

Link Assessment in an Indoor 802.11 Network

Michael R. Souryal, Luke Klein-Berndt, Leonard E. Miller, and Nader Moayeri
Wireless Communication Technologies Group
National Institute of Standards and Technology
Gaithersburg, Maryland 20899-8920

Abstract—This paper presents the empirical results and analysis of a link variability study for an indoor, stationary 802.11 network subject to external interference. The objective is to characterize the relationships between the packet loss rate on a link and the transmitter-receiver distance, the SNR as reported by the 802.11 card, and time. Experimental results confirm that distance is not predictive of link reliability. Furthermore, attempting to estimate the packet loss rate by counting missing packets can lead to large estimation errors and latency for estimation parameters of practical interest. We find, however, that the SNR measurements reported by today’s 802.11 cards can, under certain circumstances, be reliably predictive of the packet loss rate. These results suggest that SNR-based link assessment can provide for more rapid and accurate measurement of link reliability, but they point to the need for improved SNR measurements and models to adequately predict link quality in environments with significant external interference.

I. INTRODUCTION

There is a growing appreciation in the ad hoc and sensor network research community of the need for effective link assessment. The ability to distinguish between reliable and unreliable links is critical to constructing and maintaining useful routes in multihop networks. This paper presents results of a link performance study on a stationary indoor 802.11 wireless network and draws lessons from these results that inform the design and development of effective link assessment techniques.

Many existing ad hoc network routing protocols require that a node detect and maintain a list of its current neighbors for possible use in routing user traffic. For a number of reasons, the neighbor list may include links that are poorly suited for forwarding traffic, and routing through them can lead to excessive retransmissions and route errors [1]. The ability to assess or grade different links can help in avoiding links with high loss rates. In signal stability-based adaptive routing, for example, the route discovery protocol favors the selection of stronger, more stable links based on signal strength measurements observed over time [2]. The authors of [1] propose using a signal-to-noise ratio (SNR) threshold to discard “weak” control messages, thereby eliminating the consideration of low-quality neighbor links.

Others have proposed including a measure of link quality in the route metric itself. The expected transmission count metric (ETX), for example, is a function of the packet delivery ratio in each direction of a link and is used to estimate the average number of transmissions (including retransmissions) needed for a packet’s traversal of a route [3]. In the context of position-based routing, the expected progress metric takes into account

not only the progress towards the final destination offered by a relay, but also the probability of successful transmission on that link, mitigating the selection of high-progress, low-quality links (e.g., [4]).

In light of the previous and existing work incorporating link quality in ad hoc routing, there is a need for an improved understanding of real-world link dynamics as well as the issues and opportunities of quantitatively assessing link quality. A thorough study of link dynamics in the outdoor Roofnet 802.11 network concluded that links of marginal quality are the norm and that the SNR reported by the physical layer is not predictive of link reliability [5]. However, this conclusion may not be warranted because the study also demonstrated that the multipath delays experienced outdoors, which far exceeded the delays tolerated by the cards that were used, were the likely contributor to elevated loss rates.

Our study focuses on an indoor 802.11 network and sheds further light on the causes of varying link quality. In particular, it investigates the relationship of a link’s packet loss rate with link distance, SNR, and time. It confirms that link distance is not predictive of link reliability (i.e., a shorter link is not necessarily more reliable than a longer link). Moreover, attempting to estimate the packet loss rate in real time by observing recent transmissions on the link can lead to large estimation errors and latency on time-varying links. However, we find that the SNR reported by today’s 802.11 cards can, under certain circumstances, be reliably predictive of link quality, and the degree of predictability is related to the level of external interference. The analysis points to the need for and potential benefits of improved SNR measurements for rapid and accurate link assessment.

II. MEASUREMENT METHODOLOGY

A. Platform

The study was conducted using 10 transceiver nodes and a control PC. The nodes were CerfCube [6] embedded systems equipped with Compact Flash 802.11 wireless LAN (WLAN) cards with RF output power of 17 dBm. In addition to the air interface, the nodes were connected to the control PC through a wired ethernet backbone network. The PC used the wired network to control transmissions and to collect results at the end of each transmission. Nodes were located throughout the fourth floor of the NIST North office building, laid out as indicated in Fig. 1, and were positioned to ensure a variety of link conditions.

