Performance of IEEE 802.11b Wireless LAN in an Emulated Mobile Channel

Christopher Steger, Predrag Radosavljevic and J. Patrick Frantz
{cbs, rpredrag, jpfrantz}@rice.edu
Department of Electrical and Computer Engineering
Rice University – Houston, Texas 77005

Abstract— **The performance of 802.11b wireless LANs is well understood for indoor environments. However, their behavior in outdoor and mobile environments has remained largely unexplored. We have performed experiments to bridge this knowledge gap by empirically measuring the performance of an 802.11b system in a broad spectrum of emulated environments. The goal of our work is to contribute to the evolution of currently existing wireless standards by measuring the effects of different channel phenomena on 802.11b. Our results show that current implementations of the 802.11b standard may not be well suited to use in a mobile environment, but that they could conceivably be modified to have better performance in such situations.**

I. INTRODUCTION

Systems designed to operate in indoor environments are often not particularly robust against the propagation effects of mobility. As the demand for broadband mobile data service increases, wireless products must adapt and evolve to offer services that were previously restricted to indoor systems. For instance, the emerging field of telematics explores the implementation of high-bandwidth wireless networking for moving vehicles [1–3].

We tested the performance of 802.11b while varying the delay spread, velocity, and SNR of an emulated channel. Our results demonstrate the challenges that must be surmounted before an 802.11b-like system is ready to be part of a more versatile, ubiquitous network topology. Earlier work has concentrated on the performance of 802.11b with stationary nodes in an indoor environment [4]. Others have experimented with MAC (media access control) adjustments to the protocol for more hostile channel conditions [5, 6]. In contrast, we tested 802.11b links in a broad spectrum of channel conditions ranging from indoor to outdoor and mobile scenarios.

Many have studied the propagation characteristics of the indoor 2.4 GHz channel [7, 8], for which current 802.11b implementations are designed. Measurements of the outdoor 2.4 GHz channel show that channel parameters differ significantly from the indoor environment [9]. We used these studies to find a reasonable range of channel parameters, and we varied those parameters to precisely measure their relationships to system performance. In order to produce realistic, repeatable results, we used a channel emulator. Our methods are detailed in Section II, and results and analysis are featured in Section III followed by conclusions in Section IV.

II. EXPERIMENT DESCRIPTION

A. Equipment

To gauge the performance of 802.11b in more than one implementation, we constructed three different peer-to-peer links. The first consisted of two laptops equipped with Orinoco WLAN cards. In the second system we connected two PCs with D-Link WLAN cards. The third system was a combination of the first two featuring one D-Link card and one Orinoco card. One reason we chose to use these two manufacturers was to be able to observe the differences in baseband implementation by two leading chipset manufacturers, Agere and Intersil. In each case, we linked them through a channel emulator and measured the throughput and packet loss rate for a UDP data stream while systematically varying the channel parameters. Specifically, we used the following hardware and software:

- Two Orinoco PC Cards with antennas removed for isolation
- Two D-Link Air DWL-520 wireless PCI cards
- A TAS 4500 FLEX RF Channel Emulator from Spirent Comm.
- IPERF version 1.1.1 network testing software

Two of the challenges in constructing our experiments were generating the channel and connecting the WLAN cards to the emulator. We chose to use a channel emulator because it allowed us to vary individual channel parameters exactly, independently, and repeatably.

In order to use the channel emulator, we modified the WLAN cards to communicate over a wired link rather than a wireless one. It was important to insure that no RF leakage bypassed our channel, so sufficient isolation was critical. We chose our WLAN cards for their adaptability to wired operation. The D-Link card has a removable antenna, but the Orinoco cards needed to have their built-in antennas unsoldered in addition to careful shielding.

