Measurement and Analysis of the Error Characteristics of an In-Building Wireless Network

David Eckhardt and Peter Steenkiste
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213
{davide,prs}@cs.cmu.edu

Abstract

There is general belief that networks based on wireless technologies have much higher error rates than those based on more traditional technologies such as optical fiber, coaxial cable, or twisted pair wiring. This difference has motivated research on new protocol suites specifically for wireless networks. While the error characteristics of wired networks have been well documented, less experimental data is available for wireless LANs.

In this paper we report the results of a study characterizing the error environment provided by AT&T WaveLAN, a commercial product designed for constructing 2 Mb/s in-building wireless networks. We evaluated the effects of interfering radiation sources, and of attenuation due to distance and obstacles, on the packet loss rate and bit error rate. We found that under many conditions the error rate of this physical layer is comparable to that of wired links. We analyze the implications of our results on today's CSMA/CA based wireless LANs and on future pico-cellular shared-medium reservation-based wireless networks.

1 Introduction

Two major trends in networking today are support for reservation-based applications over wired networks [12, 35, 10, 1, 11, 17, 41] and making connectivity ubiquitous via wireless technologies[30, 23, 25, 14, 39, 21]. There is general belief that networks based on fiber or electrical connections have excellent error characteristics but that wireless networks typically have extremely high error rates. The high error rates in wireless LANs are considered a major challenge and research groups have considered solutions ranging from the use of Forward Error Correcting (FEC) codes that in effect improve the error rate seen at higher levels[22], to the use of special transport level protocols that treat wireless links in a special way[13, 3, 15, 4, 40]. However, while the error characteristics of wired networks have been well documented, relatively little experimental data is available for wireless LANs.

There are some obvious reasons why one would expect wireless connections to have higher error rates than wired connections. Wired connections largely isolate the signal that carries the encoded data from other signals, especially in the case of optical fiber. In contrast, wireless signals share the same propagation medium with many competing signals, and as a result, there are many more opportunities for interference that can result in bit errors. One example of wireless communication that has been studied widely is satellite communication: error rates can be very high and satellite links make extensive use of FEC to improve communication performance[31, 36]. However, wireless connections are very diverse: they differ in range, bandwidth, frequency spectrum used, modulation techniques, interference sources, and physical environment. As a result, it is difficult to generalize the results from one domain (e.g. satellite communication) to another (e.g. in-building LAN).

One reason why there is less information available on the error rates in wireless networks compared with their wired counterparts is that characterization of the error environment is much more complicated. When signals propagate through space many more factors can influence signal quality than when they propagate in an electrical conductor or fiber. However, characterizing the environment is a critical step in providing a reliable communication service to applications. Information on the frequency and nature of errors is needed to select the method of dealing with the problem. Solutions range from new transport level protocols that typically isolate wireless network segments from wired segments, to changes in the physical layer, e.g. retransmission or FEC. For the last class, the most appropriate solution depends in part on the nature of the error patterns.

In this paper we investigate the error behavior of the WaveLAN network[37], which was designed as a "wireless Ethernet" system. Our goal is to determine whether we will be able to extend the services that are currently being developed in backbones, e.g. support for real-time or near-real-time guarantees, to the wireless environment. This would make it possible to support many traffic types including video. The trends in the technologies used to implement wireless networks suggest that it will be practical to build inexpensive wireless networks with data rates of 10-40 Mbit/second in the near future[6]. Low price makes it feasible to deploy many base stations, which results in smaller coverage areas and thus fewer mobile hosts competing for that data rate. However, high error rates can significantly reduce the effective bandwidth available to users, so controlling the error rate is critical.

Wireless network designers currently have many options in the areas of frequency, modulation, framing and scrambling, addressing, and medium access control. While WaveLAN represents only a single point in a large space, we believe it is worthy of study because its Medium Access Control (MAC) protocol and radio system are similar to what we believe will be common in the future. First, we believe that a Time Division Multiple Access (TDMA) MAC layer atop a per-cell shared medium is attractive because TDMA allows flexible bandwidth sharing among stations whose needs will vary

To appear in the *Proceedings of SIGCOMM '96*, August 1996, Stanford, CA.

with time, and because a shared channel should support multicast connections efficiently. Second, direct-sequence spread spectrum (DSSS) is attractive because it provides noise tolerance and can be extended to provide sharp cell boundaries in the form of Code Division Multiple Access (CDMA)[28]. In fact, there is a similar product, Arlan[27], which uses the same frequency bands and DSSS modulation. We expect that similar systems will react in similar, though not identical, fashions when exposed to the challenges described in this paper.

The remainder of the paper is organized as follows. In Section 2 we first describe the capabilities of the WaveLAN interface. In Sections 3 and 4 we present a brief summary of the sources of wireless errors and our methodology for characterizing the error environment. We then present the results of our study: characterization of in-room, line-of-sight communication (Section 5), measurements of the errors caused by passive obstacles (Section 6) and competing radiation sources (Section 7). We discuss the implications on the architecture of wireless LANs in Section 8 and related work in Section 9.

2 AT&T WaveLAN

WaveLAN[37] is designed to be an affordable, easy-to-install wireless extension of an existing bridged Ethernet system. Products include ISA and PCMCIA network interfaces for PC-compatible computers and stand-alone WaveLAN-to-Ethernet packet bridges. They operate in either the 902-928 MHz or the 2.4-2.8 GHz ISM (Industrial, Scientific, and Medical) license-free band. The PCMCIA units comprise a Type II PCMCIA card and an external unit a little larger and heavier than a deck of cards, and are commercially available for roughly $500 in small quantity.

Internally, the WaveLAN interface contains a standard Intel 82593 single-chip CSMA/CD LAN controller, custom logic for signal processing and modem control, and a custom radio transceiver. The transmitter applies DQPSK modulation to a 2 megabit/s data stream, yielding a 1 megabaud signal. This signal is further modulated by an 11 chip per bit sequence (hence "direct sequence") to produce an 11 MHz wide signal which is transmitted with a power of 500 milliwatts. The receiver selects between two perpendicular antennas and multiple incoming signal paths to combat multipath interference.

The modem control unit prepends a 16-bit "network ID" to every packet on transmit, and can be set to reject all but one network ID on receive. In addition, it informs the host of the channel condition upon each packet arrival by reporting signal level, silence level, signal quality, and antenna selected for each packet. The signal and silence levels (5 bits) are derived from the receiver's automatic gain control (AGC) setting just after the beginning and end of the packet, respectively. Because the MAC protocol discourages simultaneous transmission and immediately consecutive transmissions, measuring the silence level during an inter-packet time is typically a good indication of the amount of non-WaveLAN background interference. The signal quality (4 bits) is sampled just after the beginning of the packet and is derived from the information the receiver uses to select between the two antennas.

As it is difficult to detect collisions in this radio environment, WaveLAN employs a CSMA/CA (collision avoidance) MAC protocol[2]. In CSMA/CD, a station which becomes ready to transmit while the medium is busy will make its first transmission attempt as soon as the medium is free, based on the optimistic assumption that it is the only waiting station. If this assumption is wrong, all waiting stations will quickly learn that when they sense a collision. Since WaveLAN cannot sense collisions, they result in packet losses which must be dealt with by higher layer protocols. WaveLAN CSMA/CA attempts to avoid collision losses by treating a busy medium as a collision. That is, any stations which become ready to transmit while the medium is busy will delay for a random interval when the medium becomes free. Aside from the modified MAC protocol and lower data rate, the 82593 performs all standard Ethernet functions, including framing, address recognition and filtering, CRC generation and checking, and transmission scheduling with exponential backoff.