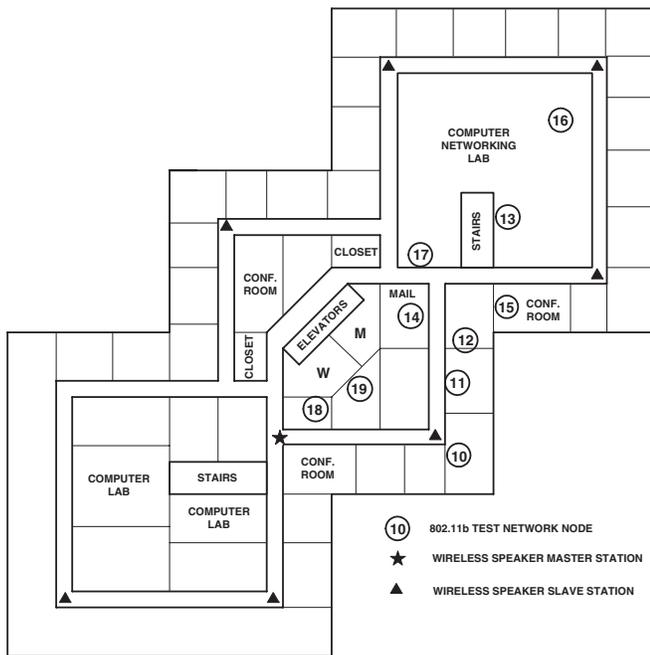


Fig. 1. Physical layout of testbed

B. Transmissions

During a test, each node takes a turn transmitting while the other nodes record the packets that they successfully receive. The sending node broadcasts IPv4, UDP packets with a data payload of 350 bytes. Each packet contains a unique identification (ID) number that is incremented with every new packet. The sending node transmits for a total of 2 minutes, pausing 50 milliseconds between each packet transmission. All transmissions are made at the 2 Mbps data rate.

Once the sending node finishes transmitting, the control PC collects the recorded packets from all other nodes. The PC then processes the collected packets, pulling out the relevant fields from the headers, and stores the collected information in a database for later inspection. This procedure continues in a round-robin fashion until each node has had a chance to send. After a round has been completed there is a pause before the next round of sending begins. For the tests conducted, there was a pause time of one hour between rounds, and each test lasted a total of 32 hours.

C. Measurements

The primary measurements recorded for each received packet are the packet ID, signal strength, and noise level. The sequential, unique ID in each packet is used during the post-processing of received packets to determine how many packets were not successfully received. Signal strength and noise measurements are delivered in additional headers provided by the 802.11 driver, in units of decibels referenced to one milliwatt (dBm). The noise measurement is the signal strength measured immediately prior to locking on to the frame preamble and is sometimes referred to as the “silence” measurement.

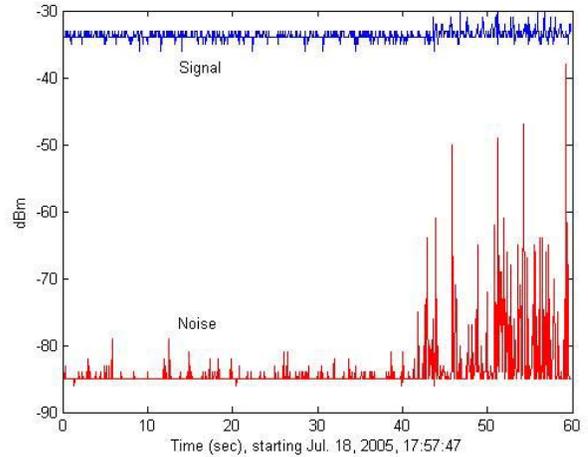


Fig. 2. Signal and noise levels reported by node 18 while node 11 was transmitting; FH speaker system turned on approximately 40 s into the measurement

Signal and noise measurements are only made available by the driver for frames that are successfully received. Furthermore, the 802.11 driver headers are only accessible if the receiving cards are placed into Monitor Mode, a special state available on cards based on the Prism chipset. In addition to these headers the 802.11 MAC, IP and UDP headers are recorded. From the collected headers it is possible to extract a wide variety of information, all of which is stored in a database for later examination.