B. Methods

1) Velocity: In our velocity measurements, we combined the effects of fast fading according to either Rician or Rayleigh distributions with a Doppler shift. In general, propagation situations in which there is no line-of-sight path and all received signal power is reflected off of local scatterers are governed by a Rayleigh distribution [11].[1]

---

[1] After conducting several experiments using both fading distributions, we observed that the effects on throughput and packet loss rate were virtually identical for both Rician and Rayleigh fading (see Figures 1 and 2). Therefore, for our later tests we chose to restrict ourselves to Rayleigh fading for the sake of simplicity.

TABLE I

CHANNEL PARAMETERS FOR THE TAS 4500 FLEX CHANNEL EMULATOR

RMS Delay Spread	Path 1 Gain (0 dB)	Path 2 Gain (-4 dB)	Path 3 Gain (-8 dB)	Path 4 Gain (-12 dB)	Path 5 Gain (-16 dB)	Path 6 Gain (-20 dB)
0 ns	0.000 ns	0.000	0.000	0.000	0.000	0.000
10 ns	0.000 ns	10.167	20.333	30.500	40.667	50.833
20 ns	0.000 ns	20.500	41.000	61.500	82.000	102.500
30 ns	0.000 ns	30.667	61.333	92.000	122.667	153.333
40 ns	0.000 ns	41.000	82.000	123.000	164.000	205.000
50 ns	0.000 ns	51.167	102.333	153.500	204.667	255.833

2) Delay Spread: We created an ad hoc multipath channel model to study the effects of delay spread. The model consists of 6 discrete paths with fixed gains and delays calculated to achieve specific values of RMS delay spread. We chose the gains to be similar to measured power delay profiles observed in literature [10]. We constrained the algorithm to fixed gains and uniform spacing between paths and used numerical methods to vary the path spacing until the desired RMS delay spread was achieved. The use of a 6 path model was dictated by the limit on the number of paths accommodated by the channel emulator and our desire to isolate the effect of delay spread as much as possible. We do not claim that our model specifically represents a particular channel, only that it is an effective and qualitatively reasonable way to generate a channel for different values of RMS delay spread. The path values for the channel emulator are shown in Table I.

It is important to note the distinction between RMS delay spread and maximum excess delay which is also often referred to as delay spread. RMS delay spread deals with the distribution of arriving power according to Equations 1 and 2 [11], while maximum excess delay spread consists of the difference between the arrival times of the first and last multipath components. It is often more revealing to study RMS delay spread because of the fact that it relies on the relative amplitudes of multipath components rather than absolute power levels [11]. Maximum excess delay depends on absolute power levels, and it is therefore highly dependent on the choice of a power threshold. Since we program our channel emulator by specifying path gains rather than powers relative to some noise floor, RMS delay spread is particularly well-suited to both our studies and our equipment.

$$\sigma_\tau = \sqrt{\bar{\tau^2} - (\bar{\tau})^2} \tag{1}$$

where

$$\bar{\tau} = \frac{\sum_k P(\tau_k)\tau_k}{\sum_k P(\tau_k)} \quad \text{and} \quad \bar{\tau^2} = \frac{\sum_k P(\tau_k)\tau_k^2}{\sum_k P(\tau_k)} \tag{2}$$

3) Variable Modulations: Systems using 802.11b are allowed to adaptively choose modulations for maximum data rates of 1, 2, 5.5, or 11 Mbps. One particularly attractive aspect of the D-Link product was the ability to restrict the cards to a single modulation. We observed the effects of different fading on each modulation, and we viewed the sometimes unpredictable ways in which the system chose its own modulation when allowed to operate in fully automatic mode. We speculated that, since the systems were not originally designed for mobile channels, they might react inappropriately to different fading conditions.

4) Signal to Noise Ratio: When setting the SNR for a given experiment, it was necessary to account for the link loss across all of our various components as well as the attenuation generated by the channel emulator. Rather than calculating a link budget from the nominal values of all of our cables, adapters, circulators, etc., we monitored the SNR from the ends of the link using the Orinoco link management utility. The Orinoco software allowed us to see the SNR at both hosts simultaneously, and we adjusted the value using a combination of hardware attenuators and the programmable attenuation of the channel emulator. In cases where we used only D-Link cards, we inserted the Orinoco card into the loop at the beginning of the experiment and then removed it after setting the SNR but before beginning to collect data.