Many of the techniques cellular radio systems [30, 23, 33] employ to re-use frequencies in nearby areas, such as power control, frequency diversity and code diversity, are easiest to employ in a point-to-point environment. Since WaveLAN follows the Ethernet protocol, where stations multicast to dynamic sets of peers rather than communicating exclusively with a central base station, it would be difficult to estimate the power required to contact a particular set of receivers for each packet transmission, and expensive or complex to synchronize with many stations each using a different spreading sequence. Perhaps for these reasons, WaveLAN does not provide the ability to vary transmit power and does not use frequency or code diversity. Instead, it gives receivers the ability to mask out weak signals through a receive threshold,[1] which improves throughput and may be sufficient to simulate cell boundaries in some cases even though WaveLAN is inherently a single shared channel. In addition, the "network ID" provides multiple logical Ethernet address spaces, which allows WaveLAN-to-Ethernet bridges to use standard bridge routing protocols. We consider the use of the receive threshold in more detail in Section 5.

3 Sources of Wireless Errors

For the purposes of this study, we organize the possible sources of errors in a wireless network into three groups.

The first group contains error sources we investigated and report on in this paper:

- *attenuation.* When electromagnetic energy encounters matter, some of it is lost in the form of heat. WaveLAN can usually penetrate several walls while maintaining a quite good error rate, but we found attenuation to be a significant source of errors.
- *front end overload.* If a very powerful transmitter of one frequency band is near a receiver of another band, the transmitter may overwhelm filters in the receiver and inject substantial noise. As we expect wireless computer networks to be employed in close proximity to microwave ovens and cellular phones, we have made an initial investigation into errors due to front-end overload.
- *narrowband interference.* This is due to an unfriendly transmitter occupying a small frequency band overlapping (perhaps totally) with the band we wish to use. We investigated the effects of two 900 MHz FM cordless phones.
- *spread-spectrum interference.* This is due to an unfriendly transmitter either switching between narrowband frequencies or spreading its energy simultaneously across a wide frequency band. We have investigated interference between competing WaveLAN transmitters and between WaveLAN and other 900 MHz spread-spectrum sources.

The second group contains error sources which might have an impact on WaveLAN performance, but which we did not or could not control the behavior of:

- *natural background noise.* For example, infrared wireless networks may perform poorly if they are near sources of direct sunlight. We did not attempt to measure or control for background noise.

[1] There is also a threshold which allows filtering based on signal quality, though we do not employ it.

Column Name	Meaning
Packets Received	Test packets received
Packet Loss	Percentage of transmitted test packets that were lost
Packets Truncated	Number of received test packets which were truncated
Bits received	Number of *body* bits received, rounded down
Wrapper Damaged	Number of packets with damaged headers or trailers
Body Bits	Total number of body bits damaged in trial
Worst Body	Number of bits damaged in most-corrupted packet body

Table 1: Column heading explanations

- *multipath interference.* When electromagnetic radiation reflects off of or diffracts around objects, it takes multiple paths between the transmitter and the receiver. Since these paths are typically of different lengths, there will be destructive interference, which can greatly reduce signal strength. We have not studied how multipath effects interfere with WaveLAN, in part because they are difficult to model and study, and in part because WaveLAN is explicitly designed to resist them.

The third group contains error sources we did not consider in this paper:

- *path loss (dispersion).* The intensity of electromagnetic energy reaching a receiver is decreased by distance even in free space. We found WaveLAN did not suffer errors due to path loss even in large lecture halls.
- *motion.* If two communicating objects are moving with respect to each other, the frequency of the electromagnetic energy changes according to the Doppler effect. While this effect may be significant in some radio environments[20], the Doppler shift due to moving a WaveLAN unit at the speed of sound would be substantially less than the inaccuracy of the clock crystals employed by WaveLAN[37]. Hence we have not investigated errors due to motion.
- *data dependent effects.* Some modulation schemes can lose clock synchronization in the face of certain long bit patterns. While we have transmitted a variety of data packets in our investigation, we have not particularly examined WaveLAN for data dependent error patterns, in part because the receiver includes active correction for clock drift.

4 Methodology

To characterize the error environment of the WaveLAN, we monitored the quality of data transfers between two identical DECpc 425SL laptops (25 MHz 80486) running NetBSD 1.0A. For different tests, we placed the PCs in different environments or added competing radiation sources.

The data transfers consisted of specially-formatted UDP datagrams. On the receiver, the kernel device driver was modified to place both the Ethernet controller and the modem control unit into "promiscuous" mode and to log, for each incoming packet, every bit and all available status information, even if the packet failed the Ethernet CRC check. We decided to collect bursts of packets at the maximum possible transmission rate (roughly 1.4 Mb/s for this machine and protocol stack), aggregating multiple bursts to form a long trial. Within each trial, packets consisted of 256 32-bit words wrapped inside UDP, IP, Ethernet, and modem framing. For each packet, the data words were identical to facilitate identification even in the face of substantial noise, and the data value was incremented between packets.

There are many reasons why a transmitted test packet might not be received, and we will consider them in order from the outside of the packet toward the center. First, certain errors might cause the modem unit to miss the beginning-of-frame marker, resulting in a slightly-damaged packet being totally lost. Second, it is possible for a packet to arrive correctly but be lost by the receiver due to unrelated system activity, even though we tried to reduce this to a minimum. Third, errors in the packet headers and trailers might lead the Ethernet or IP layers to discard the packet due to damage or misaddressing. Therefore, we enable promiscuous receive and disable automatic CRC filtering at the Ethernet level and use a heuristic matching procedure to determine whether a given packet is one of the test series. Finally, a packet may arrive but the body may be truncated and/or some bit values may be incorrect. We apply a second heuristic procedure to determine the sequence number of any packet we believe is a test packet. Since the packet body consists of a single word repeated multiple times, truncated packet bodies are ambiguous–it is not possible to know which words are missing. Therefore, we produce an estimated error syndrome (bit corruption pattern) only for those test packets which are damaged but not truncated. Furthermore, we report precise bit error figures only for errors in the data portion of the packet, as reconstructing the exact IP header (and thus the Ethernet CRC trailer) is difficult. Therefore, there are some packets which we know to be damaged but for which we cannot determine the exact number of corrupted bits. Due to these factors, our packet loss rate and bit error rate (BER) figures are necessarily only estimates.

Table 1 explains the column headings we use in most tables. When we present signal level, silence level, and signal quality, we give the minimum observation, mean, standard deviation (in parentheses), and maximum observation. Unless otherwise specified, all runs use a receive threshold of 3 and a quality threshold of 1.

In some trials we received packets from WaveLAN units in nearby rooms or in other buildings. Typically these packets were few, had poor signal characteristics, and were damaged. Frequently we could determine that they were ARP packets or inter-bridge routing packets. We present them, labelled "Outsiders," only when they are significant (due to number or signal characteristics). It is possible for reasons explained above that some packets we identify as outsiders may instead be badly corrupted test packets.