D. Channel Environments

The testbed operated on 802.11 channel 6, with other WLANs on the same or nearby floors of the building operating on other channels. Experiments were conducted for two markedly different channel environments: with and without the presence of a strong interference signal. In the former, a strong frequency-hopping (FH) interference signal is continuously generated by a wireless emergency speaker system installed in the building. The FH signal hops over 79 frequency slots spaced 1 MHz apart in the 2.4 GHz unlicensed band (similarly to Bluetooth devices and 2.4 GHz FHSS cordless phones). The locations of the FH transmitters are indicated in Fig. 1, each with RF output power of 20 dBm. The master station transmits almost continuously, and each slave station (speaker) replies with test data on a low duty cycle.

The effect of the FH interference can be seen in Fig. 2 showing signal and noise levels versus time measured at node 18 (near the FH master station) while node 11 was transmitting. Even though measurements are only reported by the card on successfully received packets, an increase in noise is clearly observed when the FH speaker system was turned on approximately 40 s into this sample.

A 32-hour test was conducted with the FH system turned on, and an identical test was conducted with it off. These results are used to assess the impact of interference on link quality as well as on link assessment procedures.

III. RELATIONSHIP OF PACKET LOSS TO DISTANCE

We expect link reliability to be loosely related to distance, since the former also depends on the local interference at the receiver and non-distance based channel propagation effects. These effects include signal loss due to obstructions and variability due to shadowing and multipath fading. The results in this section confirm the low correlation between packet loss and distance and show the extent of the variation in an indoor office environment across different interference environments.

Fig. 3 displays scatter plots of the packet loss rate versus link distance for low, moderate and high-interference environments. The low-interference results were collected when the FH system was off and during non-working hours, when other 802.11 traffic is minimal. The moderate-interference results were collected during working hours when the FH system was off, and the high-interference results were collected while the FH system was on and consist of both daytime and nighttime measurements. Each point in the scatter plots reflects the average packet loss rate measured during a two-minute sample.

Considering first the low-interference results in Fig. 3(a), packet loss rates are low for short links and high for long links, but they vary widely for link distances between 45 and 90 ft. Indeed, a 90-ft link can be more reliable than a link half that distance. The transition range between good and bad links is even wider in moderate interference (Fig. 3(b)), where it starts at 20 ft, and in high interference (Fig. 3(c)), where it starts at 10 ft. Clearly, link reliability varies widely with distance in the transition range, and the transition range increases with interference.

To isolate the effect of interference on link variability, Fig. 4 shows scatter plots of the average received signal power versus link distance for the low and high interference scenarios (results for the moderate interference scenario are similar). In each case, there is a trend of decreasing signal strength with distance, but there is also variability of up to 40 dB at a given distance, regardless of the interference environment. This variability is a result of the different obstructions, shadowing and fading characteristics of the links.

Among the implications for the analysis and design of ad hoc networks are, first, that distance cannot be used as a predictor of link reliability. Unit disk graph (UDG) models, for nodes with omnidirectional antennas, oversimplify true propagation characteristics. Improvements to UDG models that assume a monotonic relationship between link quality and distance (e.g., [7]) are also not accurate. Second, though these results are for a stationary network, they imply that small changes in distance can produce large changes in link quality. Thus, in a network with mobility, link quality can fluctuate rapidly with time. Third, the wide range of link quality implies that diversity in relay selection can significantly improve transmission reliability. Without increasing the number of hops in a multihop route, a backup relay providing comparable progress to the final destination may offer significantly better link quality.