The dependence of our measurements on SNR made it particularly important that we isolate our peer-to-peer link from any outside interference while simultaneously preventing the hosts from communicating through the air and effectively "short-circuiting" our emulated channel. Both goals were accomplished by the selection of high quality RF components and careful shielding. We verified the integrity of our shielding by attempting to detect the emulated connection with a third, unconnected WLAN card, and we insured the absence of ambient interference by selecting a channel not currently in use by any base stations in range.

III. EXPERIMENTAL RESULTS

Predictably, both implementations of 802.11b showed significant performance degradation in emulated mobile channels. However, we were surprised to observe the extreme natures and degrees of the differences in performance characteristics between the three experimental setups.

All of our data show a clear link between SNR and the susceptibility to flat fading and delay spread. Our delay spread tests demonstrate that increasing SNR yields enormous performance benefits. The effect is not so dramatic for our single-path tests, but those tests produced another interesting result. Particularly at higher SNR, the throughput and packet loss characteristics caused by Rayleigh and Rician fading were almost identical.

A. Orinoco to Orinoco

1) Velocity: Results for Rayleigh and Rician fading are shown in Figures 1 and 2. The two types of fading induced very similar performance characteristics, with Rician fading being the slightly more severe of the two. As we expected, throughput decreases monotonically with increasing velocity while the corresponding packet loss rate increases almost monotonically. The one slightly counterintuitive aspect of the data is the behavior of the packet loss rate for low velocities. In Figure 2 we observe that the packet loss rate actually decreases

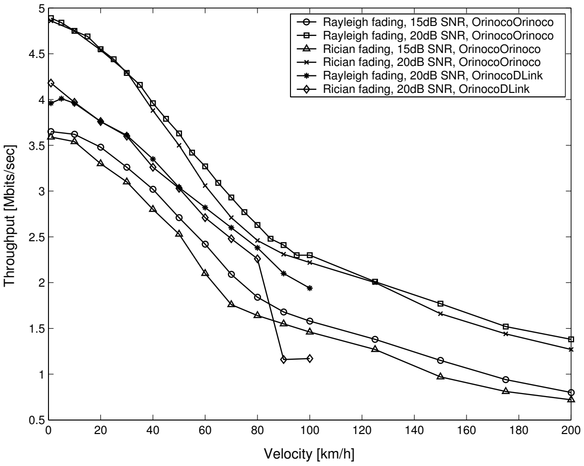

Fig. 1. Throughput vs. Velocity in Rayleigh and Rician Fading

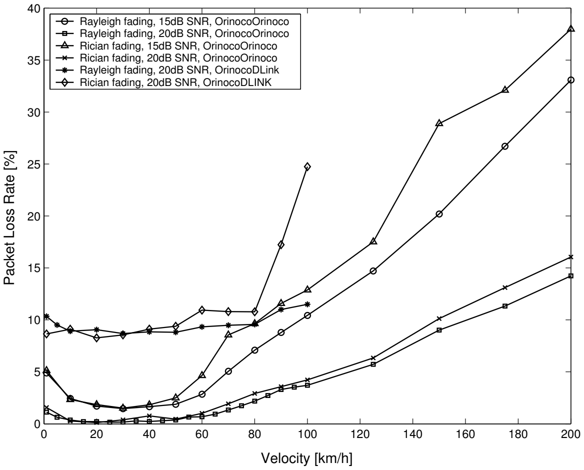

Fig. 2. Packet Loss vs. Velocity in Rayleigh and Rician Fading

as velocity increases from 0, and it does not increase again until velocity reaches approximately 40 km/h.

