

An Experimental Study on the Capture Effect in 802.11a Networks*

Jeongkeun Lee
Seoul National Univ.
jklee@mmlab.snu.ac.kr

Wonho Kim
Seoul National Univ.
whkim@mmlab.snu.ac.kr

Sung-Ju Lee
HP Labs.
sjlee@hp.com

Daehyung Jo
Seoul National Univ.
cdh@mmlab.snu.ac.kr

Jiho Ryu
Seoul National Univ.
jhryu@mmlab.snu.ac.kr

Taekyoung Kwon
Seoul National Univ.
tkkwon@snu.ac.kr

Yanghee Choi
Seoul National Univ.
yhchoi@snu.ac.kr

ABSTRACT

In wireless networks, a frame collision does not necessarily result in all the simultaneously transmitted frames being lost. Depending on the relative signal power and the arrival timing of the involved frames, one frame can survive the collision and be successfully received by the receiver. Using our IEEE 802.11a wireless network testbed, we carry out a measurement study that shows the terms and conditions (timing, power difference, bit rate) under which this *capture effect* takes place. Recent measurement work on the capture effect in 802.11 networks [10] argues that the stronger frame can be successfully decoded only in two cases: (1) The stronger frame arrives earlier than the weaker frame, or (2) the stronger frame arrives later than the weaker frame but within the preamble time of the weaker frame. However, our measurement shows that the stronger frame can be decoded correctly regardless of the timing relation with the weaker frame. In addition, when the stronger frame arrives later than the weaker frame's arrival, the physical layer capture exhibits two very distinct patterns based on whether the receiver has been successfully synchronized to the previous weak frame or not. In explaining the distinct cases we observe that the successful capture of a frame involved in a collision is determined through two stages: preamble detection and the frame body FCS check.

Categories and Subject Descriptors

C.4 [Performance of Systems]: Performance attributes

General Terms

Design, Experimentation, Measurement

*This work was supported in part by the National Information Society Agency (NIA) of Korea, the Hewlett-Packard Laboratories University Relations Program, and the Brain Korea 21 project of the Ministry of Education, Korea.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

WiNTECH'07, September 10, 2007, Montréal, Québec, Canada.
Copyright 2007 ACM 978-1-59593-738-4/07/0009 ...\$5.00.

Keywords

Capture Effect, Interference, IEEE 802.11

1. INTRODUCTION

Capture effect has been studied as an important factor in understanding the link throughput in random access (wireless) networks subject to potential collisions. With the capture effect, collision of two frames¹ in a shared channel may not destroy both frames. That is, one of the collided frames may still be successfully received by the receiver. In contention-based random access networks, it is known that the throughput with the capture effect can be more than double of the one without the capture effect [5, 14].

Wireless networks can benefit substantially from the capture effect and 802.11 networks are no exception. Understanding the capture effect will help us better analyze wireless interference, which is one of the major obstacle in achieving high capacity in wireless networks. Existing work on channel assignment, route selection, and traffic shaping models interference based on signal-to-interference power ratio (SIR) threshold between the two wireless links, but does not take the capture effect into account. With the better understanding of the capture effect, we can analyze the performance of 802.11 networks and design algorithms to improve the capacity.

This paper suggests a comprehensive set of experiments that reveals the conditions (i.e., frame timing, SIR, transmission bit rate) on which the capture effect takes place. The thorough investigation of capture effect conditions in this paper allows the research community to have more precise 802.11 wireless link models for analytic and simulation studies. To our best knowledge, this study is the first to provide the measurement-based SIR capture thresholds for the capture effect with all the 802.11a bit rates. Interestingly, for the capture effect to occur, different SIR thresholds are needed for different frame arrival timing relations between the two transmitters even when the other parameters (i.e., power, PHY bit rate) are the same.

There have been experimental studies on the capture effect in 802.11 networks [10, 8]. In the existing work, a collision is assumed to result in one of three cases: (a) both frames are corrupted, (b) the capture takes place when the stronger frame arrives at a receiver before the weaker one,

¹Two or more frames can be involved in a collision, but we limit the scope of our study to two frame collisions only.

or (c) the stronger frame can be successfully decoded when it arrives later than the weaker frame but within the weaker frame’s preamble time. We observed however that the capture effect occurs even when the stronger frame arrives later than the preamble time of the weaker frame in our Atheros chipset-based testbed. Although there have been studies [4, 18, 13] mentioning the preamble-time-independent 802.11 capture, we believe our work is the first *experimental* report on the preamble-time-independent 802.11 capture effect.