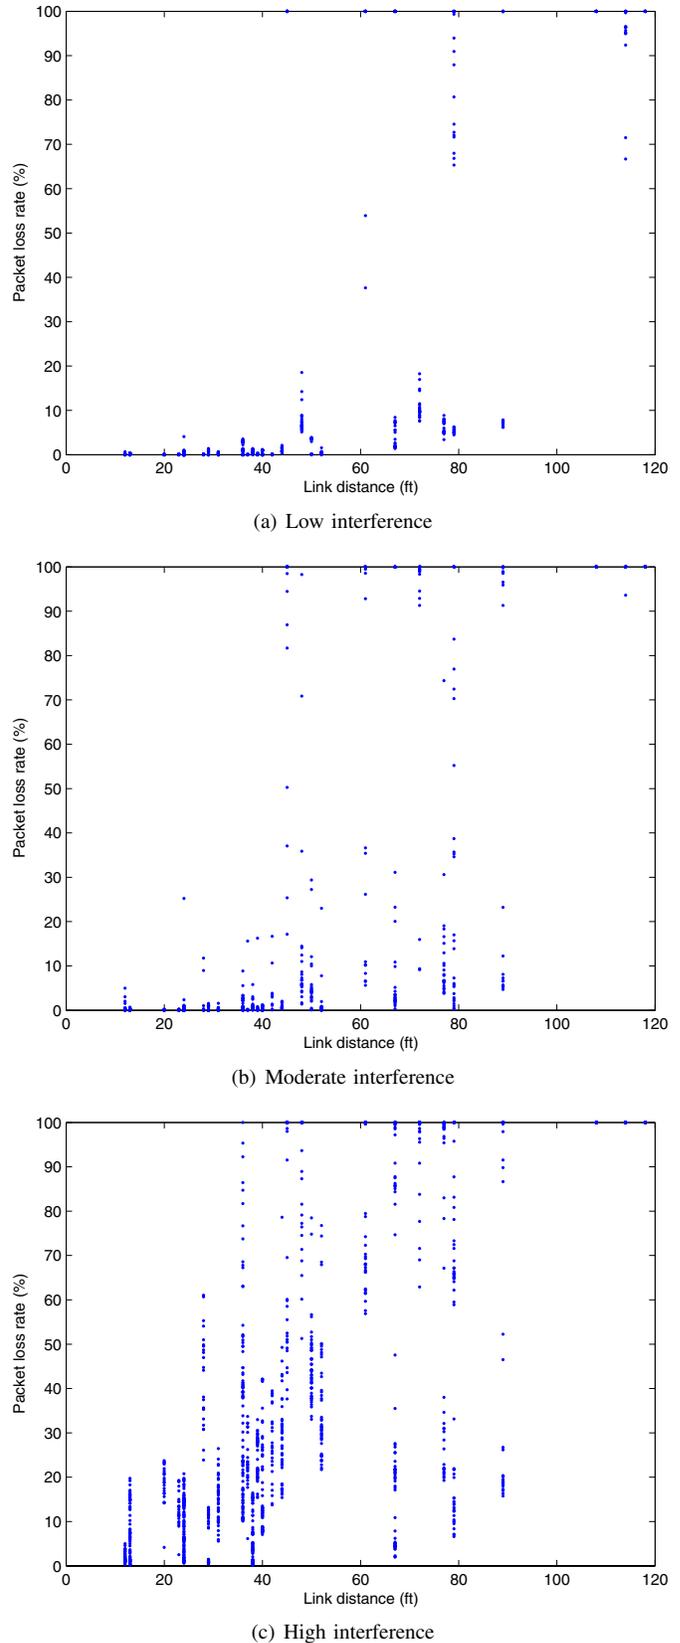
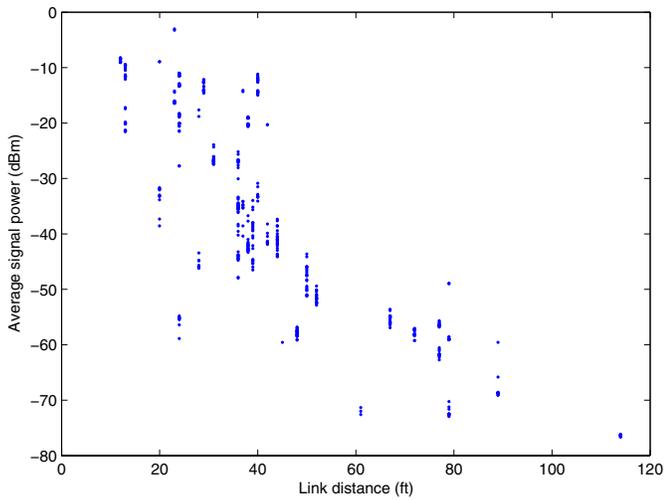
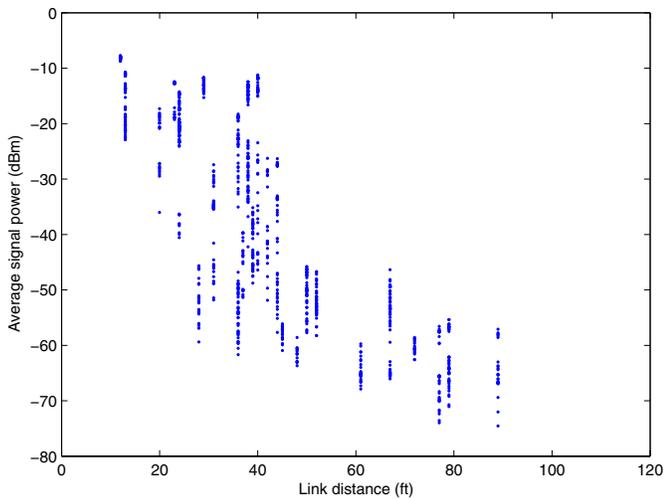


