Measured Performance of a Wireless LAN

Dan Duchamp and Neil F. Reynolds

Computer Science Department
Columbia University
500 W. 120th St.
New York, NY 10027
{djd,nfr}@cs.columbia.edu

Abstract

We have studied the performance of a high-speed commercial spread-spectrum wireless LAN that uses the CSMA/CA multiple-access strategy. Employing synthetic workloads, we measured packet capture success more so than signal propagation characteristics. Specifically, we measured throughput, packet loss rates, range, and patterns of errors within packets. We conclude that CSMA/CA is quite successful in allocating bandwidth under stress, but that packet capture rate degrades very quickly once the LAN's effective range is exceeded. Hence, network maintainers should plan the layout of wireless networks at least as carefully as they plan wired networks.

1 Introduction

Thanks to regulatory decisions and improvements in technology, the last several years have seen an explosion of developments in the commercial wireless sector of the telecommunications industry. So far the emphasis has been on low-speed, voice-oriented telephony applications.

Another promising application is data movement over a wireless LAN. Although the concept has been popular for only the past few years, already several companies have introduced products[1], thereby suggesting substantial market possibilities. The major touted advantages of a wireless LAN are:

- Provision of LAN service to hard-to-wire areas.
- Provision of continuous LAN service to mobile users who carry portable computers.
- The possibility that wireless LANs might provide lower "cost per seat," thus making them attractive for certain applications even in areas that are or can easily be wired.

If the wireless LAN is to replace or complement the traditional wired LAN, a significant question is how and to what extent wireless technologies handle packetized data traffic differently from wireline technologies. Issues such as bandwidth, effective throughput, security, range, and reliability will make or break the concept of a wireless LAN.

The question of signal propagation has received study for decades, and a great deal is known about propagation characteristics for broad ranges of values of many parameters, such as: frequency, modulation technique, multiple-access strategy, antenna design and placement, and fading environment [2, 6, 7, 5, 10]. From these studies, certain reasonably reliable predictions can be made about how a wireless LAN should behave. For example, knowledge of the fading environment, transmission power, frequency, and modulation technique should allow one to predict the effective range of a wireless LAN.

However, it is not always obvious how data about signal propagation translates into LAN behavior, and — during our two years of experience with wireless LANs — we have often seen behavior contrary to the expectations we formed based on our knowledge of signal propagation characteristics and implementation of the modem circuitry.

Accordingly, we have undertaken a "packet level" study of the behavior of wireless LAN technology. The purpose of the study is not to add to the substantial literature on signal propagation, but rather to determine the effectiveness of a wireless LAN. Studies of the suitability of wireless transmission for telephony do not necessarily apply to a LAN, for several reasons:

- Most, but not all, of these studies concentrate on outdoor, rather than indoor, signal propagation. The wide assortment of building materials used indoors, the greater variation in transmitter and receiver placement relative to obstructions, and the presence of people moving about all make the indoor environment less predictable and thus more challenging.

  We benefit from the fact that the signal propagation characteristics of our environment have al-

[1] These include: RF products (Motorola's Altair [4], NCR's WaveLAN [12], Proxim's RangeLAN and Telesystems' Arlan) as well as IR products from Spectrix and Phonotics.

ready been studied [8].

- A LAN transmits, rejects, and resends data in the form of whole packets.
- LAN traffic is burstier, and there is no "pause time" that can be used for control information.

Section 2 explains the experimental setting of our study, while Section 3 reveals and interprets our measurements. We conclude in Section 4.

2 Experimental Setting

2.1 Radio Modem

We used the NCR WaveLAN modem [12]. This device physically broadcasts every packet, with address filtering done in hardware following the Ethernet scheme (i.e., 48-bit station addresses). Transmission is via direct-sequence spread-spectrum transmission in the lowest of the three Industrial, Scientific, and Medical (ISM) bands, 902-928 MHz. The FCC limits license-free transmission power to one watt, and requires a processing gain of 10 dB.