Many independent variables could be investigated for effects on loss and error rates. Examples include: noise emitted by different models of laptop computers, relative position of transmitter and receiver, orientation of antenna, temperature, humidity, position of human bodies, different transmit and receive capabilities of individual units, different materials of walls, ceilings, floors, and furniture, and effects of very distant competing transmitters. Many of these variables have too many values to investigate with any thoroughness, may vary greatly from building to building, or, in the case of interference, may be nearly impossible to measure outside a "clean room" environment. Therefore, our approach for this study was to select certain variables, manipulate them coarsely, and focus

Trial Name	Packets Received	Packet Loss	Packets Truncated	Bits Received	Wrapper Damaged	Body Bits	Worst Body
office1	102720	.03%	1	8×10^8	0	0	–
office2	40080	0%	0	3×10^8	0	0	–
office3	102720	.01%	0	8×10^8	0	0	–
office4	122159	.02%	0	10^9	0	0	–
office5	488399	.07%	0	4×10^9	0	0	–
office6	122160	.04%	0	10^9	1	1	1
office7	122160	.02%	0	10^9	1	0	–
office8	125040	.02%	0	10^9	0	0	–
office9	122160	.02%	0	10^9	0	0	–

Table 2: Results of in-room experiment

on those with the most obvious effect, i.e. attenuation, obstacles and active radiation sources.

5 In-room line-of-sight communication

We discuss WaveLAN behavior in the best case, two stations communicating in the same room. We first present the results of a series of long-running trials designed to estimate the BER of the link under good conditions. We then briefly describe how signal propagation is affected by distance in the absence of obstacles, and finish by examining the effect of the receive threshold on in-room communication.

5.1 Base case

In Table 2 we present the results of several long trials in an office for a signal level of approximately 29.5.

Two points are worthy of note. First, the bit error rate is very low. These trials represent more than 10^{10} bits, and we have experienced very few errors. This is certainly low enough for optimism about extending even fairly error-intolerant applications to a wireless network. Second, some process is causing packets to be lost even in a near perfect environment, though at a rate of well under one per thousand. This could be due to some critical host resource being overloaded, or perhaps could indicate that the modem unit's AGC occasionally reacts too slowly and causes the beginning of a packet to be missed. In a running network, this loss would probably correspond to less than one packet per second, which would plausibly be unnoticed even by a multimedia application.

5.2 Path loss

Next we look at how the signal level changes as a function of distance. This is shown in Figure 1. In this experiment the receiver is held fixed against one wall of a large lecture hall while the transmitter is moved away from it to various distances (the zero point represents the two modem units in physical contact). To the extent that the setup approximates omnidirectional antennas in free space, one would expect to see a smooth dropoff in signal level as distance increases. Indeed, that is the dominant theme. The dips at six and thirty feet are probably due to multipath interference (similar WaveLAN non-monotonicity was observed in [16]), and are likely to be particular to the room where the measurements were taken.

Because the WaveLAN design addresses threats such as multipath effects and narrowband interference, signal level is an important predictor of error rate, and an interesting question is at what point the signal level is too low to receive packets reliably. Table 3

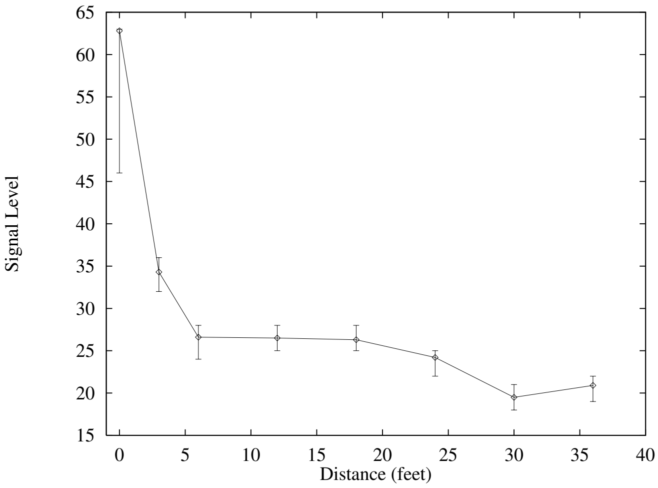

Figure 1: Signal level as a function of distance. Error bars represent minimum and maximum observed signal levels.

presents the aggregated results of several trials, with slight variations of receiver position, orientation, and obstacles within each trial. While undamaged packets may have a signal level as low as 5, and damaged packets one as high as 12, the main body of damaged packets has signal levels below 8, whereas it is well above 8 for undamaged packets. In this trial, we observed several "outsider" packets, which we believe to be from other buildings; the most striking difference between the damaged (which could be truncated and/or corrupted) and undamaged packets is their signal quality.

Figure 2 summarizes the results on signal level and errors. In general, as the distance between the transmitter and the receiver increases (x-axis) the signal level decreases (y-axis). In our in-room experiments, a signal level of roughly 10 is sufficient for the receiver to receive packets reliably with extremely low error rates. When the signal level drops below 8, the error rate becomes very high; this "error region" is shaded in Figure 2. Finally, using the receive threshold described in Section 2, the receiver can mask out packets with a signal level below a certain threshold, as described below.

5.3 Receive threshold

In order for the receive threshold to divide an indoor space into pseudo-cells, it must allow us to communicate with stations we want to include, but fully exclude more distant stations we wish to screen out. When investigating pseudo-cellular divisions we must

Packet Type	Packets Received	Level				Silence				Quality			
		↓	μ	σ	↑	↓	μ	σ	↑	↓	μ	σ	↑
All test packets	8634	4	14.15	(6.32)	27	0	2.89	(1.64)	9	2	14.80	(0.86)	15
Undamaged	7942	5	14.74	(6.23)	27	0	2.83	(1.64)	9	3	14.94	(0.37)	15
Truncated	107	4	6.20	(1.44)	10	0	3.48	(1.34)	6	6	10.07	(1.74)	15
Wrapper damaged	9	6	7.56	(0.83)	9	0	2.89	(1.79)	6	10	12.89	(2.18)	15
Body damaged	576	4	7.52	(1.45)	12	0	3.61	(1.46)	9	2	13.80	(1.75)	15
Undamaged outsiders	73	5	6.23	(0.77)	8	0	2.63	(1.12)	5	9	14.49	(0.99)	15
Damaged outsiders	867	2	5.19	(1.87)	18	0	3.26	(1.37)	17	1	7.49	(2.06)	15

Table 3: Packet error conditions versus signal metrics

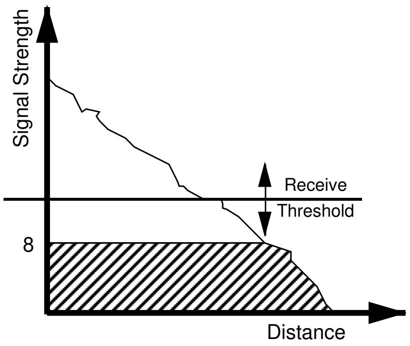

Figure 2: Signal level

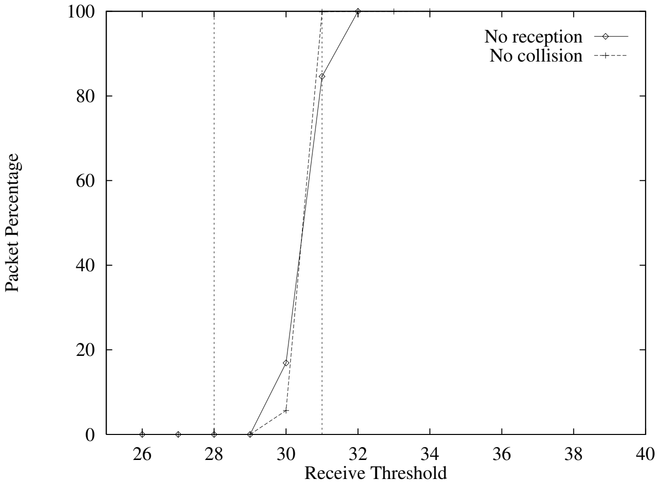

Figure 3: Effects of receive threshold.

consider two possible interactions between a system in one cell (the "victim") and a second system that we would like to be in a different cell (the "enemy"). First, packets from the "enemy" might interfere with packets the "victim" wishes to receive from nearby neighbors. Second, the carrier generated by the "enemy" system might be heard and might prevent the "victim" from transmitting; this would reduce aggregate throughput since only one system can send at any given time. Ideally, raising the threshold would filter out distant stations and hide carrier sense from the Ethernet chip, without causing packets from desirable stations to be lost or corrupted.