Another observation we make is that when the stronger frame arrives later than the weaker frame, the physical layer capture exhibits two distinct cases depending on whether the receiver has been successfully synchronized with the previous weaker frame. In explaining the two cases we learn that the successful capture of a frame involved in a collision is determined through two stages: detecting the preamble, and then checking the frame check sequence (FCS) of the payload. Although there are a few models of the two-stage capture [5, 6] for non-802.11 wireless networks, we believe that our work is the first 802.11 experimental work with the two distinct capture cases and the two capture stages.

The rest of this paper is organized as follows. Section 2 gives background of the capture mechanism. Section 3 describes how we set up a testbed to experiment the capture effect. Section 4 classifies the possible collision scenarios and reveal the SIR thresholds for the capture effect in each scenario. Section 5 investigates how the transmission bit rates of frames affect the physical layer capture. Related work is summarized in Section 6. Finally, concluding remarks are made in Section 7.

2. CAPTURE EFFECT

IEEE 802.11 does not pay special attention to the capture effect mainly to keep the design simple and cost low. Nonetheless, the capture effect occurs frequently in real deployments since all popular physical (PHY) layer technologies in 802.11 systems are non-linear modulation-based (e.g. DSSS in 802.11b, OFDM in 802.11a). In order to realize the physical layer capture, message in message (MIM) mode should be enabled in 802.11 radios [4, 18].

Suppose the first frame arrives at a receiver and after a while, the second frame arrives while the first frame reception is still ongoing. When the receiver detects the energy increase due to the second frame in the middle of the first frame’s reception, it can have two consequences depending on the energy increase level. If the energy increase is above a specified threshold, which we call *the capture threshold*, the receiver gives up the first frame and tries to receive the second frame in the MIM mode. In that case, it begins the retraining process (seeking to synchronize with and demodulate the second frame). If the energy increase is below the specified threshold, it sticks to the first frame sustaining the normal mode.

In summary, the capture effect necessitates retraining, which starts when the energy increase is the above the capture threshold. The retraining succeeds if the preamble (or carrier) of the new frame is detected. Following this *preamble detection stage*, the frame body (as well as frame header) must be successfully received in the presence of interference both from other transmissions and external noise sources. In other words, the frame check sequence (FCS) at the end of the frame must succeed: we call it a *FCS check stage*. The successful capture of a frame involved in a collision is de-

termined through two stages: preamble detection and frame body FCS check. This two-stage capture model is supported by our measurement-based observations in the following sections.

3. EXPERIMENTAL METHODOLOGY

To thoroughly investigate the capture effect, we try to observe (almost) every possible collision scenarios. The output of each collision comes from any combination of timing relation, signal strength, and PHY (or transmission) bit rate. We consider two transmitters whose transmissions create collisions and one receiver that exhibits the capture effect. The capture effect will be quantified by a frame reception ratio (FRR) which is defined as the ratio of the number of successfully received packets to the number of transmitted packets (of a particular transmitter).

Unlike other factors of the capture effect, testing every timing relation (or the difference between the arrival times of two colliding frames) is difficult. To experiment timing relation between the two transmitters, we arrange the transmitters to transmit independently (i.e., they cannot sense each other). The detailed topology design is explained in Subsection 3.2. We need a global view on transmission/reception times between the three nodes (the two transmitters and the receiver) to accurately analyze the timing relation and to distinguish between the capture cases based on arrival timing. Similar to other work [10, 8], we have two sniffer nodes, each of which is dedicated to monitor each transmitter. The observed transmission time logs of the sniffers are combined with the receiver’s reception time log, as described in Subsection 3.3.

In this section, we first introduce the 802.11 testbed used for the measurement study. We then explain the node topology and placement, followed by the description of how to achieve global timeline between the above three nodes. Finally, we describe traffic and MAC configuration.

3.1 Testbed

Our measurement study is carried out on the testbed consisting of small-form factor Soekris single-board computers running NetBSD. A mini-PCI 802.11a/b/g card [2] using Atheros AR5112 chipsets is installed in each node. Atheros cards and device drivers implement a majority of protocol functions of the 802.11 MAC protocol in the driver rather than in the hardware, so that users can specify/control most of the 802.11 MAC configuration parameters through device drivers. Using this, we have implemented a user-level utility that allows us to set various parameters such as MAC level retry limit and minimum/maximum contention window sizes, and to obtain hardware timestamps of received frames. Beacon transmission and noise-level calibration can also be enabled/disabled. Most MAC parameters/functions are accessed and controlled by *sysctl* utility. We also have implemented frame transmission power control and frame PHY bit rate control in the application program through *setsockopt* socket application programming interface (API). The received signal strength indicator (RSSI) and PHY bit rate of a received frame are reported through *getsockopt* API.