One plausible explanation for the negative slope region of the packet loss data is that the system's assumption of channel coherence time remains reasonable for low values of velocity, and the receiver is able to react to fading at a rate comparable to the speed of the fading. Clearly, the system is pushed to the limits of its channel estimation and equalization capabilities by higher velocities and the correspondingly faster fades. However, the poor performance at zero velocity may result more from the channel itself than the receiver's ability to react. At low velocities, it is possible that a user could remain in a deep fade for a long period of time, so some velocity could be beneficial in that it provides a wider variety of fading conditions over time. This explanation is supported by the fact that both 802.11b implementations exhibited higher packet loss rates at low velocity.

2) Delay Spread: As we can see in Figure 3, both throughput and packet loss remain reasonably constant for values of delay spread below some threshold. However, once the delay spread increases beyond that threshold, the system soon undergoes complete connection failure with all packets lost. The threshold appears to be highly dependent on SNR with a value of approximately 10 ns at 20 dB SNR and 40 ns at 25 dB SNR.

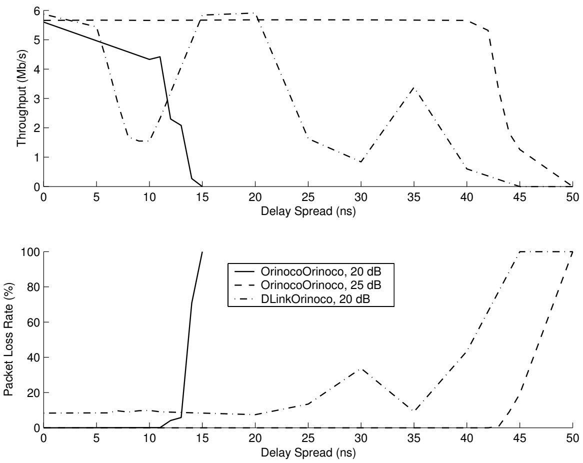

Fig. 3. Throughput and Packet Loss vs. Delay Spread

B. D-Link to D-Link

As mentioned before, one of the principle advantages of our decision to use D-Link wireless LAN cards was the fact that we were able to control the modulation on either or both sides of the peer to peer link. This provided us with two interesting insights into the performance of the system. We were able to isolate the four modulations (corresponding to the datarates 1, 2, 5.5, and 11 Mbps) and see how each performed in our mobile environment, and we compared the results with those achieved by the system when it was allowed to choose its own modulation. Thus, we were able to speculate about the strengths and weaknesses of the modulation selection algorithm and ascertain the degree to which performance characteristics were attributable to the algorithm.

1) Velocity: As expected, the general trend in Rayleigh fading measurements was for the throughput to decline monotonically while the packet loss rate increased with increasing velocity. As with the Orinoco cards, the packet loss rate at zero velocity tends to be higher than the rates at 10, 20, and even 30 km/h. We have not been able to explain this phenomenon completely, but it is consistent across both measurements and implementations, so it is not the result of aberrant data.

Comparison of the data from fully automatic operation with that of the individual fixed modulations gives us an intuition for the decisions made by the modulation selection algorithm. Taking the 15 dB throughput measurements in Figure 4 as an example, it is evident that the system reacts to increased velocity by shifting from the 5.5 Mbps modulation to the 1 or 2 Mbps modulations. There is a transitional period between 0 and 20 km/h in which the fully automatic mode appears to use a combination of 5.5 and 2 Mbps modulations, and after that point it appears to use 1 or 2 Mbps almost

