

Performance of 802.11b/g in the Interference Limited Regime

Vinay Sridhara Hweechul Shin Stephan Bohacek

vinays@qualcomm.com hcshin@udel.edu bohacek@udel.edu

Department of Electrical and Computer Engineering

University of Delaware

Newark, DE 19716

Abstract

Signal to interference plus noise ratio, SINR, is one of the main factors that affects the quality of wireless communication. While the impact of white Gaussian noise on a wireless channel is well understood, impact of interference remains one of the less explored areas. With the deployment of dense mesh networks, the interference will be a dominant factor that affects the transmission errors. This paper explores the performance of 802.11b/g when subject to interference. The findings are based on various controlled experiments in the laboratory setting. One finding of this work is that in contrast to communication over links where the noise is Gaussian, in 802.11b/g, the probability of successfully transmitting a packet is dominated by the ability of the receiver to synchronize with the carrier. As a result, changing to a lower bit-rate with same synchronization scheme will not make the transmission more resilient to interference. Moreover, we found that the impact of interference is significant even if the interference is 10 dB below the strength of the noise. The significance of this result on bit-rate selection is briefly explored.

This work was prepared through collaborative participation in the Collaborative Technology Alliance for Communications and Networks sponsored by the U.S. Army Research Laboratory under Cooperative Agreement DAAD19- 01-2-0011. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation thereon.

S. Bohacek is the corresponding author. Address: 140 Evans Hall, Dept. of Electrical and Computer Engineering, University of Delaware, Newark, DE 19716. Phone: 302-831-4274. Fax: 302-831-4316. Email: bohacek@udel.edu.

Interference, SINR, wireless mesh networks, bit-rate selection, wireless experimentation.

I. INTRODUCTION

One of the key advantages of mesh networks over cellular-based networks is that mesh nodes are relatively inexpensive. Hence, it is economically feasible to spread mesh nodes at a high enough density to provide high data rates to a dense user population. Specifically, the high node density results in a short distance between receivers and transmitters, and hence high SNR channels and high bit-rates are possible. However, an important issue of high density networks is that transmissions are significantly impacted by interference. Thus, in low density networks, transmission errors are due to low SNR, while in high density networks, transmission errors are due to low SIR. The first scenario is referred to as the *noise limited regime* and the second scenario is referred to as the *interference limited regime*. The behavior of 802.11 is well understood when transmissions are noise limited. However, performance in the interference limited case is considerably less well understood. This paper explores 802.11b/g in the interference limited regime.

The findings presented here are the result of a large number of *in vitro* laboratory experiments. Thus, the channels between transmitters and receivers were controlled and all transmissions from external sources were eliminated with EMF shielding. The experiments lead to three key findings. First, in the interference limited regime, the behavior of 802.11b/g is dominated by the ability of the receiver to synchronize to the sender's carrier. Recall that transmissions begin with the broadcast of a synchronizing preamble followed by the data. Thus, in order to decode the data, the receiver must first successfully synchronize. The experiments discussed here show that in the interference limited regime, for a large number of 802.11b/g bit-rates, transmission errors are mostly caused by synchronization errors. In 802.11b, there are two similarly performing synchronization schemes, and 802.11g provides a third scheme. Thus, while 802.11b/g provides a large number of bit-rates, it essentially only provides two synchronization schemes. As a result, in the interference regime, most of the 802.11g bit-rates have the same tolerance to interference. Therefore, in contrast to the noise limited regime, in the interference limited regime decreasing the data rate provides no added ability to decode frames. Furthermore, since lowering the data rate increases the duration of the transmission, it increases the possibility that a collision will occur. As a result, in many cases, lowering the data rate decreases the performance. This has important implications for selecting the bit-rate that minimizes the time until a transmission

is successful. This issue is further discussed in Section V.

A second finding of the experiments describe in this paper is that if synchronization is successful, then the packet error is independent of the packet size. This significantly differs from the noise limited case where the probability of successful transmission obeys $(1 - p_{BE})^Z$, where p_{BE} is the probability of bit-error and Z is the frame size. Thus, in the noise limited case, the probability of packet loss is exponential in Z . A third finding of this work is that weak interference can have a significant impact of the ability to successfully transmit a packet. For example, in some cases, an interfering signal that is nearly 9 dB weaker than noise will result in synchronization error with probability 0.5.

The remainder of this paper proceeds as follows. In the next section, the experiments and experimental set-up is described. Section II-B presents some experimental results. Section III discusses computation of transmission error probability from the data collected. Section IV explores the ability to transmit packets when both interference and noise are present. Section V describes how the models presented in Section III can be used to determine the bit-rate that minimizes the expected time to successfully transmit a frame. And finally, Section VII provides concluding remarks.

Due to lack of space, details of 802.11 and process of decoding a frame are not included; see [1] and [2] for information on these issues.

II. EXPERIMENT DESCRIPTION

A. *Experimental Setup and Protocol*

Figure 1 depicts the block diagram of the experimental setup. The main aim of the setup is to precisely control of the channels between the transmitters and receivers and eliminate external interference.