In direct-sequence spread-spectrum transmission, an information signal of a certain bandwidth is spread over a larger transmission bandwidth. WaveLAN spreads 2 Mbits/sec of information over 22 MHz (904-926 MHz). WaveLAN uses 11 *chips* per baud[2], while transmitting two bits per baud, at a rate of 1Mbaud. With 11 chips per microsecond, the WaveLAN has a *chip rate* of 11MHz and a chip period of 91ns. The many advantages of spread-spectrum transmission include superior resistance to *multipath* fading and minimal interference between the spread-spectrum signal and other transmitters occupying the same frequency range.

The WaveLAN is equipped with two antennas, with hardware dynamically choosing to receive from the one exhibiting the highest signal quality. Such "antenna diversity" is effective in combating multipath, more so as the frequency increases [11].

The multiple-access strategy is CSMA/CA, CA standing for "collision avoidance." A potential sender listens and does not send until the channel is free. Transmission may continue during a collision, because the collision may not be detected. CSMA/CA is a reasonable multiple-access strategy for a wireless LAN. Since the limited range makes "end-to-end" signal propagation time short, repeated collisions will happen only if the number of senders is large.

2.2 Signal Propagation

Multipath signal propagation leads to *delay spread* wherein some reflected signals arrive later than others. The effect is that the signal auto-correlates with a delayed version of itself, complicating the task of capturing and re-constructing a true signal. In direct-sequence spread-spectrum transmission, the algorithm for spreading a certain information bandwidth over a wider carrier bandwidth is parameterized by a *spreading code*. The extent of degradation caused by multipath is highly dependent on the particular spreading code. Many spreading codes are known to have excellent auto-correlation properties [13], and the WaveLAN uses one such code.

Nevertheless, delay spread remains harmful. Long delay spreads — on the order of the amount of time necessary to transmit a bit of information — are extremely harmful, because one spread bit will auto-correlate with the next spread bit. If the auto-correlation of spread bits is exact enough, the bits become indistinguishable and information is lost [11].

Our test environment has already been measured for delay spread and for how power drops off with distance [8]. The median delay spread in our environment at our frequency is 40ns; 75% of measured delay spreads were less than 50ns. Such delay spreads should be easily dealt with by the WaveLAN, as they are less than its chip time of 91ns, and only a small fraction of the baud time of $11 \times 91\text{ns} = 1\mu\text{s}$.

2.3 Hardware and Operating System

Transmitting computers were equipped with the Intel i486 processor and 16MB of memory. Receivers had the Intel i386 and 8MB of memory. Both types of computers had an "industry standard" (ISA) bus. Such a bus runs at 8 MHz, and is capable of transferring a byte every five cycles, yielding a maximum of 1.6 MB/sec between memory and I/O devices. All computers were running version 2.5 of the Mach operating system [1], a descendant of UNIX.

Packet transmission involves copying the packet from main memory into the WaveLAN's send buffer, where it resides until it is sent. The reverse path is taken when receiving a packet. The WaveLAN's on-board buffer space for all purposes — both receiving and holding packets awaiting transmission — is 64KB.

In order to quickly analyze packets and to inspect damaged packets, we altered Mach's WaveLAN device driver. The altered device driver recorded the following information in kernel buffers:

- For every packet, the sending address, the packet length, and the "signal quality" indicator returned by the WaveLAN hardware.
- For damaged packets, the cause of error reported by hardware and the packet itself.

At the end of test runs, we read this information out of the operating system via `/dev/kmem`. According to [12], the *signal quality* is defined as "a measure of an averaged correlation peak to side lobe ratio." In order words, it is a number (in dB) indicating the difference between the main radiation lobe and the primary side lobe. The larger this number, the clearer the signal. Signal quality is measured per-packet, with the measurement is being taken shortly after reception begins.