Fig. 3. Packet loss rate vs. link distance



(a) Low interference



(b) High interference

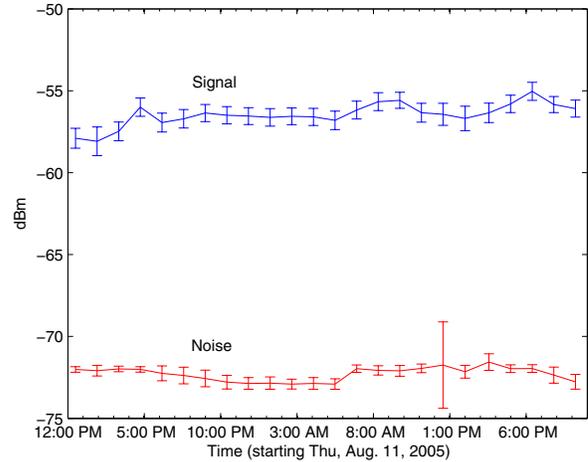
Fig. 4. Received signal strength vs. link distance

IV. PACKET LOSS VARIABILITY OVER TIME

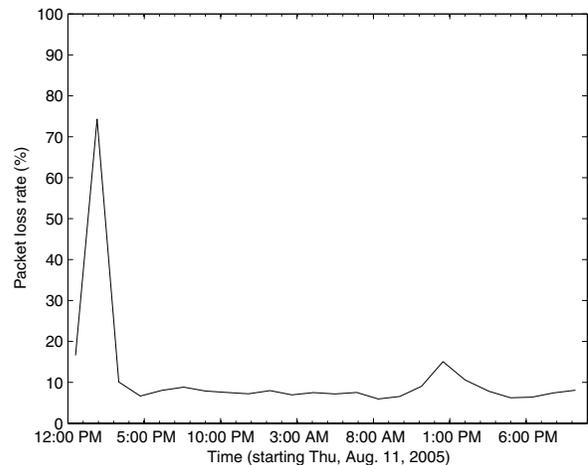
A. Long-Term Variability

Even in a stationary network, performance varies over time due to changing interference conditions as well as the movement of obstacles and reflections. Fig. 5 illustrates the long-term variability of the average signal strength, noise, and packet loss on the link from node 16 to node 12 over the entire 32-hour experiment with the FH system off. Each data point represents the average over a two-minute sample. The standard deviation of the signal and noise during each sample is indicated in Fig. 5(a) by the length of the bar at each point. Though the long-term packet loss rate is 11% on this link, it peaks to 74% around 2:00 pm of the first day, when the SNR reaches its lowest point.¹ A smaller peak occurs around 12:30 pm the following day, which is due to a burst

¹We observe in Section V that this SNR value is near the threshold at which packet loss changes rapidly with SNR.



(a) Signal and noise



(b) Average packet loss rate

Fig. 5. Long-term performance vs. time

of interference activity during that sample, as shown by the large standard deviation of the noise at that time.