Each radio card in each node is mounted with two antennas; we artificially dedicate one antenna for transmission (TX) and the other for reception (RX). We disabled RX an-




Figure 1: Tesbed topology for the capture effect.

tenna diversity² to fix the RX antenna and stably measure the RSS. Using separate TX and RX antennas also helps us embody a sophisticated topology in our testbed.

We use 802.11a channel 52, which is verified to be free from external interferences. The channel status is monitored and analyzed by Airopeek [3] for verification. Our RSS measurement data exhibit negligible variation over time (mostly smaller than 1 dB), which is smaller than that of an 802.11b channel.

By default, Atheros cards report RSSs in terms of signal-to-noise ratio (SNR) and the noise power level is periodically calibrated. To further reduce the instability in the reported SNRs, we assume a fixed noise power value, -90 dBm and calculate the RSS by adding the fixed noise level to the reported SNR (in dBm scale). For that purpose, we disable the periodic noise calibration through our modified *sysctl* command. Given RSSs (in dBm scale) from the sender and the interferer, the signal-to-interference ratio (SIR) at the receiver is defined in dB scale as “RSS from the sender minus RSS from the interferer.”

3.2 Topology Design

In most experiments, we use five nodes: two transmitters, one receiver, and two sniffers. The sniffers are used to monitor the exact timing of the frame transmission from the transmitters. Fig. 1 depicts the placement of the five nodes on the testbed. One transmitter is assigned a role of an interferer. We focus on whether the receiver can successfully receive a frame from the sender despite the interferer’s concurrent transmission. For the ease of topology setup, all nodes are placed on a $5m \times 1m$ table. We place the nodes and the obstacles, configure antenna attachment/detachment, and control transmission power to create a topology to observe the capture effect. Specifically, the following conditions should be taken into account.

1. The condition of the links from the two transmitters³ to the receiver should be good enough to allow the receiver to successfully receive a frame from any transmitter even with the minimum transmission power. The reason is that we need to know the RSS of the interferer’s frame at the receiver. To do so, the interferer’s frame should be suc-

²This refers to the antenna switching capability by which a radio dynamically selects the better antenna for the frame reception.

³Depending on contexts, we use “the sender and the interferer” interchangeably with “the two transmitters.”




Figure 2: TSFT drift for 10 seconds.

cessfully received over the whole range of the transmission power.

2. In Fig. 1, S_1 is the sender and S_2 is the interferer. To observe the capture effect under various situations, the topology should generate a wide range of SIR values by controlling the transmission powers of the sender and the interferer. In our experiments, [-5 dB, 25 dB] SIR range can exhibit SIR thresholds in all possible capture scenarios.
3. The sender and the interferer should not sense each other in any case (i.e., even at their maximum transmission power). To confirm that they do not sense each other, we check the sender (or the interferer) alone can broadcast the same traffic rate (say, frames per second) compared with when both broadcast simultaneously.
4. The sniffers SN_1 and SN_2 must receive its corresponding transmitter S_1 and S_2 ’s frames correctly in all cases. In other words, SN_1 must receive all the S_1 ’s frames transmitted at the S_1 ’s minimum transmission power despite S_2 ’s simultaneous transmission at the S_2 ’s maximum transmission power.

In order to meet the first and the second conditions, the receiver must be placed close to both transmitters. The third and the fourth conditions however, require enough separation of the two transmitter-sniffer groups. That is why we placed an obstacle between the two transmitter-sniffer groups as shown in Fig. 1. The obstacle is composed of iron plates and piles of books that cause sufficient propagation loss between two groups. Despite the obstacle, both transmitters sense each other as long as the RX antenna is equipped. Therefore the RX antennas of the sender and the interferer are detached for the third condition (no carrier sensing between the transmitters). For the fourth condition, we place the S_1 ’s TX antenna right next to the SN_1 ’s RX antenna and place S_2 and SN_2 in the same manner. This way, we have an extremely high SIR (above 40 dB) between the sender (S_1) and the interferer (S_2) at SN_1 . Because of the high SIR ratio and the capture effect, SN_1 and SN_2 receive most of the frame transmissions (about 99.9%) from its corresponding transmitter S_1 and S_2 despite simultaneous transmissions.

