference on. (LCN ’99).* p. 49-50, 1999.

[10] Kamerman, A.; Aben, G. "Throughput performance of wireless LANs operating at 2.4 and 5 GHz," *Personal, Indoor and Mobile Radio Communications, 2000, The 11th IEEE International Symposium on, (PIMRC 2000)*, Vol 1, p. 190-195, 2000.

[11] Kim, S-C.; Bertoni, H. L.; Stern, M. "Pulse Propagation Characteristics At 2.4 GHz Inside Buildings," *IEEE Transactions on Vehicular Technology*, Vol. 45, Iss. 3, p. 579-592, Aug. 1996.

[12] Khurana, S.; Kahol, A.; Jayasumana, A. P. "Effect of Hidden Terminals on the Performance of IEEE 802.11 MAC Protocol," *IEEE Conference on Local Computer Networks Proceedings, Piscataway, NJ*, p. 12-20, 1998.

[13] Leskaroski, D; Mikael, W. B. "Frequency Planning and Adjacent Channel Interference in a DSSS Wireless Local Area Network," *Wireless Personal Communications: Bluetooth Tutorial and Other Technologies*, p. 169-180, 2001.

[14] Maeda, Y.; Takaya, K.; Kuwabara, N. "Experimental Investigation Of Controlling Coverage Of Wireless LAN By Using Partitions With Absorbing Board," *Electromagnetic Compatibility, 1999, International Symposium on*, p. 674 -677, 1999.

[15] Maeda, Y.; Takaya, K.; Kuwabara, N. "Experimental Investigation of Propagation Characteristics of 2.4 GHz ISM-Band Wireless LAN in Various Indoor Environments," *IEICE Transactions in Communications*, Vol. E82-B, No. 10, Oct 1999.

[16] MATLAB 6.0 Product Documentation. Supplied with the MATLAB Product from Mathworks, Inc. 2001. Available at: http://www.mathworks.com/access/-helpdesk/help/helpdesk.shtml.

[17] Messier, A.; Robinson; Pahlavan, K. "Performance Monitoring of a Campus Area Network." *22nd Annual Conference on Local Computer Networks, Proceedings*, p. 232-238, 1997.

[18] Motley, A. J.; Keenan, J. M. P. "Personal Communication Radio Coverage in Buildings at 900 MHz and 1700 MHz," *Electronics Letters*. Vol 4, Issue 12, p. 763-764, Jun 1998.

[19] http://www.netiq.com/Products/Network_Performance/Chariot/Default.asp. "NetIQ Products - Chariot" Downloaded on: June 14, 2001.

[20] Papoulis, A. *Probability, Random Variables, and Stochastic Processes*. 3rd Ed. New York: WCB/McGraw-Hill, 1991.

[21] Peterson, L. L.; Davie B. S. *Computer Networks, A Systems Approach*. 2nd Ed. San Francisco, California: Morgan Kaufmann, 2000.

[22] Prasad, A.R.; Prasad, N.R.; Kamerman, A.; Moelard, H.; Eikelenboom, A. "Indoor Wireless LANs Deployment" *Vehicular Technology Conference Proceedings, 2000. IEEE 51st, (VTC 2000-Spring Tokyo.)*, Vol. 2, p. 1562-1566, 2000.

[23] Prasad, N. R. "IEEE 802.11 System Design" *Personal Wireless Communications, 2000, IEEE International conference on*, p. 490-494, 2000.

[24] Rappaport, T. S. *Wireless Communications, Principles and Practice*. 7th Ed. Upper Saddle River, NJ: Prentice-Hall, 1996.

[25] Seidel, S. Y., Rappaport, T. S., "914 MHz Path Loss Prediction Models for Indoor Wireless Communications in Multifloored Buildings." *IEEE Transactions on Antennas and Propagation*, Vol. 40, No 2, p.207-217, Feb 1992.

[26] SitePlanner 2000 Product Documentation. Supplied with the SitePlanner Product from Wireless Valley Communications, Inc. 2000.

[27] Takaya, K.; Maeda, Y.; "Experimental And Theoretical Evaluation Of Interference Characteristics Between 2.4-GHz ISM-Band Wireless LANs," *IEEE International Symposium on Electromagnetic Compatibility, 1998*, Vol. 1, p. 80-85, 1998.

[28] Tarng, J. H.; Liu, T. R. "Effective Models in Evaluating Radio Coverage on Single Floors of Multifloor Buildings," *IEEE Transactions on Vehicular Technology*, Vol. 48, No. 3, May 1999.

[29] WaveLAN Access Point 500 and ORiNOCO Gold PCMCIA Card pDocumentation. 2001. Available at: http://www.wavelan.com/template.html?section=m116&-envelope=205

[30] "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications" *IEEE 802.11 Standard*, 1999.

[31] "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications: Higher Speed Physical Layer Extension in the 2.4 GHz Band", *IEEE 802.11b Standard*, 1999.