We performed an experiment to verify this. One station, the "enemy," was configured to transmit packets continuously. As the "victim" station varied its receive threshold through a window around the received packets' signal level, we observed both the packet loss rate from the "enemy" and the collision rate[2] when the "victim" attempted to transmit. The results are presented in Figure 3. Each packet loss figure is based on at least 1,400 transmitted packets, and each collision figure is based on at least 10,000 transmission attempts. The vertical lines represent the minimum and maximum signal level received during the trial, and the two curves represent the percentage of packets which are filtered out and the percentage of transmissions completed without collision. Ideally, both curves would range from 0% at the left line, representing no filtering when the threshold is set to the signal level of received packets, to 100% at the right line. As the figure shows, the threshold is not perfect, and we have observed that it is wise to allow a margin of several units when choosing a threshold. A useful feature of the receive threshold is that it seems to cleanly filter packets. That is, we did not receive any damaged or truncated packets in the course of the trial. These results indicate that the receive threshold enables frequency re-use in situations where WaveLAN hosts are clustered with significant signal attenuation between clusters.

[2]Recall that WaveLAN considers "medium busy" a collision.

5.4 Summary of in-room operation

Evidence we have presented in this section gives us reason to hope that WaveLAN could be employed to experiment with pseudo-cellular networks. In typical cases, the bit error rate and packet loss rate of the network are very low, and there is reason to believe this is true throughout even large rooms. Furthermore, there seems to be a plausible mechanism for excluding interfering stations if they are are separated from a cluster of cooperating stations by reasonable distance.

6 Errors due to passive obstacles

In this section, we focus on mundane obstacles which we expect WaveLAN stations to typically encounter. We present the results of three experiments. In the first, we measure the effect created by a single wall. In the second, we hold a receiver in a fixed position and measure the propagation environment to it from four transmitter locations. Because interposing a wall requires moving one unit, and because small position changes may result in noticeable propagation changes, we cannot be certain how much the wall itself is to blame. However, we have observed this correlation in several other cases we do not present. Finally, we interpose a human body between two WaveLAN units and observe the effects on data transport and signal metrics.

Trial Name	Packets Received	Level				Silence				Quality			
		↓	μ	σ	↑	↓	μ	σ	↑	↓	μ	σ	↑
Air 1	12720	29	30.58	(0.70)	32	0	1.80	(1.38)	7	15	15.00	(0.00)	15
Wall 1	12720	24	25.78	(0.67)	28	0	1.25	(1.27)	7	14	15.00	(0.03)	15
Air 2	12715	25	28.58	(0.60)	30	0	3.35	(1.11)	13	15	15.00	(0.00)	15
Wall 2	12720	25	26.66	(0.59)	28	0	3.25	(1.10)	8	15	15.00	(0.00)	15

Table 4: Signal metrics with a single wall

Trial Name	Packets Received	Packet Loss	Packets Truncated	Bits Received	Wrapper Damaged	Body Bits	Worst Body
Tx1	12715	0%	0	10^8	0	0	–
Tx2	12720	.007%	0	10^8	0	0	–
Tx4	1440	.07%	0	10^7	0	0	–
Tx5	1440	.07%	1	10^7	0	82	7

Table 5: Results of multi-room experiments

Trial Name	Packets Received	Level				Silence				Quality			
		↓	μ	σ	↑	↓	μ	σ	↑	↓	μ	σ	↑
Tx1	12715	25	28.58	(0.60)	30	0	3.35	(1.11)	13	15	15.00	(0.00)	15
Tx2	12720	25	26.66	(0.59)	28	0	3.25	(1.10)	8	15	15.00	(0.00)	15
Tx4	1440	11	13.81	(0.66)	15	0	4.03	(1.27)	9	15	15.00	(0.00)	15
Tx5	1440	7	9.50	(0.93)	12	0	4.08	(1.29)	8	11	14.99	(0.15)	15

Table 6: Signal metrics for multi-room experiment

Packet Type	Packets Received	Level				Silence				Quality			
		↓	μ	σ	↑	↓	μ	σ	↑	↓	μ	σ	↑
All	1440	7	9.50	(0.93)	12	0	4.08	(1.29)	8	11	14.99	(0.15)	15
Error-Free	1414	7	9.51	(0.92)	12	0	4.08	(1.28)	8	14	15.00	(0.03)	15
Truncated	1	10	10.00	(0.00)	10	3	3.00	(0.00)	3	12	12.00	(0.00)	12
Body Damaged	25	7	8.72	(0.87)	10	1	4.56	(1.27)	8	11	14.72	(0.87)	15

Table 7: Signal metrics for multi-room scenario Tx5

6.1 Single wall

In the first scenario a transmitter and receiver are separated by approximately 7 feet, and then further separated by approximately 6 inches of wall (in the second case, approximately four feet of free space were added in addition to the wall). We did experiments with two different types of wall. In each location we collected 10^8 bits with no loss or error whatsoever.

However, the WaveLAN reports slightly different propagation quality numbers for the different scenarios (Table 4). We see that the wall affects the signal *level* in a way that is similar to a noticeable move across a room, though the signal *quality* is not significantly reduced. The first wall is plaster with a wire mesh core and it reduces the signal level by about 5 points. The second wall consists of concrete blocks and reduces the signal level by only 2 points, i.e. concrete walls seem to be less of a hindrance for these signals than plaster over wire mesh walls.

6.2 Multiple obstacles

In the second experiment we use a more complex setup in the building with concrete block walls (wall 2 of the previous experiment). The layout is shown in Figure 4, and the results of the test are summarized in Table 5. The receiver and the first transmitter location are at diagonally opposite sides of a single office, and we obtain results similar to the previous in-office case. The second transmitter location is approximately four feet away through a single concrete block wall, and corresponds to the single wall experiment of the previous section. The other two transmitter locations are more distant, with several intervening walls and metal objects. They are at distances of roughly 45 and 30 feet respectively from the receiver. The fourth transmitter location shows us our first corrupted packet bodies. Twenty-five of the received packets have a total of 82 bit errors, with the worst packet containing seven bit corruptions. While this number is trivial to correct using error coding, the existing WaveLAN system does not include such a mechanism. The fourth transmitter location demonstrates a single packet truncation of roughly 10% of the packet body.