[2] This definition of a "chip" is from the jargon of spread-spectrum technology, and referrs to a partial bit

3 Measurements

In addition to obvious concerns such as the range and signal quality of a wireless LAN, we were interested in other "packet level" issues such as:

1. the effect of packet size on loss rate and throughput
2. what types of errors occurred
3. the effectiveness of the CSMA/CA multiple-access strategy in coping with load generated by multiple transmitters

Accordingly, we carried out the experiments described below.

All measurements were conducted along the "long axis" of the Columbia Computer Science department — a hallway 56 meters long, terminated by walls at either end. We placed transmitters and receivers at various points along the hallway. The hallway acted as a waveguide, and presumably enhanced the range of the LAN.

A sending station ran a user-level process that remained in a tight loop, first sending a packet of some size and of predetermined contents, then pausing for a while. The size of the packet and the length of the pause were parameters. All actions at a receiving station took place entirely within the operating system; at the end of a test, the information accumulated by the receiver was drained from the kernel, but during a test no processes ran on the receiver.

Although the senders had considerably faster CPUs than the receivers, both types of machines had the same bus and hence the same I/O bandwidth. This fact coupled with the fact that a receiver never had to context switch meant the receiver was faster than the sender.

3.1 Baseline

The baseline measurement consisted of a single sender transmitting 10,000 packets to a single receiver, with packets spaced every 10ms. Packet sizes were 100, 500, 1000, and 1500 bytes; 1500 is the maximum size. The distance was also varied: from 10 meters to 56 meters, in gradations of 10 meters. The results are presented in Table 1. We use these results only for comparison later on.

In Table 1, for each combination of distance and packet size, four results are given. The two percentages are, respectively, the percentage of transmitted packets that were received and the percentage of received packets that were damaged. The two numbers in parentheses are the signal quality of good and bad packets, respectively.

Several facts are apparent from this table. First, almost all packets were received at distances between 10 and 50 meters. At 56 meters, adjacent to a wall, the receipt rate drops noticeably. At all distances, the signal quality recorded for damaged packets is significantly lower than for good packets. Finally, signal quality is roughly correlated with distance, with a large drop apparent between 20 and 40 meters. All of the observations held true over all other tests we performed.

3.2 Range

The results in Table 1 indicate a sharp dropoff in capture success between 50 and 56 meters, where the hallway terminates in a wall. To better understand the effective range of the LAN as well as the effect of the wall, we performed two more tests. In the first, we varied the distance between 50 and 55 meters, a meter at a time. The second test was identical to the first except that the transmitter was moved 25 meters down the hallway; the receiver remained near the wall. We wanted to determine what was responsible for the poor capture: reflection from the wall, or the distance, or both. The results are summarized in Tables 2 and 3.

Table 2 indicates the instability of the signal in the 50-meter range: packet capture results change significantly from meter to meter, with closer not necessarily being better. At this range, the signal is sufficiently attenuated that multipath effects — not just the delay spread mentioned above but also *local scattering* — weaken the signal enough so that it cannot be reliably captured and de-spread. Local scattering refers to reflections off objects located near the receiver. The effect of local scattering is to superpose a collection of signals with random, hence cancelling, phase components.

The results achieved by moving the transmitter closer while keeping the receiver near the wall show that the wall is a problem only when the line-of-sight signal is weak.

3.3 Throughput

Further tests showed the CSMA/CA multiple-access strategy quite effective in delivering a substantial fraction of the rated bandwidth.

A throughput measurement employing a single sender-receiver pair moved 10,000 1500-byte packets in 104 seconds, working out to 1.15 Mb/sec. When the number of sender-receiver pairs is increased, the packet size must be reduced or greater delay added between packets, since two pairs operating at 1.2 Mb/sec substantially exceed the link capacity.

Two pairs moved 10,000 1150-byte packets in 117 sec, yielding a throughput of 1.57 Mb/sec. Three pairs moved 10,000 1000-byte packets in 140 sec, providing a throughput of 1.71 Mb/sec. As with the previous tests, well over 99% of packets were received without error within the effective range of the LAN (under 30 meters).