B. Short-Term Variability

The instantaneous link quality can change due to small-scale variations in the multipath propagation profile and due to bursty interference. For routing protocols or link adaptation techniques to adapt to changes in link quality, a node must be able to estimate link reliability and respond to it on a short timescale. A straightforward method for estimating a link's packet loss rate is to count missing packets over some period of time. These packets could be periodic beacons generated by each node at a known interval. The accuracy of the estimate is affected by the length of the period of observation, the number of packets transmitted during that period, and the time-variability of the channel. Clearly, numerical granularity of the estimate increases with the number of packet transmissions being counted. On the other hand, a longer observation period loses temporal granularity for capturing time-varying

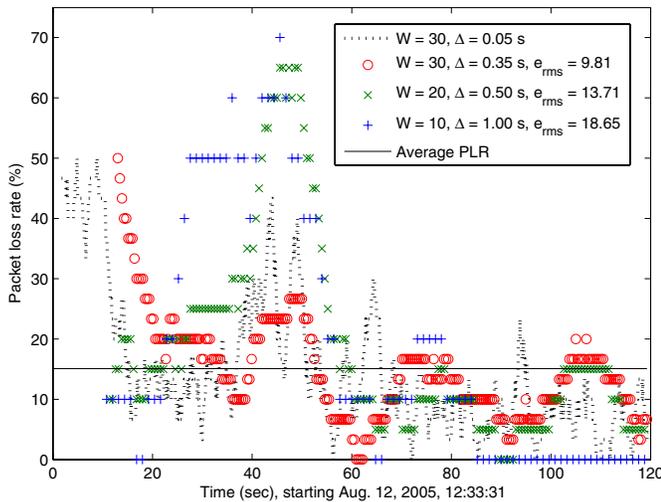


Fig. 6. Example instantaneous packet loss rate estimates vs. time (W = window size, Δ = inter-packet interval, e_{rms} = root-mean-square error)

effects. Another tradeoff is the increased estimation accuracy afforded by more transmissions versus the increased load on the network imposed by these transmissions.

Data obtained through our experiments can be used to analyze the impact of various packet loss rate estimation parameters. Fig. 6 displays estimates of the instantaneous packet loss rate versus time for the link from node 16 to node 12 during the sample with bursty interference noted in Section IV-A. This example is illustrative of links with moderate time-varying packet loss. The average packet loss rate over the entire two-minute sample is 15%, indicated by the solid line in the figure. A “best estimate” of the instantaneous packet loss rate is shown by the dotted line and is based on a sliding observation window (W) of 30 packets with inter-packet interval (Δ) of 50 ms (the minimum available interval in our study). It shows that the packet loss rate fluctuates between zero and 50% on this link. Three additional estimates are shown for what might be considered typical beacon intervals. The observation period in each case is about 10 s, but the number of transmissions observed during the period ranges over $W = 10, 20$ and 30. The root-mean-square error (e_{rms}) between these estimates and the “best estimate” is shown in the legend. It is clear that estimates using window sizes of 10 and 20, with inter-packet intervals of 1 s and 0.5 s, respectively, do not track the instantaneous packet loss rate well, at times largely overestimating or underestimating it. However, the estimate based on a window size of 30 transmissions (and transmission interval of 350 ms) follows the general shape of the instantaneous curve, but even then it experiences a measurement latency on the order of the observation period.

This example highlights some of the inherent challenges in assessing link quality through direct measurements of packet loss. Basing packet loss rate measurements on the periodic broadcast of beacon packets, with typical beacon intervals being on the order of one second, may lead to large estimation errors. Furthermore, the latency imposed by measurement

windows that provide sufficient numerical granularity may render adaptive techniques ineffective in attempting to adapt to the current link state. Future work should examine these tradeoffs more fully, and should investigate faster and more accurate techniques for estimating packet loss. The basis for one such technique is discussed in the following section.