The setup included of two access points, three laptops, and a controller computer. The two access points were the Cisco 1240 a/b/g [3] with Broadcom chip-sets. Prior work has shown that unlike some PCMCIA-based transmitters, the Cisco 1240 provides good transmit power stability. The sender and interferer controllers were used to adjust transmission bit-rate of the corresponding APs, transmit packets to the AP via the Ethernet (the AP then broadcasted these packets via the wireless transmitter), and to receive frame via the wireless transmitter (which was used for reference purposes). The receiver laptop was used to log all received frames. A modified MadWifi [4] driver was

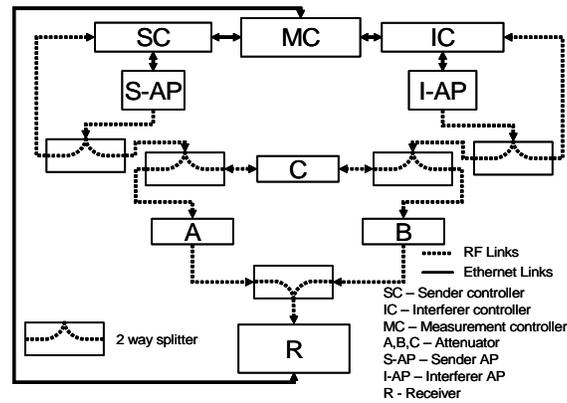


Fig. 1. The figure shows the block diagram of the experimental setup for all sets of measurements collected for analysis. The red lines indicate the RF cables used and the green thick line indicate the ethernet cable connections used for controlling the experiment.

used to collect all frames received, including those received with bit-errors. The sender and interferer controllers and the receiver were equipped with Proxim Orinoco b/g Gold Cards with the Atheros AR5212 chip-set [5], [6].

In order to conduct the experiments in a controlled and repeatable fashion, the receiver, the transmitter, and the interferer were isolated from each other. The isolation was achieved by using shielded wires to carry the "wireless" signals and by using attenuators. In order to prevent the RF leakage from the devices, all of the access points, laptops, splitters and connectors are wrapped with an RF resistant cloth. It was found that a single layer of the RF resistant cloth provides at least 40dB of attenuation, which was found to be sufficient to keep the weak RF leakage from affecting the experiments. Furthermore, the wireless transmissions used channel 1, while there were no nearby transmitters on channels 1-10.

Each attenuator was composed of calibrated Agilent 8495A and 8494A attenuators, providing repeatability within 0.3dB. These attenuators were used to control the received signal strengths, hence, the AP's transmit power was not used. In all cases, the combined attenuation of attenuator A and B exceeded 130 dB. This is critical since the signal transmitted by the sender will partially reflect off of the receiver and pass through the splitters to arrive at the interferer. Thus, a combined attenuation of 130 dB ensured that the interferer is unable to detect the sender and vice versa. The level of attenuation and the inability to communicate between sender and interferer when $C = -\infty$ dB was verified through experiments. Thus, the experimental setup resulted in the hidden node topology.

The objective of the experiment was to determine the probability of receiving a frame when the transmission is subject to interference. Since transmitting frames at precise times is difficult and error prone, frames were transmitted

at random times so that collisions occurred at random. Specifically, the sender transmitted packets at a fixed interval of approximately 12.4 msec between transmissions, while the interferer transmitted packets randomly with the time between transmissions exponentially distributed with mean 31 msec. The number of interferer transmissions was recorded, and hence the total duration that the channel was occupied by the interferer could be determined. With this duration and the transmission duration, the probability that the frame experienced interference can be determined. The details of this are provided in Section III.

Finally, each trial consisted of the sender broadcasting 10,000 frames with RTS/CTS disabled. The frames were broadcasted, and hence, there were no retransmissions nor ACKs. It was found that 10,000 frames resulted in a suitably small confidence interval. (The analysis with confidence intervals is not included in this paper due to lack of space.)

B. Experimental Results

As discussed above, the performance in the interference limited regime was investigated by transmitting packets so that collisions occurred at random. Some of the results of these experiments are shown in Figure 2. In the cases shown, the average time between the beginning of the interferer's frames was approximately 31 msec. The interferer frames were 576B and sent at 1Mbps. The sender's frames were 1464B.

At high SIR, transmission errors occur with low probability. As expected, at low SIR, transmission errors occur regularly. It is important to note that the x-axis Figure 2 is the SIR when a collision occurs, but collisions do not always occur. Therefore, the fraction of transmissions errors at low SIR is the probability of a collision occurring, i.e., the SIR is so low that if a collision occurs, then there is always an error. Since the SNR is high, when a collision does not occur, the frame is decoded with a high probability.

For 11, 12, and 24 Mbps, there is a plateau between the low SIR and the high SIR regions. Although not shown, this plateau also occurs at 5.5, 6, 9, and 18 Mbps. In the case of 6, 9, 12, 18, and 24 Mbps, this plateau ends at around 12 dB. For SIR above 12 dB, the observed probability of error is nearly the same for these bit-rates. Bit-rates 5.5 and 11 behave similarly, but the plateau ends at 7 dB. While not shown in Figure 2, 1 and 2 Mbps are similar to 11 Mbps in that they transition between a non-zero probability of transmission error and (nearly) zero probability of transmission error at 11 dB.