3.3 Global Timeline

Our measurement study on the capture effect requires an accurate timing analysis of the two transmitters’ transmissions and the receiver’s receptions. We need to analyze whether any two frames from the two transmitters collide,

and if so, the difference of the arrival times of the colliding frames. We try to achieve microsecond-level accuracy in timing analysis. In our testbed, each sniffer records the timestamps of the received frames transmitted from its corresponding transmitter and the receiver also records the timestamps of the received or captured frames. We use Time Synchronization Function Timer (TSFT) in 802.11 as a timestamp. The TSFT is a 64-bit hardware counter that ticks every microsecond. According to our TSFT log data analysis, TSFT is stamped on each received packet after the arrival of the last bit of the frame. Thus, when we combine TSFT logs from the two sniffers and the receiver, a TSFT value must be interpreted as the end of transmission time in the air, and the start of transmission time must be calculated backward from the reported TSFT. Propagation time is negligible considering short distances.

According to the 802.11 standard, nodes in the same network (with the same Basic Service Set Identifier (BSSID)) try to keep synchronizing their TSFTs through beacon broadcasting.⁴ In each beacon message, the current TSFT of the beacon transmitter is recorded. In the infrastructure mode, only an AP sends beacons and all clients are synchronized to the AP's clock. On the other hand, nodes in the ad hoc mode send beacons in a contention-based manner, which does not guarantee TSFT synchronization between nodes in the same network. There are a number of solutions for time synchronization in 802.11 ad hoc networks [16, 15]; however, it is extremely difficult to apply those solutions to our testbed because synchronization functions (which are part of the MAC management sublayer) are not implemented in the Atheros driver.

Our approach hence is to disable beacons. To do so, after every node joins the same network in an ad hoc mode, we turn off beacon transmissions. In such environments (no synchronization among nodes), each node maintains its clock speed that is different from others. Based on our measurements however, the clock speed of each node is almost constant during one experiment session time of 30 seconds. Therefore, we measure the clock drift between any pair of the two sniffers and the receiver before each capture experiment session, and approximate it as a linearly increasing/decreasing function of time depending on clock speed difference. This way, we construct the global timeline of the frame transmissions/receptions among the nodes.

Fig. 2 plots the TSFT drift between SN_1 and the receiver measured during the first session and the last session of an experiment that took about seven hours in total, where each session took 10 seconds. Note that the drift value is in the order of microsecond while the measurement time unit is second. We measure the initial time gap at the beginning of each session and extract it from the following TSFT drift values, so that the TSFT drift in the graph begins at zero. From the plots, we verify the linear function approximation but we also learn that the clock drift speed must be measured periodically to mitigate the change of speed although the change is small ($1\mu\text{s}/\text{sec}$ change during seven hours).

In order to test the accuracy of this global timeline estimation, we measured the distributed inter-frame spacing (DIFS) time between two carrier-sensing senders and compared it with the DIFS time of 802.11a ($34\mu\text{s}$). We had

⁴The main purpose of TSFT synchronization in 802.11 systems is to enable synchronized sleep/wake among stations in power saving mode.

two senders transmitting the same sized (1000 bytes payload) broadcast packets at the same rate of 6 Mbps while two sniffers records TSFTs of their corresponding sender's transmissions into trace files. The traces are combined to make global timeline of all transmitted frames (about 20,000 frames) from both senders. From the global timeline data, we extracted DIFS time distribution, in which DIFS time appears mostly at $34\mu\text{s}$ and $35\mu\text{s}$ and only 0.3% of DIFS distribution is observed at $36\mu\text{s}$ and $37\mu\text{s}$. Considering $34\mu\text{s}$ DIFS time of 802.11a, we can conclude that the global timeline synchronization error is about $1\mu\text{s}$.

3.4 Traffic Generation and MAC Setting

We generate UDP application traffic at various rates (between 2.5 and 7 Mbps) and with different payload sizes (between 1000 and 1500 bytes) at each transmitter. Recall that we must produce a wide range of arrival time difference. In each experiment run, each transmitter is scheduled to transmit frames whose transmission times occupy between the quarter and the half of the airtime. This enables collisions to frequently occur and generates diverse overlapping transmission times. As the packet generation rate is fixed in each run, we add a random interval before delivering each packet to the MAC layer protocol for the same purpose.