In these throughput tests, the maximum packet size was determined empirically. As we increased the number of sender-receiver pairs, we started with the packet size used in the previous test. The result was that one or more of the senders would abort the test when the onboard buffer of its WaveLAN overflowed. This occurs because, when the channel is busy, the collision avoidance circuitry causes packets to queue for later sending.

Eventually the WaveLAN's 64KB buffer fills, as packets are queued faster than they are sent.

"Eventually" can be almost immediately, suggesting that the timing can work out such that one sender can be starved of channel access. Furthermore, occurrence of this overflow event is not easily predictable — it is not the case that the time to overflow has a linear relationship with packet size, as might be supposed.

3.4 Error Characterization

When a packet is transmitted, there are three possible outcomes:

1. The receiver captures the packet without error.
2. The receiver captures the packet, but with an error.
3. The receiver fails to capture the packet.

During our tests the transmitters and receivers recorded some data about packets in the latter two categories.

Transmitters sent packets whose (known) contents contained an approximately equal number of 0's and 1's, smoothly distributed throughout the packet. When the hardware signaled an error, our altered device driver retained the packet. The damaged packets were processed after the test to determine the number and types of errors. Every damaged packet experienced an "alignment error," which for 802.3 networks means that the packet is not a whole number of bytes. Since months of experiments saw many thousands of erroneous packets, such uniformity is suspicious — the 802.3 standard defines a few other types of errors; for example, failure to pass the CRC check. We cannot know whether the hardware is designed to always report alignment error or whether in fact only alignment errors do occur. However, on occasion the hardware reports an alignment error for an undamaged packet, which leads us to believe that the hardware is not very careful about discriminating among types of errors, and that it is possible that any type of error is reported as an alignment error.

So as to detect possible burst errors as well as single-bit errors, we define an error as a "run of bad bytes." That is, we compare a damaged packet byte-by-byte with what it should be. Whenever consecutive bytes are found to be damaged, the run is defined as starting with the first bad bit in the first bad byte and ending with the last bad bit in the last bad byte. For example, a run of length 76 consists of some bad bits in each of 10 or 11 consecutive bytes, beginning and ending with a bad bit, and containing 74 bits in between which may be good or bad. A similar definition applies when several bad bits appear within a single isolated bad byte.

Our post-processing collected the following information: histogram of run length, histogram of the number of separate runs per packet, and the number of incorrect bits in each run. A few clear patterns emerge from the data:

- Errors tend to be non consecutive bit errors. Computation of the number of bad bits per run yielded, in almost all cases, exactly the minimum number of bad bits necessary to define a run.
- Multi-byte runs are not uncommon, and occasional very long runs do occur. The longest we saw was 848 bits, meaning that 106 consecutive bytes each had at least one bad bit. Nevertheless, the number of runs per packet drops off steeply: by far the most common number of runs in a damaged packet is one.
- Certain run lengths are strongly preferred at all distances and packet sizes. In particular, in some tests close to 90% of the runs are of length 13 or 14. This clustering is more pronounced in tests that suffer the fewest errors.

  Another pattern that is evident is that, for runs of length 16 and over, the number of runs rises steadily up to a multiple of 8, then drops precipitously. The effect is that of a sawtooth pattern with peaks at multiples of 8. For example, runs of length 48 are much more common than runs of length 49.
- The number of errors per bit (or the "error density") is roughly constant over all packet sizes and distances.

4 Summary and Conclusions

We measured a single wireless LAN in an environment that has been previously found not to present challenging multipath fading characteristics. We found that the LAN is very reliable within a certain range, and when under stress can deliver a substantial fraction of the rated bandwidth.

In our environment, the reliable range of WaveLAN is about 40 meters. Beyond this limit, its behavior becomes unstable: the relative placement of sender and receiver becomes extremely important, with movements of even a single meter in a unexpected direction making an enormous difference in the success of packet capture. Instability outside of the preferred operating regime is not unique to wireless LANs; technologies such as Ethernet behave similarly [3]. When such unstable conditions do occur, it appears that packet length is a key parameter, with the smallest packets being captured much more successfully than the largest packets.