V. RELATIONSHIP OF PACKET LOSS TO SNR

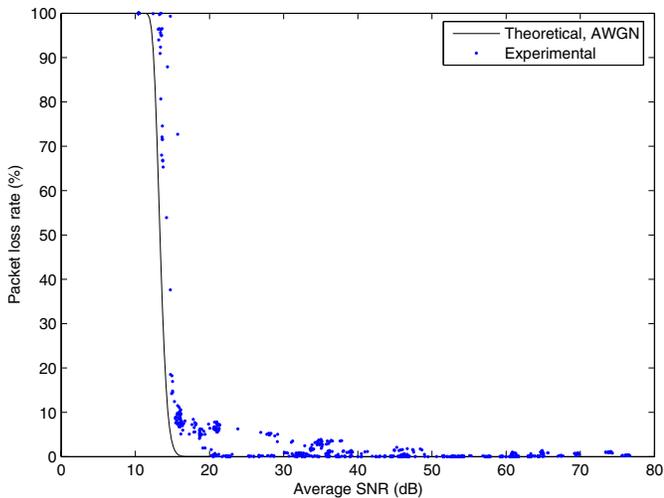
To investigate the relationship between packet loss and SNR, we compare the theoretical relationship with that observed experimentally. The 2 Mbps data rate of 802.11 utilizes differential quadrature phase shift keying (DQPSK). The theoretical probability of bit error, P_b , of DQPSK in additive white Gaussian noise (AWGN) is given by [8, eq. (5.2-70)]. Since the 2 Mbps data rate of 802.11 includes no forward error correction, the probability of packet success is simply the probability that no bit errors occur. Assuming bit errors are independent, the packet error probability for a packet of length N bits, then, is $P_p = 1 - (1 - P_b)^N$.

Fig. 7(a) is a plot of P_p versus the symbol SNR along with a scatter plot of experimental results for the low-interference environment. The theoretical packet loss rate exhibits a sharp transition with SNR. We observe that the experimental results closely track the theoretical result, exhibiting the same sharp transition. The fact that the experimental results are close to the theoretical result is indicative that the channel environment in this case is well modeled by AWGN.

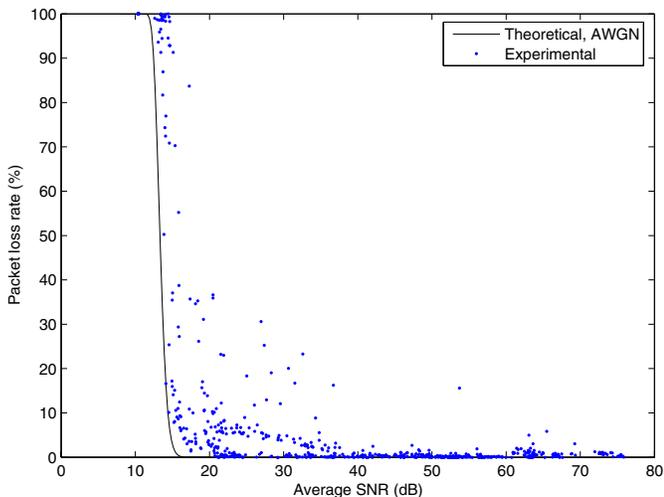
Fig. 7(b) shows results for the moderate-interference environment. The key difference relative to the low-interference results is the greater spread in packet loss with respect to SNR. For example, links with measured SNR in the neighborhood of 30 dB exhibit packet loss rates as high as 20%, whereas in the low-interference environment such links exhibit much lower packet loss. This difference is attributable to the greater level of interference. Note that the effect of the interference is largely hidden from the SNR measurements reported by the 802.11 card, since these measurements are reported only for *successfully received* packets. The sharp distinction between reliable and unreliable links as a function of SNR is clearly absent here, as the channel now includes non-negligible interference in addition to the thermal noise in the receiver.

Fig. 7(c) illustrates results for the high-interference environment, that is, in the presence of the FH interferer. Here, we observe an even larger spread in packet loss, and one which is present over a broader range of SNR. Even links reporting SNR as high as 50 dB are losing one out of five packets due to collisions with the interferer. In both the high and moderate-interference environments, the theoretical curve provides a lower bound on packet loss, but the amount of variation above that bound is related to the level of interference present on the channel.