The reason that 6, 9, 12, 18, and 24 all have a plateau that ends at the same SIR is that they all use the same synchronization scheme and, apparently, this synchronization performs poorly when subjected to interference such that the SIR is less than 12 dB. Similarly, 1, 2, 5.5, and 11 all use similar synchronization schemes¹ that perform poorly when the SIR is less than 7 dB. Indeed, the height of each plateau is the probability that the sender's synchronization phase overlaps with the interferer's transmission (See the next section for details). Thus, for 802.11g (802.11b), if the SIR during synchronization is below 12 dB (7dB), then synchronization will fail nearly every time. If the sender's synchronization phase does not overlap with the interferer's transmission, then there still is a possibility that the sender's data transmission will overlap with the interferer's transmission. The probability of incorrectly decoding the data part of the frame depends on the bit-rate. Figure 2 indicates that for many bit-rates, the SIR that results in errors in decoding the data part of the packet is considerably smaller than the SIR required to synchronize. For example, at 12Mbps, few data errors occur if the SIR exceeds 5 dB, but synchronization will always fail unless the SIR exceeds 12 dB. In summary, the region where synchronization always fails and data decoding mostly succeeds is exactly the plateau. Thus, the plateaus end at similar points when the synchronization schemes are the same. Since different modulation schemes have different tolerance to interference, the plateaus begin at different points. Unlike the height of the plateaus, the fraction of lost frames at low SIR depends on the duration of the sender's transmissions, which, of course, depends on the bit-rate used. Note that in the case of 36Mbps (and also 48Mbps and 54Mbps), synchronization is at least as tolerant to interference as decoding data is. Thus, in these cases, there is no plateau.

III. PROBABILITY OF TRANSMISSION ERROR IN THE INTERFERENCE LIMITED REGIME

A. Analysis

When the received signal strength is sufficiently high, a transmission error can only occur when the sender's transmission overlaps with an interferer's transmission. In this case, there are three ways that the frame is lost. First, it is possible that synchronization will fail; we refer to this type of loss as a *sync error*. The event where synchronization fails is denoted with *SE*. Second, a frame can be lost if there are bit-errors that cannot be recovered from FEC (if FEC is used) . Since these bit-errors are detected with CRC errors, the event that a frame is lost

¹Our experiments have found no performance difference between 802.11b long and short preamble.

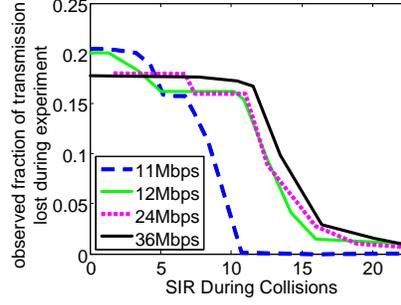


Fig. 2. The fraction of transmissions lost during the experiments described in Section II-B. Note that not every frame experiences a collision, and hence, even at very low SIR, not all frames are lost. The fraction of frames lost depends on the probability that a frame experiences interference, the synchronization scheme used by the physical layer, and the bit-rate used. For example, the probability of observing a transmission error at low SIR is equal to the probability of a collision occurring, i.e., every frame that experiences a collision is lost.

due to a CRC error is denoted with $CRCE$. The third way that packets are lost is that during the decoding of the packet header and payload, synchronization lock is lost. A loss of lock event is denoted with LL . Note that $CRCE$ can only occur if synchronization succeeded. Similarly, a loss of lock can only occur if the initial lock occurs. Furthermore, bit errors are irrelevant if there was a loss of lock. Thus, the events, $CRCE$, SE and LL are mutually exclusive. The objective of this section is to determine $P(CRCE|SIR)$ and $P(SE|SIR)$, where these probabilities denote the probability of the event occurring when a collision occurs. Prior work has found that when the SIR is low enough that $P(LL|SIR) > 0$, bit errors are common and hence $P(CRCE|SIR) \approx 1$. Thus, $P(LL|SIR)$ has no impact on the probability of frame error. For this reason, $P(LL|SIR)$ is not investigated.

Since the focus is on the interference limited regime and not on the noise limited regime, the received signal strength is high. Thus, an error can occur only when the sender is transmitting at the same time as the interferer. Specifically, a SE can only occur if the sender's synchronization phase overlaps with the interferer's transmission. There are different types of overlap. Specifically, the sender's synchronization phase could completely or partially overlap with the interferer's transmission. In the case of the experiments described in Section II, the probability that the sender's synchronization phase is entirely overlapped by the interferer's transmission is $(T_{\text{Sync,Interferer}} + T_{\text{Data,Interferer}} - T_{\text{Sync,Sender}}) / T_{\text{InterfererInterval}}$, where $T_{\text{Sync,Interferer}}$ is the duration of the interferer's synchronization, $T_{\text{Data,Interferer}}$ is the duration that the interferer's transmits data, $T_{\text{InterfererInterval}}$ is the time between interferer's transmissions. A partial overlap occurs when n synchronization symbols overlap with the interferer's

transmission. In the experiments described in Section II, the probability that exactly n symbols of the synchronization are overlapped with the interferer's transmission is $T_{\text{SyncSymbol,Sender}}/T_{\text{InterfererInterval}}$, where $T_{\text{SyncSymbol,Sender}}$ is the duration of a synchronization symbol. Thus, let $P(SE|SIR, n)$ be the probability of a sync error when n of the synchronization symbols are overlapped with the interferer's transmission. Then, the probability of observing a sync error in the experiments described in Section II is

$$\begin{aligned} P(\text{observed SE} | \text{SIR}) &= P(SE|SIR, N_{\text{SyncSym,Sender}}) \\ &\times \frac{T_{\text{Sync,Interferer}} + T_{\text{Data,Interferer}} - T_{\text{Sync,Sender}}}{T_{\text{InterfererInterval}}} \\ &+ \sum_{n=1}^{N_{\text{SyncSymbol,Sender}}-1} P(SE|SIR, n) \frac{T_{\text{SyncSymbol, Sender}}}{T_{\text{InterfererInterval}}}, \end{aligned}$$