The signal quality measure returned by the WaveLAN modem exhibits a substantial dynamic range, and is roughly correlated with distance throughout and beyond the WaveLAN's range. This fact plus the fact that most small packets can be captured even during periods of unstability, suggests that dynamic cell-to-cell handoff — such as that described in [9], which requires exchange of a few small control packets — can be done efficiently with this sort of hardware. This conclusion is significant for wireless mobile computing.

When packets are damaged (but still possible to capture), they tend to suffer from a small number of isolated bit errors. The fact that run length distribution fits a very clear pattern suggests that errors result from a specific underlying physical process, which presumably can be masked with proper hardware design; the scope of our study does not permit identification of the underlying physical process. The complete absence of long bursts dense with bad bits suggests that transmitter and receiver do successfully maintain synchronization, even in circumstances where the signal is weak and capture is difficult. (If there were such a phenomenon as a "lost bit," then all the bits that follow would be detected as bad.)

In the near future we hope to carry out similar measurements on a few other wireless LAN technologies (including IR technologies and RF technologies employing CDMA) and to examine the behavior of several wireless LANs operating in the same time and space.

5 Acknowledgements

This work was supported in part by National Science Foundation grants CDA-90-22123, NCR-91-04905, IBM Corporation, the New York State Science and Technology Foundation Center for Advanced Technology in Computer and Information Sciences, and the Center for Telecommunications Research, an NSF Engineering Research Center funded by grant ECD-88-11111.

References

[1] M. Accetta et al. Mach: A New Kernel Foundation for UNIX Development. In *Proc. of Summer Usenix*, pages 93–112, July 1986.

[2] S. E. Alexander. Characterizing Buildings for Propagation at 900 MHz. *Electronics Letters*, 19(20), September 1983.

[3] D. R. Boggs, J. C. Mogul, and C. A. Kent. Measured Capacity of an Ethernet: Myths and Reality. In *Proc. SIGCOMM '88*, pages 222–234. ACM, August 1988.

[4] D. Buchholz, P. Odlyzko, M. Taylor, and R. White. Wireless In-Building Network Architecture and Protocols. *IEEE Network Magazine*, 5(6):31–38, November 1991.

[5] D. M. J. Devasirvatham, R. R. Murrayand, and C. Banerjee. Time Delay Spread Measurements at 850 MHz and 1.7 GHz Inside a Metropolitan Office Building. *Electronics Letters*, 25(3), February 1989.

[6] D. M. J. Devasirvatham. Time Delay Spread Spectrum Measurements of Wideband Radio Signals Within a Building. *Electronics Letters*, 20, November 1984.

[7] D. M. J. Devasirvatham. A Comparison of Time Delay Spread Spectrum Measurements Within Two Dissimilar Office Buildings. In *Proc. of the 1986 Intl. Conf. on Communications*. IEEE, June 1986.

[8] C. Huang and R. Khayata. Wideband Indoor Radio Propagation Measurements in the ISM Bands. In *Proc. of the 1992 Intl. Conf. on Communications*. IEEE, June 1992.

[9] J. Ioannidis, D. Duchamp, and G. Q. Maguire Jr. IP-based Protocols for Mobile Internetworking. In *Proc. SIGCOMM '91*, pages 235–245. ACM, September 1991.

[10] M. Kavehrad and G. E. Bodeep. Design and Experimental Results for a Direct Sequence Spread-Spectrum Radio Using Differential Phase-Shift Keying Modulation for Indoor, Wireless Communication. *IEEE Journ. Selected Areas in Communications*, SAC-5(5), June 1987.

[11] J. E. Mitzlaff. Radio Propagation and Anti-Multipath Techniques in the WIN Environment. *IEEE Network Magazine*, 5(6):21–26, November 1991.

[12] B. Tuch. An Engineers Story of WaveLAN. In *First Virginia Tech Symp. on Wireless Personal Communications*. IEEE, June 1991.