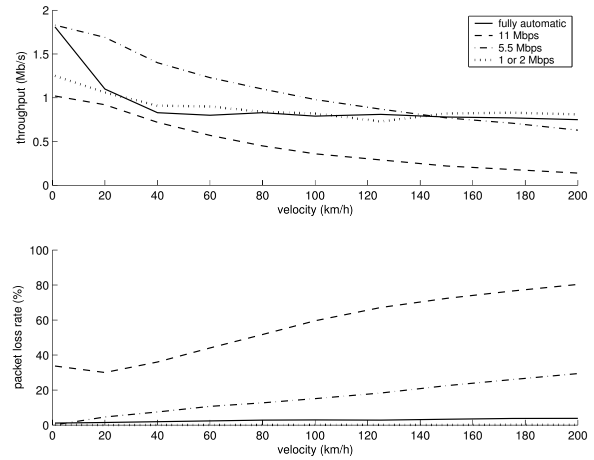

Fig. 4. Throughput and Packet Loss vs. Velocity in Rayleigh Fading for 15 dB SNR in D-Link Setup

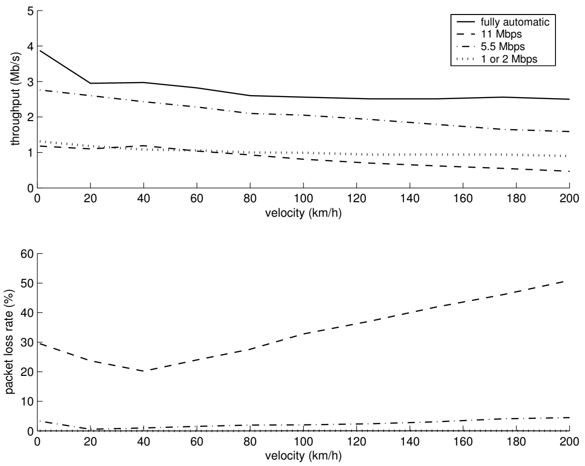

Fig. 5. Throughput and Packet Loss vs. Velocity in Rayleigh Fading for 20 dB SNR in D-Link Setup

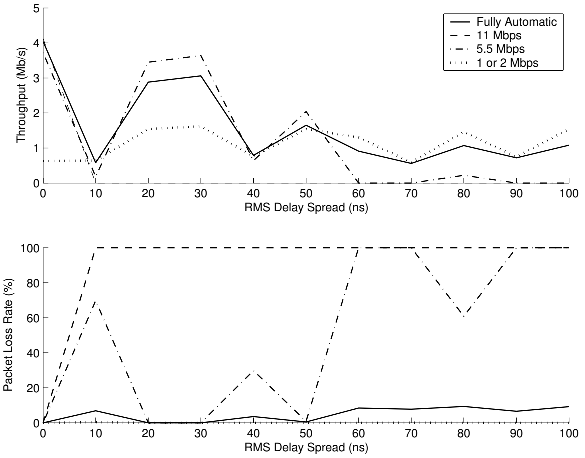

Fig. 6. Throughput and Packet Loss vs. Delay Spread

exclusively. Discrepancies between the packet loss rates of the fully automatic mode and 1 or 2 Mbps mode in Figure 4 indicate that the automatic mode continued to attempt the use of higher-order modulations that experience higher packet loss rates. Returning to the throughput data, it appears that the system would have been better served if it had continued to use the 5.5 Mbps modulation until velocity reached approximately 140 km/h. However, since the physical layer is unaware of the nature of the data traffic that it is carrying (in this case UDP), it may have decided that the packet loss rate of the 5.5 Mbps modulation was unacceptably high. Indeed, TCP traffic is particularly sensitive to dropped packets [10], and that may be the reason for the conservative choice of modulation.

In contrast, the 20 dB data in Figures 5 shows extremely effective performance of the fully automatic modulation selector. It actually achieves higher throughput than any single modulation for all values of velocity, and me must conclude that this is the result of effective adaptation to fading. In other words, the system was able to take advantage of favorable channel conditions by using higher order modulations, and it switched back to lower order modulations in adverse channel conditions to preserve its extremely low packet loss rate.