where $P(\text{observed SE} | \text{SIR})$ is the fraction of frames transmitted in the experiment that resulted in synchronization error and $N_{\text{SyncSymbol,Sender}}$ is the number of symbols in the sender's synchronization. Since the synchronization phase has a short duration, the probability of a partial overlap is much smaller than the probability of complete overlap. Thus,

$$\begin{aligned} P(\text{observed SE} | \text{SIR}) & \tag{1} \\ & \approx P(SE|SIR) \frac{T_{\text{Sync,Interferer}} + T_{\text{Data,Interferer}} - T_{\text{Sync,Sender}}}{T_{\text{InterfererInterval}}}, \end{aligned}$$

where we drop the $N_{\text{SyncSymbol,Sender}}$ from $P(SE|SIR, N_{\text{SyncSymbol,Sender}})$ since a complete overlap is the only type of collision that is significant.

$P(SE|SIR)$ was estimated for 11 Mbps with long and short preamble and for 12 Mbps. For verification purposes, $P(SE|SIR)$ was estimated for several other bit-rates. Figure 3 shows $P(SE|SIR)$ for various types of synchronization types. It can be seen that 802.11b short preamble and long preamble behave nearly the same. This is surprising since the long preamble has 128 bits, while the short preamble has only 56. On the other hand, this behavior was also detected in the noise limited case (an analysis of the noise limited case is not included due to space limitations).

Next, the probability of bit-errors occurring is examined. When the channel is noise limited, the probability of successfully transmitting a packet is given by $(1 - p_{BE})^Z$, where p_{BE} is the probability of bit error and Z is the frame size. However, the conclusion of extensive measurements is that when the transmission is interference

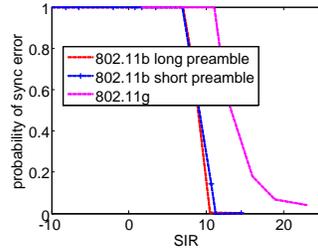


Fig. 3. Probability of synchronization failing as a function of SIR.

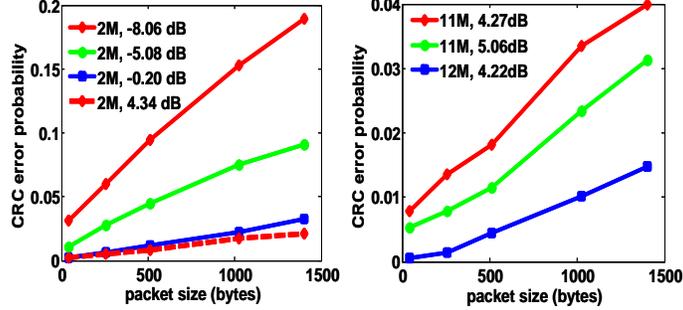


Fig. 4. Probability of CRC error observed during the experiments described in Section II as a function of frame size. The left-hand size is for 2Mbps while the right-hand size is for 11 and 12 Mbps, as indicated. Note that the probability of CRC error occurring is linear in the frame size.

limited, the probability of successfully transmitted a packet does not obey $(1 - p_{BE})^Z$. Rather, $P(CRCE|SIR)$ is independent of the frame length. On the other hand, if interference occurs randomly, then a longer frame is more likely to experience interference. Specifically, in the experiments described in Section II, the probability that a frame experiences interference is $(Z_{\text{Sender}}/R_{\text{Sender}})/T_{\text{InterfererInterval}}$, where Z_{Sender} is size of the sender's frame and R_{Sender} is the sender's transmitted bit-rate. Thus, if $P(CRCE|SIR)$ is independent of frame size, then

$$P(\text{Observed CRCE}|SIR) = P(CRCE|SIR) \frac{Z_{\text{Sender}}/R_{\text{Sender}}}{T_{\text{InterfererInterval}}}, \quad (2)$$

which is linear in Z . Figure 5 shows $P(\text{Observed CRCE}|SIR)$ as a function of Z for several bit-rates and confirms that $P(\text{Observed CRCE}|SIR)$ is linear in Z and hence $P(CRCE|SIR)$ is independent of the frame length.

Figure 5 shows the $P(CRCE|SIR)$ for $Z_{\text{Sender}} = 1464\text{B}$ and several modulation schemes. For reference, Figure 5 also show $P(SE|SIR)$.

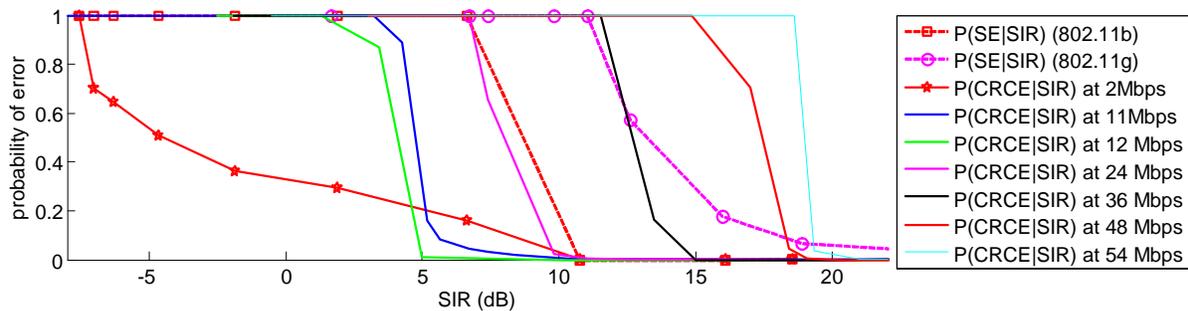


Fig. 5. Probability of error as a function of SIR for different types of modulation.