It is instructive to make a more detailed comparison of the propagation environments along these four paths (Tables 6 and 7). First we see that moving through a single wall from Tx1 to Tx2 has not significantly changed the propagation environment. Next we observe that passing through more walls and a door has a more noticeable effect. Finally, if we look in more detail at the Tx5 location, we see that the corrupted packets have noticeably reduced

Trial Name	Packets Received	Packet Loss	Packets Truncated	Bits Received	Wrapper Damaged	Body Bits	Worst Body
No body	1440	0%	0	10^7	0	0	-
Body	1442	.14%	3	10^7	1	224	27

Table 8: Effects of human body on packet loss and errors

Trial Name	Packet Type	Packets Received	Level ↓	Level μ	Level σ	Level ↑	Silence ↓	Silence μ	Silence σ	Silence ↑	Quality ↓	Quality μ	Quality σ	Quality ↑
No body	All Packets	1440	11	12.55	(0.60)	15	0	4.23	(1.47)	13	15	15.00	(0.00)	15
Body	All Packets	1442	5	6.73	(0.88)	10	0	1.90	(1.64)	9	11	14.95	(0.27)	15
	Undamaged	1214	5	6.73	(0.87)	10	0	1.73	(1.59)	9	14	14.99	(0.09)	15
	Truncated	3	7	7.67	(0.47)	8	0	3.00	(2.16)	5	11	12.33	(1.25)	14
	Wrapper damaged	1	5	5.00	(0.00)	5	1	1.00	(0.00)	1	15	15.00	(0.00)	15
	Body damaged	224	5	6.74	(0.95)	9	0	2.85	(1.61)	7	13	14.77	(0.51)	15

Table 9: Effect of human body on signal measurements

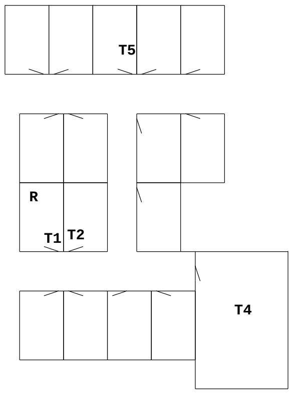

Figure 4: Layout of multi-room experiment

signal *level*, while the sole truncated packet has a noticeably reduced signal *quality*. We will see similar correlations in other scenarios, so they may point to some feature of the WaveLAN hardware. For example, it is possible that data decoding and clock recovery are impaired by different signal features.

From this brief investigation of obstacles, we learn that different construction materials have noticeably different effects on WaveLAN propagation. Furthermore, it seems unlikely that there are many cases where a single building wall can be pressed into service as a cell boundary. The receive threshold experiment in Section 5 showed that packets from a transmitter to a receiver can vary in signal strength by several points and that the threshold is not perfect. In other words, if we want to use the receive threshold to shut out senders in an adjacent cell, the difference in average signal level for senders inside and outside of the cell should be at least 6, although 8-10 would be more desirable. The results of this section predict that it will typically require multiple walls to safely isolate two transmitters in different offices. Unfortunately that introduces large "border zones" in which mobile clients would disrupt multiple pseudo-cells.

6.3 Human Body

After noticing in a few informal trials that WaveLAN signals seemed significantly attenuated by a human body, we decided to investigate further. In order to obtain a path with significant attenuation, we separated two WaveLAN units by placing them in two rooms across a hallway. The direct path was approximately 56 feet long and passed through two concrete block walls and some classroom furniture. We collected two packet streams, with the second impaired by the presence of a person bending over as if to examine the laptop screen closely. Tables 8 and 9 contain the results. Interposing a person has induced packet loss, truncation, and packet body damage. Furthermore, we observe a noticeable reduction in signal level. These results could be significant if similar networks were deployed in crowded lecture halls, though the effect would probably be mitigated if a WaveLAN base station were located near the ceiling.

7 Errors due to competing active radiation sources

We briefly discuss four different types of interference: front end overload, narrowband interference, spread spectrum interference, and competing WaveLAN units.

7.1 Front end overload

If a very powerful transmitter is close to a receiver, the early filter stages of the receiver, which are designed to reject out-of-band signals, may be overwhelmed. We tested two sources of front end overload. The first was a 144 MHz Amateur Radio Service FM transmitter emitting roughly two watts while in physical contact with the receiver's modem unit. During this test we observed no bit errors unless we separated the two WaveLAN units far enough for the signal level to be severely attenuated, which is itself a source of errors.

Microwave ovens are powerful sources of potential interference which are common in an office environment. Though we expect them to be well shielded, even a small leakage percentage would be significant. We made a crude test of a single microwave oven. The transmitter was placed at varying distances from the receiver, which was in physical contact with a microwave oven (operating with the door closed), and no errors were observed. Since most microwave ovens operate in the 2 GHz range, it is possible that 2.4 GHz WaveLAN units would receive more interference.

7.2 Narrowband interference

We briefly investigated the effects of two narrowband 900 MHz cordless phones (AT&T 9100, Panasonic KX-T9500) on Wave-

Trial Name	Packet Source	Packets Received	Level ↓	Level μ	Level σ	Level ↑	Silence ↓	Silence μ	Silence σ	Silence ↑	Quality ↓	Quality μ	Quality σ	Quality ↑
Phones off	Test	1432	25	26.71	(0.66)	28	0	2.40	(1.32)	6	14	14.98	(0.14)	15
	Outsiders	330	2	4.69	(1.56)	25	0	2.57	(1.37)	6	1	6.67	(1.99)	15
Cluster	Test	1440	25	26.89	(0.61)	28	8	15.45	(1.22)	19	15	15.00	(0.00)	15
Handsets nearby	Test	1440	25	26.62	(0.62)	28	3	11.33	(2.62)	15	13	14.93	(0.26)	15
Handsets nearby talking	Test	1431	25	26.63	(0.59)	28	0	6.11	(1.98)	10	14	14.96	(0.20)	15
	Outsiders	213	6	7.85	(0.87)	11	2	6.26	(1.73)	11	1	9.93	(1.93)	15
Bases nearby	Test	1440	25	26.31	(0.65)	28	17	19.32	(1.18)	21	15	15.00	(0.00)	15

Table 10: The effects of narrowband 900 MHz cordless phones

LAN reception. We placed our WaveLAN transmitter and receiver approximately 20 feet apart in a large lecture hall and subjected them to various telephone interference. The results of this investigation are summarized in Table 10. In the "phones off" trial, the handsets and base units were turned off; in the "cluster" trial, both handsets and base units were activated and placed a few inches from the receiver's modem unit. These represent the extreme cases. For the other three trials, we moved either the handsets or the base units to an office across the hall from the lecture hall to investigate whether the phones were using some form of power control.[3] In the "handsets nearby talking" case, a human body was necessarily near the receiver, but in all other cases nobody was in the immediate vicinity of the WaveLAN units.

Except for the "cluster" trial, both handsets received nearly solid static, and the Panasonic unit beeped, presumably to indicate poor signal conditions. In sharp contrast, the WaveLAN experienced no damaged test packets, and only background levels of packet loss (the "phones off" and "handsets nearby talking" trials each experienced a single packet loss). The telephones affected the silence level to varying degrees; in the two cases where the silence level was lowest, we observed packets from stations not participating in the test (they were relatively weak and all were damaged in some way, so we suspect they were leakage from one or more nearby buildings which we know contain WaveLAN units). The fact that the highest silence level is in the "bases nearby" trial, when the handsets were distant, rather than the "cluster" trial suggests that the cordless phones may be using power control, perhaps to extend handset battery life.

We also briefly investigated the effects of an Advanced Mobile Phone Service (AMPS) narrowband FM cellular phone on WaveLAN reception. Once again, at varying distances, the WaveLAN seemed immune to bit errors. On the other hand, the cellular phone received significant amounts of white noise, totally overwhelming the audio signal, when it was close to an operating WaveLAN transmitter. We did not control for the effects of power control or channel selection by the cellular phone system.

WaveLAN's resistance to these interference sources is probably due to the DSSS modulation, which is known to be resistant to narrowband sources[28].