We carry out the experimental study with both unicast and broadcast traffic. In the case of unicast, there are retransmissions, which make it hard to construct the global timeline. Hence, we set the MAC retry count to one, which eliminates any retransmission. Another reason to avoid retransmissions is to make application-level throughput equal to the MAC-level throughput.

4. TAXONOMY OF PHYSICAL LAYER CAPTURE

In this section, we provide the taxonomy of comprehensive capture scenarios in 802.11a networks based on our experiments with Atheros cards. Fig. 3 illustrates the four timing relations that characterize distinct capture cases. Throughout this section, 802.11a 6 Mbps PHY bit rate and 1000 byte payload size are used in both senders. Both the sender and the interferer send frames to the broadcast address. We analyze measurement logs and classify all the captured packets into one of the four capture cases below. Note that an FRR (Frame Reception Ratio) is measured by taking into account only colliding packets.

4.1 Sender First Capture (SF)

We first look at the case when a sender's frame arrives first at the receiver ahead of an interferer's frame as illustrated in Fig. 3.(a). We call this Sender First (SF) capture. Fig. 5.(a) plots the FRR of the sender's frames at the receiver with various SIR. The FRR transits sharply from zero to one near 0 dB SIR. This means that a sender's frames are received correctly when the two conditions are met: (1) it arrives ahead of the interferer's, and (2) its RSS is stronger than the interferer's frame at the receiver. This result is consistent with previous studies [10, 19]. Note that the sender's preamble is almost always successfully detected because the interferer's frame arrives later than the sender's in this SF case. Thus, the preamble detection stage is passed without interference but the FCS check stage suffers from interference. In Section 5, we show that the SIR threshold for the




Figure 3: Four capture cases.

SF capture increases as the PHY bit rate of the sender's frame increases. Hence the frame payload encoding rate (PHY rate) determines the SIR threshold.

Timing relation (i.e., the difference of the arrival times of the frames) does not make a significant difference in the SF case as shown in Fig. 4 which plots the FRR of sender's frames versus arrival time difference. Regardless of the arrival time difference, the first arriving sender's frame is captured by the receiver as long as the above two conditions are met. Here $\Delta t = (\text{the arrival time of the interferer's frame}) - (\text{the arrival time of the sender's frame})$. Although the transmission time of a frame is $1444 \mu\text{s}$, we plot FRR only up to $50 \mu\text{s}$ at the granularity of $5 \mu\text{s}$ due to space limit. We verify that FRR remains nearly one after $50 \mu\text{s}$. The FRR is averaged over the frames whose SIR range falls between 5 and 15 dB.⁵

4.2 Sender Last Capture with Interferer's Clear Frame (SLC)

When a sender's frame arrives later than an interferer's frame, we observe two distinct capture scenarios: (1) the first arriving frame (from the interferer) is being decoded until the sender's frame arrives as illustrated in Figs. 3 (b) and 3 (c), and (2) the first arriving frame (from the interferer) was garbled, which means the receiver deems the garbled interferer's frame as noise. We discuss case (1), called SLC here and case (2) in Section 4.3.

In the SLC case, the receiver requires a higher SIR value than the SF case to capture the sender's frame. The SIR threshold is around 10 dB as shown in Fig. 5.(b). Recall

⁵There is a small dip around $10\text{-}20\mu\text{s}$. As the preamble time of 802.11a is $16 \mu\text{s}$, the sender frame's reception seems a bit vulnerable to interference before the receiver completes its synchronization to the sender's preamble. However, it is somewhat strange that FRR is one in $0\text{-}5\mu\text{s}$ Δt




Figure 4: FRR vs. arrival time difference Δt for SF and SLC cases. Δt is defined as “the arrival time of the interferer's frame *minus* the arrival time of the sender's frame.”

that in the MIM mode, the receiver switches to the newer frame if the energy increase is above a specified threshold.

Fig. 4 shows the impact of the timing relation on the FRR. It shows that the stronger frames that arrives later than the preamble time ($16 \mu\text{s}$ in 802.11a) as shown in Fig. 3 (c) can be captured. We deem that this is due to Atheros chipsets which are believed to implement the MIM mode.

4.3 Sender Last Capture with Interferer's Garbled Frame (SLG)

To the best of our knowledge, this capture case has never been investigated or reported in the literature, where the first arriving frame (from the interferer) is already garbled. While the garbled frame's reception is ongoing, the sender's frame arrives at the receiver.