B. Discussion

Perhaps the most significant implications of Figure 5 is that when subject to interference, synchronization is considerably less robust to interference than decoding data. If the SNR is high and the short preamble is used, then 2, 5.5, and 11 Mbps all use the same synchronization. Thus, when subjected to interference, these modulation schemes will all have the same probability of error. Similarly, when subject to interference, 6, 9, 12, 18, 24, and 36 Mbps will all have the same probability of error. This behavior has implications for selecting a bit-rate. This issue is examined in more detail in the next section.

It is often assumed that the relationship between bit error and SNR is the same as the relationship between bit error and SIR or SNIR. To test this assumption, a larger number of experiments were performed to determine the noise limited performance of 802.11b/g. Table I used the results of these experiments to compare the SIR such that $P(CRCE|SIR) = 0.5$ to the SNR such that $P(CRCE|SNR) = 0.5$ under the assumption that the noise factor is 7dB, and hence the total noise (i.e., thermal noise plus the noise factor) is -93dB . Thus, excluding 2Mbps, 48Mbps, and 54Mbps, the assumption that relationship between SNR and bit-error is the same as the relationship between SIR and bit-error appears to approximately hold. Therefore, on the one hand, SIR can be treated as SNR when considering decoding the data. On the other hand, since synchronization is sensitive to interference, the performance of 802.11 frame decoding significantly depends on whether the noise is Gaussian or is interference. Thus, the relationship between SNR and transmissions error is not the same as the relationship between SIR and transmissions error.

The behavior at 2Mbps is intriguing and required a huge number of experiments to verify. On the one hand, at high SIR (e.g., 7dB), 2Mbps performed worst than 11Mbps. However, at -5dB , 2Mbps significantly outperforms

Bit rate (Mbps)	SNR	SIR
2	0	-5
11	5	4.75
12	5	4
24	7	8
36	12	13
48	15	18
54	16	19

TABLE I

SIR AND SNR REQUIRED FOR PROBABILITY OF ERROR OF 1/2.

11Mbps. The causes of this behavior are currently under investigation.

IV. PACKET ERROR WHEN SUBJECT TO INTERFERENCE AND NOISE

The discussion above focused on the interference limited regime. However, in some case, both interference and noise are important. For example, as will be discussed in Section V, even in the interference limited regime, the bit-rate that results in the shortest time to deliver a packet across a link is the fastest one that the link SNR can support. As a result, even in the high interference case, bit-rates should be selected so that noise is important. Intuitively, it is expected that as the strength of the received signal decreases, a higher SIR is required to correctly synchronize and decode the packet. It will be shown that this is often the case.

In order to explore the joint impact of noise and interference, the same experiments described in Section II were performed. However, in this case, an attenuator was adjusted to vary the received signal strength of the sender. Specifically, the received SNR of the sender was 21 dB, 12 dB, 9 dB, and 8 dB. Without interference, the packet error probability for a 1464 B packet transmitted at 12 Mbps with these SNRs was found to be 10^{-4} , 0.005, 0.01, and 0.02, respectively. Thus, at 21 dB, errors are quite rare, whereas at 8 dB, the probability of error is beginning to quickly increase. Using the techniques described in Section III, the probability of synchronization error and the probability of data error during interference was determined. Figure 6 shows these probabilities for the four received signal strengths. It can be observed that a higher SIR is required to synchronize when the received signal strength

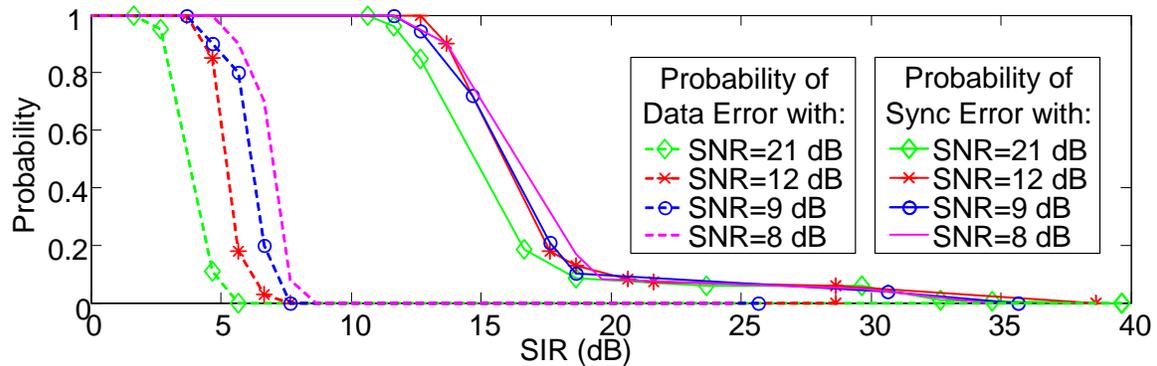


Fig. 6. Probability of synchronization error and the probability of data error as a function of SIR for different signal strengths (SNR).