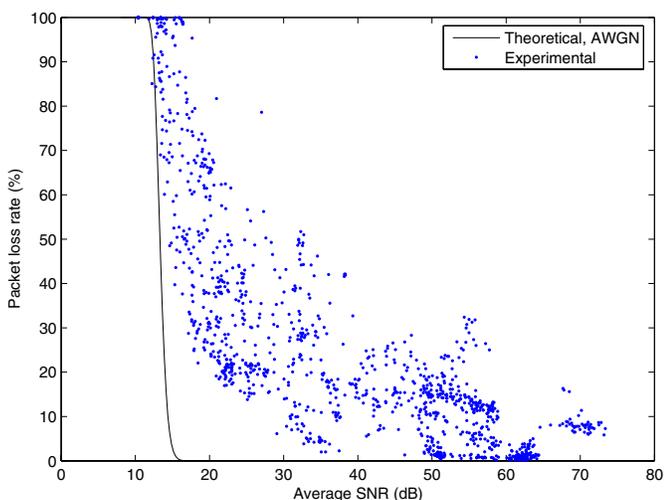
These observations have implications for the usefulness of the SNR measurements reported by today’s 802.11 cards. In low-interference conditions, the SNR serves as a good indicator of link reliability; in fact, the relationship of packet loss to SNR in this case closely follows that predicted for



(a) Low interference



(b) Moderate interference



(c) High interference

Fig. 7. Packet loss rate vs. reported SNR

AWGN channels. Link assessment, therefore, can be based on relatively few SNR observations compared with schemes that estimate packet loss rate by counting missing packets. As a result, links can be assessed more quickly, which is particularly important when the channel state is time-varying. With rapid link assessment, multihop routes can be constructed using more timely and accurate information on the quality of neighbor links, and advanced channel-adaptive techniques such as adaptive routing and adaptive modulation can be used to exploit temporal and spatial channel diversity.

As external interference increases, however, the relationship between SNR and link reliability becomes less predictable, and the SNR measurements of today's cards are less useful for link assessment. Reporting measured SNR on both corrupted and successfully received packets, as well as better models relating SNR to packet loss in channels with interference, should improve the usefulness of these measurements for link assessment.

VI. CONCLUSIONS

A link performance study of a stationary indoor 802.11 network was conducted to better understand link dynamics in such a network and to investigate the opportunities and challenges for real-time link assessment, a critical feature for the adequate performance of many ad hoc network routing protocols. The study confirms that distance plays a weak determinative role in predicting link reliability. Furthermore, link assessment based on the direct measurement of packet loss through packet counting can result in unacceptable estimation error and latency. However, the analysis of data collected with and without the presence of externally generated interference reveals that the SNR reported by today's 802.11 cards can, in the case of low interference, be a predictive measure of link reliability and, therefore, provide a more rapid means of link assessment. Future development should focus on improved SNR measurements and models that adequately account for possible external interference.

REFERENCES

- [1] H. Lundgren, E. Nordström, and C. Tschudin, "Coping with communication gray zones in IEEE 802.11b based ad hoc networks," in *Proc. 5th ACM Int. Workshop on Wireless Mobile Multimedia*, 2002, pp. 49–55.
- [2] R. Dube, C. D. Rais, K.-Y. Wang, and S. K. Tripathi, "Signal stability-based adaptive routing (SSA) for ad hoc mobile networks," *IEEE Pers. Commun.*, vol. 4, no. 1, pp. 36–45, Feb. 1997.
- [3] D. S. J. De Couto, D. Aguayo, J. Bicket, and R. Morris, "A high-throughput path metric for multi-hop wireless routing," in *Proc. MobiCom*, 2003, pp. 134–146.
- [4] M. R. Souryal and N. Moayeri, "Channel-adaptive relaying in mobile ad hoc networks with fading," in *Proc. Second Annual IEEE Commun. Society Conf. on Sensor and Ad Hoc Commun. and Networks (SECON)*, Sept. 2005.
- [5] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris, "Link-level measurements from an 802.11b mesh network," in *Proc. Special Interest Group on Data Communication (SIGCOMM)*, 2004, pp. 121–132.
- [6] Intrinsyc Software, "Cerfcube 255," Available online at http://www.intrinsyc.com/products/mob_ref_sys/cerfcube_255/.
- [7] R. Negi and A. Rajeswaran, "Physical layer effect on mac performance in ad-hoc wireless networks," in *Proc. Commun., Internet and Info. Tech. (CIIT)*, Nov. 2003.
- [8] J. G. Proakis, *Digital Communications*. 4th ed., New York: McGraw-Hill, 2001.