7.3 900 MHz spread spectrum cordless phone

We also investigated the effects of two 900 MHz spread spectrum cordless phones, an AT&T 9300 and a Radio Shack ET-909.[4] The two phones were quite similar in user interface and effects on WaveLAN reception. Unfortunately, if the handsets were within approximately three feet of each other, they could not simultaneously maintain connections to their respective base stations, even when multiple channels were employed. Our WaveLAN transmitter and receiver were approximately 25 feet apart in a conference room. The "near" location used for these trials was several inches from the receiver's modem unit, and the "far" location was approximately 14 feet from the receiver and 20 feet from the transmitter. For all runs we collected enough received packets to yield roughly 10^7 bits of packet body. Throughout the tests the telephones maintained acceptable audio signals. There were occasional clicks, but no actual outages, and the handsets did not complain about poor signal conditions.

The results of this investigation are summarized in Table 11. Three cases indicate that these phones can severely damage the WaveLAN environment: half of the packets are totally lost, while every packet that arrives is truncated. On the other hand, the "RS remote cluster" case indicates that reasonable separation between the WaveLAN and telephone leaves the link unharmed (though we will see below that the signal characteristics change noticeably). Finally, the "AT&T handset" case demonstrates that there is a significant intermediate effect: while a small number of packets are lost or truncated, nearly two thirds of the remainder contain correctable errors (the worst corruption of a packet body observed was 5% of the bits).

The signal information presented in Tables 12 and 13 suggests the following observations:

- The phones add a significant amount of noise to the environment. Every case except for "phones off" has a very high silence level.
- Very low signal quality seems to be a good predictor of truncation.
- If the signal level is high but signal quality is not outstanding, bit errors are likely.
- Based on signal information, in some trials some "outsiders" may be test packets corrupted beyond recognition.

In summary, ISM band spread spectrum cordless phones, depending on their location, can be harmless, severely disrupt a WaveLAN, or inject mild, plausibly correctable errors.

7.4 Competing WaveLAN units

The single-channel design of the WaveLAN suggests that non-cooperating WaveLAN units could provide significant mutual interference. Using the experimental layout described in Section 6.2, we placed additional WaveLAN transmitters at the Tx4 and Tx5 locations, and raised their receive threshold to 35, thus ensuring they would transmit continuously, and not defer to any nearby stations,

[3] Power control is the practice of reducing transmitter power as long as it does not induce errors; this saves battery life and allows frequency re-use.

[4] The Radio Shack ET-909 uses DSSS modulation, and we suspect the AT&T 9300 does as well.

Trial Name	Phone Type	Handset Location	Base Location	Packet Loss	Packets Truncated	Wrapper Damaged	Body Bits	Worst Body
Phones off	-	-	-	.5%	0%	0%	0%	-
RS base	RS	far	near	52%	100%	0%	0%	-
RS cluster	RS	near	near	51%	100%	0%	0%	-
AT&T cluster	AT&T	near	near	52%	100%	0%	0%	-
RS remote cluster	RS	far	far	0%	0%	0%	0%	-
AT&T handset	AT&T	near	far	1%	4%	1%	59%	4.9%

Table 11: Summary of spread spectrum cordless phones

Trial Name	Packet Source	Packets Received	Level				Silence				Quality			
			$\downarrow$	μ	σ	$\uparrow$	$\downarrow$	μ	σ	$\uparrow$	$\downarrow$	μ	σ	$\uparrow$
Phones off	Test	1389	28	29.63	(0.60)	31	0	2.20	(1.10)	7	15	15.00	(0.00)	15
	Outsiders	619	2	7.35	(5.54)	29	0	2.51	(1.17)	7	1	9.84	(4.41)	15
RS base	Test	1597	28	31.54	(2.57)	38	1	32.73	(5.44)	38	2	10.07	(3.29)	15
	Outsiders	316	4	31.03	(5.77)	37	2	27.07	(8.06)	37	1	7.10	(4.07)	15
RS cluster	Test	1488	28	32.01	(1.73)	37	0	30.73	(6.58)	37	1	8.91	(3.09)	15
	Outsiders	1818	3	32.57	(2.71)	37	0	28.97	(6.84)	37	1	5.48	(2.98)	15
AT&T cluster	Test	1766	29	32.52	(4.18)	41	1	38.96	(3.16)	41	1	7.45	(5.42)	15
	Outsiders	157	4	28.69	(8.95)	40	1	37.77	(7.13)	41	1	8.30	(5.43)	15
RS remote cluster	Test	1440	28	29.83	(0.65)	32	0	21.81	(5.91)	27	15	15.00	(0.00)	15
	Outsiders	9	3	4.22	(0.63)	5	21	24.33	(1.41)	26	7	8.11	(1.10)	11
AT&T handset	Test	1456	28	30.04	(1.04)	33	0	23.52	(6.57)	32	1	13.46	(2.14)	15
	Outsiders	265	4	31.31	(6.23)	33	3	23.29	(6.33)	32	1	6.15	(1.57)	12

Table 12: Signal measurements for spread spectrum phones

Packet Type	Packets Received	Level				Silence				Quality			
		$\downarrow$	μ	σ	$\uparrow$	$\downarrow$	μ	σ	$\uparrow$	$\downarrow$	μ	σ	$\uparrow$
All test	9136	28	31.01	(2.57)	41	0	25.78	(12.67)	41	1	11.44	(4.35)	15
Undamaged	3341	28	29.81	(0.72)	32	0	13.96	(10.97)	30	1	14.81	(0.77)	15
Truncated	4911	28	32.02	(3.09)	41	0	34.27	(6.35)	41	1	8.76	(4.28)	15
Wrapper damaged	21	29	30.62	(0.65)	32	3	24.71	(5.79)	29	9	12.24	(1.63)	15
Body damaged	863	28	29.89	(1.00)	33	0	23.26	(6.79)	30	3	13.62	(1.98)	15

Table 13: Signal breakdown for spread spectrum phone test packets

Trial Name	Packet Source	Packets Received	Level				Silence				Quality			
			$\downarrow$	μ	σ	$\uparrow$	$\downarrow$	μ	σ	$\uparrow$	$\downarrow$	μ	σ	$\uparrow$
Without interference	Test	12715	25	28.58	(0.60)	30	0	3.35	(1.11)	13	15	15.00	(0.00)	15
With interference	Test	12717	28	28.65	(0.49)	30	0	13.62	(3.39)	21	15	15.00	(0.00)	15
	Outsiders	31	24	28.65	(0.51)	32	0	13.62	(3.39)	28	15	15.00	(0.00)	15

Table 14: Signal metrics with and without interfering WaveLAN transmitters

including each other. We then attempted to transmit from location Tx1 to the receiver. Using the standard receive threshold value of 3, the link was completely unusable. For example, hundreds of invalid Ethernet addresses appear in the receiver's log, many associated with our test packets, indicating that the Ethernet station address field was frequently corrupted. Packet loss rates were high and collision-free transmissions very rare.

However, raising the receive threshold to 25, a value safely above the signal levels we had measured from the hostile transmitters' locations (Table 6), allowed the communicating stations to completely mask out the competition. We collected 10^8 bits with and without competition. In the case with competition, we experienced a .02% packet loss rate, which is not clearly significant, and no bit errors. If we compare the signal statistics (Table 14), we see that the background ("silence") level has increased significantly, but that the signal level and quality are essentially unchanged.