is lower. However, while the received signal strength drops from 21 dB to 12 dB, the SIR required to synchronize with probability of 0.5 only increases by 1 dB. Moreover, when the signal strength drops from 12 dB to 8 dB or to 9 dB, the SIR required to synchronize with probability 0.5 changes by a small amount. On the other hand, when considering the ability to decode data in the face of interference, a decrease in the received signal strength has a noticeable impact on the SIR required to decode the data. In fact, the required SIR grows more quickly as the signal strength is decreased. For example, when the SNR drops by 9 dB from 21 dB to 12 dB, the SIR required to decode data with a probability of 0.5 increases by about 1.5 dB, whereas, when the SNR drops by only 1 dB, from 9 dB to 8 dB, the SIR required increases by nearly 1 dB.

Figure 6 shows that weak interference was a significant impact on transmission error. For example, observe that when the SIR is 17 dB and the SNR is 8 dB, the probability of synchronization error is approximately 0.5. However, if the SNR is 8 dB and the SIR is 17 dB, then the interference to noise ratio is -9 dB. That is, although the interference is 9 dB weaker than the noise, it is still has a significant impact on the ability to synchronize. To some extent, this behavior is surprising. Since the interference is weaker than the noise, one might expect that the synchronization error would be dominated by the SNR, not SIR. However, noise and interference are different. Specifically, noise is Gaussian and interference is not. Consequently, even if the noise power is the same as the interference power, the distributions are different. For example, the interference distribution might have a longer tail than a Gaussian distribution with the same variance, implying that the interference signal is more likely to reach large values than the noise. These large values can disrupt the synchronization process.

Figure 7 provides another view of the transmission error probability at different signal strengths. In this case,

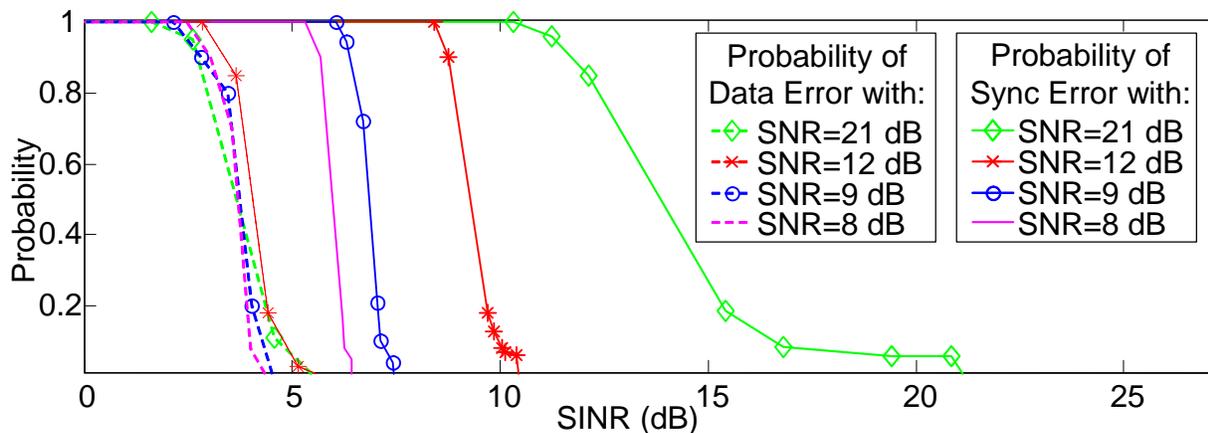


Fig. 7. Probability of synchronization error and the probability of data error as a function of SINR for different signal strengths (SNR).

the x-axis is the SINR (signal to interference plus noise ratio). It is widely believed that the ability to correctly transmit a packet is a function of the SINR. However, Figure 7 shows that synchronization is clearly not a function of SINR; if SNR is large, then a very large SINR is required to correctly synchronize, whereas if SNR is low, then the SINR required for synchronization is also low. On the other hand, as can be observed in Figure 7, the probability of correctly decoding data can be approximated by a single function of SINR. However, a more accurate model would have the SINR required for decoding data increasing as the SNR decreases.

V. BIT-RATE SELECTION IN THE INTERFERENCE LIMITED REGIME

The poor performance of 802.11b/g synchronization in the interference limited regime has significant implications for bit-rate selection. For example, if a channel has a SNR that can support 36Mbps, then reducing the bit-rate to 6, 9, 12, 18, or 24 will not impact the ability to synchronize and not increase robustness to interference. Furthermore, decreasing the bit-rate increases the transmission time and hence increases the possibility of experiencing interference. Hence, arguably, the highest bit-rate that the channel can support should be used regardless of the frequency of transmission errors. This contradicts the common approach to ARF that reduces bit-rate when packet losses are detected [7]. This section examines bit-rate selection in more detail.