2) *Delay Spread:*

Our results for the D-Link cards in a multipath channel with varying delay spread are particularly interesting. Whereas the Orinoco implementation seen in Figure 3 exhibited relatively constant performance before undergoing link failure, D-Link's implementation has wildly fluctuating throughput and packet loss statistics. Thus, we conclude that the D-Link system is particularly sensitive to certain values of delay spread in addition to the general trend of decreased performance with increasing delay spread. Close examination of Figure 6 reveals that all of the modulations suffered from poor performance at delay spreads of 10 and 40 ns. However, performance improved dramatically before, after, and between the two points. In general, the 1 and 2 Mbps modulations performed better than the other modulations at the pernicious points, but the other modulations recovered well for more favorable values of delay spread. The fully automatic mode reacted effectively by favoring the best modulation at each of the delay spread values. The fully automatic throughput and packet loss data resides between the data from the 5.5 and 1 or 2 Mbps modulations. It appears to approximate the most favorable curve at each point, but its performance rarely exactly matches that of either fixed modulation which reflects the fact that it continuously attempts to use a mixture of the two modulations.

Without detailed knowledge of the D-Link hardware and receiver algorithms, it is difficult to isolate the exact cause of the performance fluctuations with delay spread. Nonetheless, we speculate that the root cause has to do with the receiver's ability to distinguish different multipath components from one another. When individual multipath components can be resolved from one another, maximal ratio combining increases the received signal power [12]. However, when two components are too close to be distinguished from one another, only the signal power of the dominant component is used, and the other component is treated as noise. Thus, inability to resolve different multipath components can lead to a reduction in SNR at the detector and correspondingly decreased performance. Another possible explanation involving the resolution of the RAKE receiver involves the granularity of the delay tolerated by the RAKE"fingers". It may be that there are some values of delay spread for which the multipath components do not fall neatly into individual bins and the signal power of those paths is split between multiple fingers. The end result is the same in that some signal energy cannot be resolved and becomes self-interference, thus decreasing the SNR.

C. D-Link to Orinoco

1) Velocity: The performance of the D-Link to Orinoco link exhibits characteristics similar to both of its constituent implementations. In addition, some of the apparent shortcomings may indicate some incompatibility between manufacturers rather than a weakness of either implementation. Indeed, Figure 1 shows that the D-Link to Orinoco link failed completely for SNR of 15 dB and for velocity greater than 100 km/h, so we were unable to even measure throughput and packet loss for the entire velocity range.

2) Delay Spread: Similar to the all D-Link configuration, the D-Link to Orinoco link exhibited wild throughput fluctuations for different values of delay spread. However, the mixed link was different in that its packet loss statistics did not reflect the same fluctuation. Therefore, we conclude that the D-Link card was transmitting with unwarranted conservatism and sacrificing throughput unnecessarily. If the receiving card had been another D-Link, then our data from the previous section indicates that the sender would have been correct in using a lower order modulation. However, the Orinoco receiver structure was not yet experiencing significant performance degradation. Therefore, we conclude that the D-Link card selected its modulation based on its perception of the symmetrical emulated channel and that the control system at work was open-loop rather than closed-loop.

IV. DISCUSSION AND FUTURE WORK

Our results clearly demonstrate that typical implementations of 802.11b are far less effective in mobile environments. However, we have no reason to believe that it would be impossible to adapt the implementations for mobile systems, while still retaining compliance to the 802.11b standard. In fact, by quantifying the performance degradation due to mobile channel effects, our work contributes a starting point for such adaptations.

In particular, we have shown that RAKE receivers intended for low-mobility indoor applications are insufficient for use in outdoor channels with high values of RMS delay spread. Useful adaptations might include the addition of more " fingers" to track multipath components. We also demonstrated that when the cards are permitted to select their own modulations, their decisions do not always maximize throughput. This could be remedied with a reevaluation of the selection metric and perhaps some awareness of the nature of the traffic. Finally, we observed that incompatibilities between different implementations of the 802.11b protocol can contribute to significant performance degradation.