[32] Xylomenos, G.; Polyzos, G. C. "Internet Protocol Performance Over Networks With Wireless Links," *IEEE Network*, Vol. 13, Iss 4, p. 55-63, July-Aug 1999.

[33] Xylomenos, G.; Polyzos, G. C. "TCP And UDP Performance Over A Wireless LAN," *Eighteenth Annual Joint Conference of the IEEE Computer and Communications Societies, Proceedings. (INFOCOM '99)*, Vol. 2, p. 439-446, 1999.

[34] Xylomenos, G.; Polyzos, G. C.; Mahonen, P.; Saaranen, M. "TCP Performance Issues Over Wireless Links," *IEEE Communications Magazine*, Vol. 39, No. 4, p. 52-58, April 2000.

[35] Yang, S-H.; Lee, Y-H.; Yen, R.Y.; Zheng, Y-J; Shue, S-T; Ko, C-H; Chen; M-H. "A Wireless LAN Measurement Method Based on RSSI and FER," *Fifth Asia-Pacific Conference on Communications and Fourth Optoelectronics and Communications Conference, Proceedings, (APCC/OECC '99)*, Vol. 1, p. 821-824, 1999.

[36] Zyren, J; Petrick, A. "Brief Tutorial on IEEE 802.11 Wireless LANs," Intersil Application Note 9829 (AN9829). http://www.intersil.com/data/an/an9/an9829/-an9829.pdf Downloaded: Fall 1999, Written: Feb 1999.

Vitae

Benjamin E. Henty

Benjamin Henty began his college career as an Electrical and Computer Engineering student at Carnegie Mellon University as an Andrew Carnegie Scholar. He attended Carnegie Mellon from August 1995 through May of 1999, at which time he graduated with University and College Honors.

At Carnegie Mellon, Benjamin performed research as an undergraduate in the area of IEEE 802.11 wireless LAN design. The research work involved investigating the feasibility and performance of using heating and air conditioning ventilation ducts for distributing wireless LAN signals throughout a building. This research included a large number of RF and network measurements as well as antenna design work. This work was presented in a conference paper at the 8th Annual Virginia Tech Symposium on Wireless Personal Communications in 1998 and at the Carnegie Mellon Meeting of the minds in 1999.

While at Carnegie Mellon, Benjamin worked as a sound engineer as part of the student activities program. From April 1998 through April 1999 he was the head supervisor for the Activities Board Technical Committee (AB Tech). The responsibilities of this position included directing 15 students in the safe use of professional sound and light systems for campus events. These events varied between national label entertainers to small student group presentations. In this position, Benjamin worked with Carnegie Mellon organizations, outside artists, and agents to organize technical details of over 150 events, or an average of one show every two days. In addition, Benjamin oversaw the repair and maintenance of all sound, light and power equipment and was fiscally responsible for a $35,000 budget. Benjamin received a Student Leadership Award for his work in this position.

Benjamin worked as an RF Engineering intern at Daimler Chrysler Rail Systems (ADTranz) in the summer of 1999. At Daimlet Chrysler, Benjamin developed and performed several test specifications on various different wireless modems. Benjamin analyzed and measured the performance of several wireless data modems, microwave and ultrasonic train detection sensors, and voice radios. These measurements were performed at the companies test track and on the Pittsburgh Airport people mover system. All results were written up in internal test specifications and presented to other engineers at Daimler Chrysler.

Benjamin began his masters work at Virginia Tech University in August of 1999 as a Bradley Fellow. Benjamin again performed research with IEEE 802.11 and IEEE 802.11b wireless LANs. This research involved using simulation tools to design indoor wireless LAN networks. The work was used as the basis for his masters thesis. The goal of his research was to develop throughput prediction models that would simplify computer-aided design of wireless LAN systems. This research was performed at the Mobile and Portable Radio Research Group (MPRG) of Virginia Tech under the supervision of Dr. Theodore S. Rappaport.

Benjamin became a Bradley Industrial Fellow in June of 2000 when he began to work at a Virginia Tech spinoff company, Wireless Valley Communications, Inc. The company was a start-up company consisting of approximatly 10 people. At Wireless Valley, Benjamin worked as a Wireless Network Engineer. Benjamin was the product manager of the companies wireless LAN measurement and design software tools, LANFielder and SiteSpy. These products allow location based, site-specific measurement of wireless LAN throughput, and signal strength statistics. In addition to developing, supporting and marketing these products, Benjamin supported the companies propagation channel simulation tools, SIRCIM and SMRCIM. Further, Benjamin submitted a patent, performed consulting work, made customer visits, attended conferences and gave presentations and talks on wireless LAN system design, while working at Wireless Valley. Benjamin left Wireless Valley at the end of April 2001 to complete the remainder of his Masters Degree at Virginia Tech University. This he did in August of 2001.