These results suggest that it is possible to use the receive threshold to shut out more distant systems without having high bit error rates when communicating with nearby systems. However, the fact that WaveLAN supports only a receive threshold and does not have transmit power control or multiple spreading sequences makes it difficult to create fully isolated pseudo-cells. In most environments, cells will be separated by "border zones" in which mobile clients will have poor performance and can easily disrupt communication in adjacent pseudo-cells. The reason is that hosts in the border zone can hear and be heard by hosts in multiple pseudo-cells, while the hosts in the different pseudo-cells cannot hear each other. This creates two types of problems. First, if a mobile host in the border zone communicates with a host in a cell, the carrier will be sensed in other cells, thus preventing communication in those other cells and reducing overall throughput. Second, if there is simultaneous communication in more than one cell (which is possible since the pseudo-cells are isolated), then a mobile host in the border zone may receive badly damaged packets. This is a special case of the classical "hidden transmitter" problem. We have observed, though not experimentally verified, that, when operated without thresholding, WaveLAN is fairly resistant to errors caused by hidden transmitters. We conjecture that this is because WaveLAN seems to be able to sense carrier even when it cannot receive complete packets, and because of a "capture effect" inherent in its multipath-resistant receiver design.

8 Implications for Wireless Network Architectures

While current wireless networks focus on connectivity rather than per-station bandwidth availability, our results suggest that future networks are likely to provide substantial bandwidth to individual client machines.

Our experiments indicate that "radio Ethernet" systems can provide good connectivity despite a wide variety of environmental hazards. The DSSS modulation and multipath resistant receiver design seemed to confer resistance to many environmental hazards. Another factor is the CSMA/CA Medium Access Control mechanism: within a collision domain, transmitters defer when they sense a carrier, so there should only be a single transmitter at any point in time. However, the shared nature of the channel also means that one should expect individual mobile machines to see only limited throughput. Furthermore, the fact that WaveLAN supports a receive threshold instead of transmit power control or multiple spreading sequences makes it difficult to create fully isolated cells.

In many cases, we observed a near-perfect link, arguing that FEC would be useless overhead in most situations. However, there were other situations, some plausibly predictable by signal measurements, in which there is frequent but minor packet corruption. Our observations, especially the spread spectrum phone results in Section 7.3, argue that the errors we did observe might be recoverable through a variable FEC mechanism.

The measurements we have presented provide reason to believe that it is feasible to construct an affordable wireless network that will provide bandwidth equivalent to what many wired computers currently enjoy, along with reasonable loss and error characteristics. For example, a WaveLAN-like device including multiple spreading sequences for sharp cell boundaries and transmitter power control to reduce unnecessary interference seems plausible, and would allow the construction of truly cellular network. The current WaveLAN design takes advantage of a particular spreading sequence with a very low self-correlation. While it is difficult to construct large sequence families which simultaneously have low self-correlation and low cross-correlation, and the effect of higher correlation would be more errors, the current WaveLAN seems to have processing gain to spare, especially since some robustness could be recovered via power control and adaptive multi-rate error coding[19, 31].

9 Related work

9.1 High Speed Radio LAN Error Environment

Duchamp and Reynolds have observed signal quality, throughput and error characteristics of an earlier model of 900MHz WaveLAN installed in ISA bus PCs, subjected to various challenges such as distance and multipath propagation[16]. Their testing regime included a propagation environment impeded by distance and local scatter induced by reflections from a wall. In this environment they observed packet loss and corruption rates both typically below 1%, except when a combination of attenuation and local scatter produced packet loss rates in the vicinity of 10% with a peak around 15% and packet corruption rates ranging as high as 40%. In the difficult environment, both rates varied nonmonotonically with distance, making it very unstable and unpredictable in the face of small motions. We observed similar loss and error rates due to attenuation and/or obstacles. Our work extends theirs by considering more error sources and by our investigation of the feasibility of pseudo-cellular WaveLAN operation.

Lewis and Guy measured the performance of the Arlan 610 system[26] and evaluated its utility for mobile multimedia applications. An interesting study [6] suggests it may be possible to obtain as much as 14 Mb/s from a single European DECT-style phone channel, albeit with a BER of 10^{-5}.

9.2 Olivetti Wireless ATM LAN

A joint project between Olivetti Research Labs and Cambridge University is constructing a wireless ATM LAN[29]. The architecture consists of 10 Mb/s dual-antenna radios operating in the 2.4GHz band, a ten-meter transmission radius leading to geographical re-use of a single frequency, and a CSMA-based MAC protocol, CSMA/AED[7]. The unit of transmission and acknowledgement is a standard ATM cell augmented with headers for wireless connection management.

9.3 Mobile IP Community

The IP networking community is investigating the effects on a reliable transport protocol connection of traversing a mixture of wireless and wired IP links, focusing on ameliorating the effects of small MTUs, high error rates, and handoff-related packet delays and losses[8, 13].

In response to these challenges, several groups have examined *indirect* approaches, in which a backbone host communicates with a mobile host via a transparent protocol translator[3, 4, 9, 5]. For example, in an indirect implementation of a transport protocol, a fixed host communicates with a mobile host via a proxy that transparently terminates the transport-level connection and uses another transport-level connection (and potentially another protocol) to reach the mobile host [4, 40]. This approach is able to improve throughput during handoff-related packet loss and reduce small-MTU inefficiency. Of course, the applicability of this approach to a particular network depends on its characteristics. Our initial experience suggests that there may be a class of high-performance wireless networks for which less aggressive approaches may suffice.

9.4 Adaptive Forward Error Correction

If channel conditions such as error patterns or bandwidth available for redundancy change over time, a single error correction algorithm may be inappropriate. While it is possible to use several different FEC codes, each with dedicated hardware resources, more attractive schemes exist. Hagenauer presents a family of codes called rate-compatible punctured convolution codes[19] which use the popular Viterbi decoding algorithm[38, 18, 34]. One example code family has 13 codes with redundancy overhead varying from 12.5% to 300%. Qualcomm, Inc., a manufacturer of digital radio systems, offers a single-chip encoder/decoder designed for satellite channels which operates at up to 25 Mb/s with four levels of redundancy [32]. Karn presents a variable-rate FEC system designed for implementation by general-purpose microprocessors[24] and evaluates its operation on a link subject to short but frequent interference from radar.

10 Conclusions

We used detailed packet tracing to investigate the effects of distance, obstacles, and different interference sources on the error and loss rates of WaveLAN, a 2 Mb/s wireless LAN designed specifically for an indoor fading environment.

Distance alone seemed to have little effect in a fairly large area and WaveLAN can frequently penetrate formidable obstacles, although it eventually succumbs to attenuation. In general we observed very few bit errors, and the receive threshold was very effective at shutting out distant sources while allowing reliable communication with nearby sources. The worst errors were induced by spread spectrum cordless phones operating in the same frequency band, and even these were substantially reduced by distance.

While WaveLAN proved resistant to several sources of radio interference, self-interference is substantial enough to impede building a robust cellular architecture. The reason is that WaveLAN lacks transmitter power control and multiple spreading sequences, which are needed to completely isolate adjacent cells. On the other hand, its isolation capabilities are probably sufficient for researchers wishing to experiment with cellular wireless architectures.

Overall, our measurements suggest that low-cost, low-error-rate, high-speed, shared-channel wireless networks may be commonplace in the near future. Further investigation may determine whether variable forward error correction is useful in such networks.

Acknowledgements

This work was partially supported by the Defense Advanced Research Projects Agency under contract number F19628-92-C-0116, and David Eckhardt was supported by an Intel Foundation Graduate Fellowship. Additional support was provided by Digital Equipment Corporation. Robert Baron of the Coda distributed filesystem group wrote the WaveLAN device driver which we extended to provide packet traces, members of the Coda group lent their time and laptops to provide interference, and obliging people lent me their cellular and cordless phones. This paper was significantly improved by the reviewers' comments.

References

[1] ATM Forum. *ATM User-Network Interface Specification: Version 3.0*, 1993.