New technologies are increasingly blurring the line between WLAN and cellular technologies by taking data services outdoors [13] and bringing mobile telephony indoors to wireless LANs [14]. In fact, the integration of multiple interfaces into a single device has long been an active research area [15–17]. Overall, we conclude that, with some relatively minor structural modifications, 802.11b or a similar protocol could serve as an important component of future mobile wireless systems. It is neither intended nor suited for the replacement of cellular systems, but our results indicate that 802.11b could extend its domain well beyond the semi-stationary, indoor propagation environments that it currently occupies.

ACKNOWLEDGMENT

This work was partially funded by an NSF CISE Research Infrastructure grant, EIA-0224458.

Christopher Steger receives funding from a 2001 Natl Defense Science and Engineering Graduate (NDSEG) fellowship.

The authors would also like to thank Spirent Communications for their frequent loans of test equipment.

REFERENCES

[1] http://www.atheros.com/news/mercedesdemo.html, "Atheros chipsets used in Mercedes-Benz future-technology demonstration."

[2] http://www.meshnetworks.com, "MEA Mesh Enabled Architecture."

[3] http://www.roamad.com, "RoamAD."

[4] A. Kamerman and G. Aben, "Throughput performance of wireless LANs operating at 2.4 and 5 GHz," in *Proc of The 11th IEEE International Symposium on Personal Indoor and Mobile Radio Communications (PIMRC)*, vol. 1, pp. 190–195, September 2000.

[5] Z. Hadzi-Velkov and B. Spasenovski, "The influence of flat rayleigh fading channel with hidden terminals and capture over the IEEE 802.11 WLANs," in *Proc of The 54th IEEE Vehicular Technology Conference*, vol. 2, pp. 972–976, Fall 2001.

[6] B. Bing, "Measured performance of the IEEE 802.11 wireless LAN," in *Proc of The Conf On Local Comp Networks*, pp. 34–42, Fall 1999.

[7] H. Zepernick and T. Wysocki, "Multipath channel parameters for the indoor radio at 2.4 GHz ISM band," in *Proc of The 49th IEEE Vehicular Technology Conference*, vol. 1, pp. 190–193, Spring 1999.

[8] C. Huang and R. Khayata, "Delay spreads and channel dynamics measurements at ISM bands," in *Conference record of the IEEE International Conference on Communications, SUPERCOMM/ICC '92, Discovering a New World of Communications*, vol. 3, pp. 1222–1226, June 1992.

[9] G. Woodward, I.Oppermann, and J. Talvitie, "Outdoor-indoor temporal and spatial wideband channel model for ism bands," in *Proc of The 50th IEEE Vehicular Technology Conference*, vol. 1, pp. 136–140, Fall 1999.

[10] H. Chaskar, T. Lakshman, and U. Madhow, "On the design of interfaces for TCP/IP over wireless," in *Proc of The IEEE Military Communications Conference (MILCOM)*, vol. 1, pp. 199–203, October 1996.

[11] T. Rappaport, *Wireless Communications*. Upper Saddle River, NJ; Prentice Hall, 1996.

[12] J. Proakis, *Digital Communications*. New York, NY; McGraw Hill, 4th ed., 2001.

[13] http://www.nokia.com/phones/nokiad211/, "Nokia D211."

[14] http://www.motorola.com/wlan/index.html, "WLAN seamless voice and data solution."

[15] E. A. B. amd R. H. Katz and Y. C. et al., "Network architectures for heterogeneous mobile computing," *IEEE Personal Communications*, vol. 5, pp. 8–24, October 1998.

[16] B. Aazhang and J. R. Cavallaro, "Multitier wireless communications," *Kluwer Journal on Wireless Personal Communications*, vol. 17, pp. 323–330, June 2001.

[17] P. Frantz, "A prototype platform for a multitier network interface card (mNIC)," in *Proc of The Communications Design Conference*, September 2002.