Since interference can result in transmission errors, the behavior of the 802.11 backoff algorithm must be included. Specifically, for transmission at bit-rate R of a frame of size Z with probability of success p , the expected time to

transmit is

$$\begin{aligned}
 ETx(Z, R, p) & \\
 &= T_{slot} \left(\sum_{j=1}^6 \frac{(2^{j+4} - 1)}{2} (1 - p)^j + \frac{1023}{2} \sum_{j=7}^{\infty} (1 - p)^j \right) \\
 &+ \frac{1}{1 - p} (DIFS + SIFS + T_{ACK} \\
 &+ T_{Sync} + T_{PLCPHeader} + \frac{Z}{R}), \tag{3}
 \end{aligned}$$

where T_{slot} is the duration of a time-slot, DIFS and SIFS are the durations of the DCF and short interface frame spacings, respectively, T_{Sync} is the time it takes to synchronize, $T_{PLCPHeader}$ is the time to transmit the PLCP header, T_{ACK} is the duration required to send an ACK or, in case that the transmission failed, $T_{ACK} + DIFS$ is the time that the transmitter waits before beginning to decrement the backoff timer. Here, it is assumed that the ACK is transmitted at 2Mbps, which is the default value used 802.11 APs such as the Cisco 1240. The constants T_{Sync} and $T_{PLCPHeader}$ depend on whether 802.11b or g is used and in the case of 802.11b, they depend on whether the long preamble or short preamble is used. Here, it is assumed that the short preamble is used with 802.11b and the 802.11g PLCP header is used with 802.11g bit-rates. In (3), it is assumed that the initial value of the contention window is 31 (i.e., $2^5 - 1$) and the maximum value is 1023, i.e., $CW_{min} = 31$ and $CW_{max} = 1023$.

Under the assumption that RTS/CTS eliminates the majority of interference, the time to transmit a frame when RTS/CTS is used is

$$ETx(Z, R, 1) + 2 \times SIFS + T_{RTS} + T_{CTS}, \tag{4}$$

where T_{RTS} and T_{CTS} are the times to transmit the RTS and CTS, respectively. Again, in the results that follow, it is assume that control packets are sent at 2Mbps.

When the transmission is interference limited, the probability of successful transmission depends on how often a collision occurs. Suppose that the interference is such that the fraction of time that the channel is occupied is ρ and the duration of the silent times between interferer transmissions is exponentially distributed with mean λ . Thus, we have the following.

Proposition 1: The probability of a transmission failure is

$$P(\text{frame error} | SIR) = \tag{5}$$

$$(\rho + (1 - \rho)(1 - \exp(-\lambda T_{Sync}))) P(SE | SIR) \tag{6}$$

$$+ (1 - \rho) \exp(-\lambda T_{Sync}) \tag{7}$$

$$\times (1 - \exp(-\lambda(Z/R))) P(CRCE | SIR) \tag{8}$$

$$+ (1 - (1 - \rho) \exp(-\lambda T_{Sync})) \tag{9}$$

$$\times (1 - P(SE | SIR)) P(CRCE | SIR) \tag{10}$$

Proof: There are three ways in which a failure can occur. First, there could be a sync error. A sync error can only occur if the channel is busy during synchronization, which occurs with probability $(\rho + (1 - \rho)(1 - \exp(-\lambda T_{Sync})))$. To see this, note that the channel is busy at the beginning of the synchronization phase with probability ρ . If the channel is not busy at the beginning of the synchronization phase (which occurs with probability $1 - \rho$), then it will become busy during synchronization with probability $1 - \exp(-\lambda T_{Sync})$. Thus, (6) is the probability of synchronization failure.

The second way a transmission failure can occur is if the channel is not busy during synchronization, but becomes busy during the transmission of the data, and then a bit-errors occur. The probability of this type of error occurring is given in (7-8).

The final way that a transmission error can occur is when the channel is busy during synchronization, but the synchronization succeeds. However, the bit-errors occur during data transmission phase. The probability of this type of error is given in (9-10). ■

With (3) and (5), the time required to transmit a frame in the face of occasional interference can be determined. Figure 8 shows the average time to transmit a 40B and 1400B frame when the channel is purely interference limited and the channel utilization is $\rho = 0.3$. As expected, the fastest bit-rates are either 11Mbps, 36Mbps, 48Mbps, or 54Mbps. Which of these rates is best depends on the packet size and SIR. Note that 48Mbps is nearly the same as 54Mbps, hence, little performance is lost if only 11Mbps, 36Mbps and 54Mbps are considered. We see that for small size frames, and $SIR = 10$ dB, 11 Mbps is the best. Furthermore, at this SIR, 11Mbps is faster than 36 and

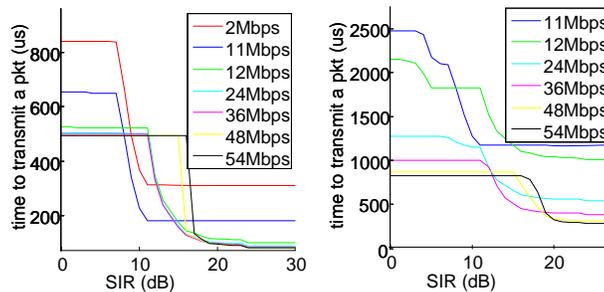


Fig. 8. Average time require to complete a transmission of a frame when $\rho = 0.3$ and $\lambda = 1/10 \text{ m.sec}^{-1}$ and when the frame is 40B (left-hand side) and 1400B (right-hand side).

54Mbps by a factor of five, or $300\mu s$. Similarly, when the frame is 1400B and the $SIR = 16 \text{ dB}$, 36 Mbps is the fastest and is faster than 54Mbps by a factor of 2 or $400\mu s$.