[2] AT&T Wireless Communications and Networking Division. *WaveLAN Air Interface*, July 1995. Document 407-0024785, Revision 2 (draft).

[3] B. R. Badrinath, A. Bakre, T. Imielinski, and R. Marantz. Handling mobile clients: A case for indirect interaction. In *Proceedings of IEEE WWOS-IV*, pages 91–97, Napa, CA, October 1993. IEEE Press.

[4] A. Bakre and B. R. Badrinath. I-TCP: Indirect TCP for mobile hosts. In *Proceedings of the 15th International Conference on Distributed Computing Systems*, pages 136–143, May 1995.

[5] H. Balakrishnan, S. Seshan, and R. H. Katz. Improving reliable transport and handoff performance in cellular wireless networks. *Wireless Networks*, 1(4):469–481, 1996.

[6] M. Barton, J. McGeehan, A. Nix, and M. Lawton. Error rate prediction for high data rate short range systems. In *Wireless Personal Communications*, pages 251–279. Kluwer, 1993.

[7] S. K. Biswas, J. D. Porter, and A. Hopper. Performance of a multiple access protocol for an ATM based pico-cellular radio LAN. In *Proceedings of The Third IEEE International Symposium on Personal, Indoor and Mobile Radio Communications*, pages 139–144. IEEE, 1992.

[8] R. Cáceres and L. Iftode. The effects of mobility on reliable transport protocols. In *Proceedings of the 14th International Conference on Distributed Computing Systems*, pages 12–20, June 1994.

[9] R. Cáceres and L. Iftode. Improving the performance of reliable transport protocols in mobile computing environments. *IEEE Journal on Selected Areas in Communications*, 13(5):850–857, June 1995.

[10] CCITT. Proposed recommendation I.311, B-ISDN general network aspects, June 1991.

[11] D. D. Clark, S. Shenker, and L. Zhang. Supporting real-time applications in an integrated services packet network: Architecture and mechanism. In *Proceedings of ACM SIGCOMM '92*, pages 14–26. ACM SIGCOMM, August 1992.

[12] J. Crowcroft, Z. Wang, A. Smith, and J. Adams. A rough comparison of the IETF and ATM service models. *IEEE Network*, 9(6):12–16, Nov/Dec 1995.

[13] A. DeSimone, M. C. Chuah, and O.-C. Yue. Throughput performance of transport-layer protocols over wireless LANs. In *Proceedings of IEEE GLOBECOM 1993*, pages 542–549, December 1993.

[14] P. Dryden. Infrared connections get red-hot. *Communications Week*, page 1, November 21 1994.

[15] D. Duchamp and A. Athan. Agent-mediated message passing for constrained environments. In *Proceedings of the USENIX Mobile and Location-Independent Computing Symposium*, pages 103–107. USENIX Association, August 1993.

[16] D. Duchamp and N. F. Reynolds. Measured performance of a wireless LAN. In *Proceedings of the 17th Conference on Local Computer Networks*, pages 494–499. IEEE, Sep 1992.

[17] D. Ferrari, A. Banerjea, and H. Zhang. Network support for multimedia–a discussion of the TENET approach. *Computer Networks and ISDN Systems*, pages 1267–1280, July 1994.

[18] G. Forney. The Viterbi algorithm. *Proceedings of the IEEE*, 61(3):268–278, March 1973.

[19] J. Hagenauer. Rate-compatible punctured convolutional codes (RCPC codes) and their applications. *IEEE Transactions On Communications*, 36(4):389–400, April 1988.

[20] K. Hamied and G. L. Stuber. A non-iterative algorithm for estimating the impulse response of ISI channels. In *Wireless Personal Communications*, pages 175–186. Kluwer, 1993.

[21] A. Hills and D. B. Johnson. A wireless data network infrastructure at Carnegie Mellon University. *IEEE Personal Communications*, pages 56–63, January 1996.

[22] S. Kallel. Efficient hybrid ARQ protocols with adaptive forward error correction. *IEEE Transactions on Communications*, 42(2/3/4):281–289, February/March/April 1994.

[23] P. Karn. The Qualcomm CDMA digital cellular system. In *Proceedings of the USENIX Mobile and Location-Independent Computing Symposium*, pages 35–39. USENIX Association, August 1993.

[24] P. Karn. Toward new link-layer protocols. *QEX*, pages 3–10, June 1994.

[25] K. Keeton, B. A. Mah, S. Seshan, R. H. Katz, and D. Ferrari. Providing connection-oriented network services to mobile hosts. In *Proceedings of the USENIX Mobile and Location-Independent Computing Symposium*, pages 83–102. Usenix Association, 1993.

[26] A. Lewis and C. Guy. Assessing the radio performance of wireless LANs for mobile multimedia applications. In *Proceedings of the 2nd International Workshop on Mobile Multimedia Communications*, pages A5/6/1–6. Hewlett-Packard Laboratories, April 1995.

[27] S. Loudermilk. Hardware extends wireless LAN mileage. *LAN Times*, January 1995.

[28] R. L. Pickholtz, D. L. Schilling, and L. B. Milstein. Theory of spread spectrum communications–a tutorial. *IEEE Transactions On Communications*, 30(5):855–884, May 1982.

[29] J. Porter and A. Hopper. An ATM based protocol for wireless LANs. Technical Report ORL–94–2, Olivetti Research Limited, Cambridge, CB2 1QA, England, April 1994.

[30] Qualcomm, Inc. Proposed EIA/TIA interim standard: Wideband spread spectrum digital cellular system—dual-mode mobile station–base station compatibility standard. Technical Report 80-7814 Rev DCR 03567, Qualcomm Inc., San Diego, California, April 1992. Accepted as TIA IS-95; available as ftp://lorien.qualcomm.com/pub/cdma.

[31] Qualcomm, Inc. Q1650 k=7 multi-code rate Viterbi decoder technical data sheet. 6455 Lusk Boulevard, San Diego, CA 92121-2779, 1992.

[32] Qualcomm, Inc. Master selection guide. 6455 Lusk Boulevard, San Diego, CA 92121-2779, 1994.

[33] K. Raith, E. Lissakers, J. Uddenfeldt, and J. Swerup. Cellular for personal communications. In *Wireless Personal Communications*, pages 1–20. Kluwer, 1993.

[34] M. Ryan and G. Nudd. The Viterbi algorithm. Technical Report RR–238, Department of Computer Science, University of Warwick, Coventry, CV4 7AL, United Kingdom, February 1993.

[35] S. Shenker and C. Partridge. Specification of guaranteed quality of service, July 1995. Internet Draft, IETF Integrated Services WG.

[36] F. A. Tobagi, R. Binder, and B. Leiner. Packet radio and satellite networks. *IEEE Communications*, pages 24–40, November 1984.

[37] B. Tuch. Development of WaveLAN, an ISM band wireless LAN. *AT&T Technical Journal*, pages 27–37, July/August 1993.

[38] A. J. Viterbi. Error bounds for convolutional codes and an asymptotically optimum decoding algorithm. *IEEE Transactions on Information Theory*, 13(2), April 1967.

[39] M. Weiser. Some computer science issues in ubiquitous computing. *Communictions of the ACM*, 36(7):75–84, July 1993.

[40] R. Yavatkar and N. Bhagawat. Improving end-to-end performance of TCP over mobile internetworks. Technical report, Department of Computer Science, University of Kentucky, Lexington, KY 40506, 1994.

[41] L. Zhang, S. Deering, D. Estrin, S. Shenker, and Z. Zappala. RSVP: a new resource ReSerVation Protocol. *IEEE Network*, pages 8–18, September 1993.