[13] N. Zhang and S. W. Golomb. Sixty-Phase Generalized Barker Sequences. *IEEE Trans. Information Theory*, 35(4), July 1989.

DIST/SIZE	100	500	1000	1500
10	99.87% 0.07% (21.94) (19.43)	99.96% 0.08% (20.51) (24.75)	99.95% 0.01% (21.10) (24.00)	99.96% 0 (24.69) (–)
15	99.28% 0 (22.09) (–)	99.53% 0 (23.28) (–)	99.39% 0 (23.11) (–)	99.29% 0.01% (22.41) (23.00)
20	99.60% 0.02% (19.34) (16.75)	99.88% 0.03% (22.98) (15.33)	99.65% 0.17% (21.55) (18.00)	99.85% 0 (23.27) (–)
30	99.88% 0 (14.96) (–)	99.88% 0.03% (15.02) (9.33)	99.87% 0.09% (14.49) (11.00)	99.93% 0.21% (15.38) (11.10)
40	99.82% 0.01% (9.55) (12.00)	99.76% 0.17% (9.82) (8.24)	99.75% 0.18% (10.04) (9.78)	99.77% 0.25% (10.21) (8.20)
50	99.71% 0.41% (10.91) (6.90)	99.76% 0.56% (11.39) (8.64)	99.71% 1.21% (11.14) (7.75)	99.74% 1.82% (10.99) (7.16)
56	82.99% 13.77% (8.31) (5.23)	88.83% 22.36% (7.99) (5.66)	77.39% 31.14% (7.72) (6.16)	68.78% 48.87% (7.35) (5.72)

Table 1: Baseline Results Between 10 and 56 Meters

DIST/SIZE	100	500	1000	1500
50	99.87% 0.30% (7.89) (7.83)	99.72% 2.92% (7.87) (6.09)	99.91% 0.27% (7.81) (6.63)	99.77% 99.52% (7.87) (6.48)
51	87.48% 12.49% (7.83) (4.40)	90.89% 6.15% (8.70) (4.01)	85.76% 13.18% (8.55) (4.87)	87.76% 22.99% (8.26) (4.21)
52	99.66% 0.6% (8.74) (5.65)	99.74% 0.87% (9.15) (6.18)	99.89% 0.96% (9.37) (6.47)	99.52% 4.33% (8.97) (5.88)
53	99.69% 0.71% (10.49) (5.01)	99.55% 2.31% (10.72) (5.94)	99.63% 3.91 % (10.73) (5.93)	99.68% 5.83% (11.10) (5.83)
54	99.84% 0.49% (9.64) (5.65)	99.79% 0.12% (10.22) (6.75)	99.51% 0.98% (9.92) (5.73)	99.29% 5.70% (9.37) (5.91)
55	91.96% 11.91% (7.46) (4.34)	95.15% 17.54% (7.40) (6.02)	97.63% 18.77% (6.92) (6.69)	96.83% 40.73% (6.74) (6.52)

Table 2: Results Between 50 and 55 Meters

DIST/SIZE	100	500	1000	1500
25	99.82% 0.03% (20.17) (6.67)	99.93% 0.10% (19.73) (7.70)	99.77% 0.03% (20.01) (23.33)	99.88% 0.06% (19.57) (19.50)
27	99.89% 0.12% (23.85) (3.92)	99.73% 0.07% (23.70) (4.29)	99.72 0.06% (23.34) (23.83)	99.76% 0.03% (23.56) (24.67)
29	99.91% 0.34% (20.18) (6.06)	99.91% 0.16% (20.31) (7.06)	99.81% 0.11% (19.50) (10.64)	99.86% 0.14% 918.38) (15.14)
30	99.77% 0 (21.26) (–)	99.86% 0.06% (20.88) (6.50)	99.80% 0.11% (20.79) (5.64)	99.73% 0.09% (21.06) (20.22)

Table 3: Transmitter Moved 25 Meters Closer