Figure 8 shows that in the interference limited case, the bit-rate that results in the smallest expected time to successful transmission depends on the frame size, the SIR, and the channel utilization. Figure 9 shows the bit-rate for a wide range of frame sizes, SIRs, and channel utilizations. Figure 9 also shows where RTS/CTS leads to the smallest expected transmission time. In 802.11, frames larger than $RTSThreshold$ use RTS/CTS. However, the optimal value the $RTSThreshold$ is unknown. Figure 9 shows that the optimal value of $RTSThreshold$ depends on the channel utilization (at $\rho \leq 0.5$, $RTSThreshold > 1500B$; $\rho = 0.2$, $RTSThreshold = 500B$ and at $\rho = 0.3$, $RTSThreshold = 0$).

VI. RELATED WORK

There has been considerable effort focused on understanding the behavior of 802.11 in the noise limited setting. In [1], measurements are used to explore the types of transmission errors at various bit-rates. Furthermore, [1] provides a useful explanation of many subtleties of transmission and decoding in 802.11. However, [1] does not examine the interference limited regime. In [8], the 802.11b was examined in a rooftop network setting where it was found that packet error was not closely correlated with SNR. One possible explanation of this behavior is that the channel suffered from delay spread. However, another explanation is that the transmissions also suffered from interference. One drawback of the work presented in this paper is that the combined impacts of delay spread and interference was not studied. Further investigation in this issue is required. In [9], a PHY receiver model of

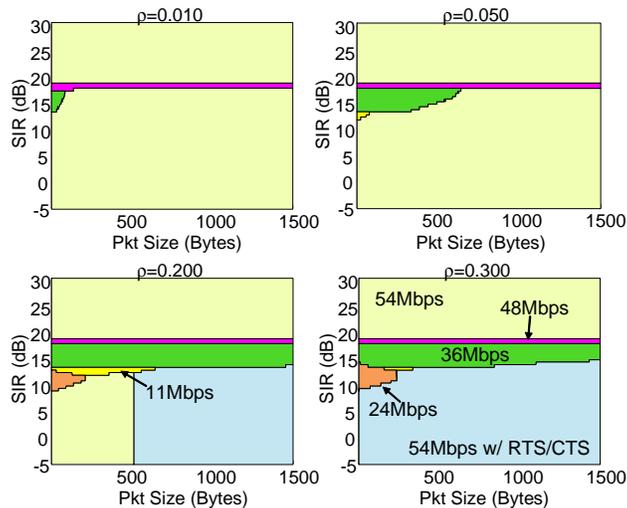


Fig. 9. Optimal Rate Regions. For a given combination of channel utilization, ρ , frame size, and SIR, there exists a particular bit-rate that results in the smallest delay. The above shows the optimal bit-rate in each region of the packet size/SIR plane. The region where a particular bit-rate is optimal is marked with a particular color. In these plots, $\lambda = 1/10 \text{ msec}^{-1}$.

decoding frames in the presence of interference was developed from measurement. While the developed model had some predictive power, it did not include the impact of synchronization, but rather relied on the replacing SNR with SNIR in the relationship between SNR and bit-error (See Section III-B above). In [10], capture was studied and a simple model for capture was developed from measurements. Contrary to findings presented here, [10] found that synchronization and decoding is possible as long as $SIR > 0 \text{ dB}$.

VII. CONCLUSIONS

This paper explored the behavior of 802.11b/g in the interference limited regime through a large number of laboratory experiments. Three observations were made. First, in the scenarios considered, synchronization error plays a critical role in the performance. As a result, in many cases, lowering bit-rate will not improve tolerance to interference. The implications that this observation has on bit-rate selection was investigated. A second observation is that during a collision, the probability of a bit error is independent of the frame size. This differs from the noise limited case where the probability of bit-error grows exponentially with the frame size. The third observation is that weak interference (e.g., with a signal power that is more than 10 dB weaker than the noise power) can have a significant impact on the ability successfully transmit packets.

DISCLAIMER

The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U. S. Government.

REFERENCES

- [1] G. Bianchi, F. Formisano, and D. Giustiniano, "802.11b/g link level measurements for an outdoor wireless campus network," in *WOWMOM*, Washington, DC, USA, 2006.
- [2] IEEE Std 802.11b, *Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications: Higher-Speed Physical Layer Extension in the 2.4 GHz Band*, 1999.
- [3] Cisco Systems, "Aironet 1240AG Series 802.11A/B/G Access Point," available at: <http://www.cisco.com/en/US/products/ps6521/products-data-sheet0900aecd8031c844.html>.
- [4] MADWiFi, "Multiband atheros driver for wifi," available at: <http://madwifi.org>.
- [5] Proxim Wireless, "ORiNOCO 11b/g Card," available at: <http://www.proxim.com/products/cp/pc.html>.
- [6] A. Communications, "AR5212," available at: <http://www.atheros.com/pt/AR5002GBulletin.htm>.
- [7] A. Kamerman and L. Monteban, *WaveLAN-II: A High-performance wireless LAN for the unlicensed band*. Bell Lab Technical Journal, 1997, pp. 118–133.
- [8] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris, "Link-level measurements from an 802.11b mesh network," in *SIGCOMM*, 2004.
- [9] C. Reis, R. Mahajan, M. Rodrig, D. Wetherall, and J. Zahorjan, "Measurement-based models of delivery and interference in static wireless networks," in *SIGCOMM*, 2006.
- [10] A. Kochut, A. Vasan, A. U. Shankar, and A. Agrawala, "Sniffing out the correct physical layer capture model in 802.11b," in *ICNP*, Washington, DC, USA, 2004.