The interferer's frame can be garbled in the following two cases: (i) the interferer's frame is garbled by the previous frame⁶ of the sender as illustrated in Fig. 3 (d), and (ii) the interferer is located outside the communication range of the receiver and the receiver cannot correctly decode the interferer's frame. Recall from Section 3 that the first condition of the capture experiment is to have a good link condition from the interferer to the receiver. But case (ii) can frequently occur in wireless networks. We denote case (i) as SLG1 and case (ii) as SLG2.

In both cases (i) and (ii), as the frame is garbled, the receiver cannot decode it but is able to distinguish the garbled (interferer's) frame signal from white noise. In other words, the receiver can recognize the garbled frame's signal as an 802.11 signal although the receiver is not trained and synchronized to the garbled frame. The receiver's phase-locked loop continues to attempt synchronizing itself to the garbled signal but fails because its preamble is already missed. While the receiver is struggling to find and lock onto the interferer's garbled frame, the sender's frame arrives. These continuous synchronization attempts hinder the receiver from detecting (retraining) the preamble of newly arrived sender's frame. Thus as shown in Fig. 5 (c), the SIR for the SLG1 capture needs to be higher than that of the SLC case to correctly decode the preamble of sender's frame (last frame in Fig. 3 (d)). We also reason that in the SLC case, the receiver can cancel the interferer's frame more easily since it has already synchronized with the sender's frame.⁷ Hence,

⁶Actually, it could be transmitted by any nearby node.

⁷This “known-interference cancelation” has been well stud-




Figure 5: FRR vs. SIR with 6Mbps sender/interferer's transmissions. SIR is defined as "RSS from sender minus RSS from interferer."

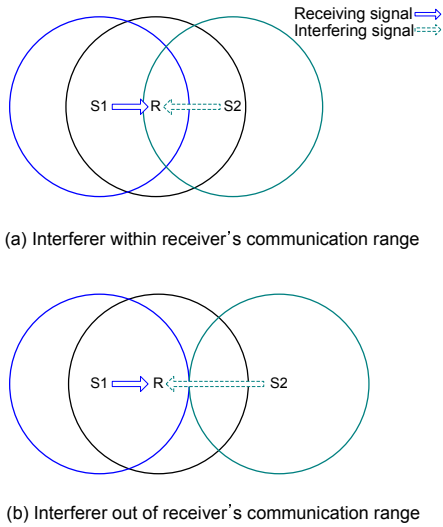


Figure 6: Two examples of hidden interference.

smaller SIR threshold is required in the SLC case compared to the SLG1: 10 dB for SLC, 24 dB for SLG1.

To test the SLG2 case, we arrange the nodes to form a topology sketched in Fig. 6 (b). Sender S_1 transmits a frame to receiver R while interferer S_2 also transmits a frame (destined to R or another receiver). Each circle represents the carrier sense range as well as the communication range assuming 802.11a radio. Because the senders are hidden from each other, their transmissions can take place simultaneously and collide. In this topology, only the SF and SLG2 cases can happen. Meanwhile, only SF, SLC, and SLG1 cases can take place in Fig. 6 (a). The capture experiments in both SLG1 and SLG2 cases produce the same pattern as shown in Fig. 5 (c).⁸

⁸ied [7].

⁸Note that S_2 in Fig. 6 (b) can cause more significant interference than S_2 in Fig. 6 (a) even though S_2 is placed farther from R in (b). Because SLG2 may be expected to happen more frequently in (b) than SLG1 in (a) and SLG capture requires higher SIR values (up to 24 dB) than SF and SLC, S_1 's transmissions are more difficult to be captured by R in (b) than in (a). The use of RTS/CTS can reduce hidden

5. MULTIPLE PHY BIT RATES

In all scenarios in our experiments, the PHY bit rate of the interferer does not make a difference on capture effect as long as other factors (e.g. the PHY bit rate of the sender) remain fixed. We thus focus on the impact of the bit rate of the sender's frame. Fig. 7 provides the FRR and SIR graphs with various sender rates for the three capture cases. FRR of each point in the graphs of Fig. 7 corresponds to one experiment session which lasts 30 to 50 seconds. The number of sender's attempted transmissions in one session varies over different bit rates but is at least 12,000.

In 802.11a mode of Atheros cards in use, as the PHY bit rate increases, the maximum transmission power decreases [2] which results in narrower SIR range we can generate. The minimum powers are the same for all bit rates. Based on our measurement, the same maximum transmission power limit is used for 6~24 Mbps bit rates; maximum powers for 36, 48, and 54 Mbps rates are 2, 4, and 6 dB lower than the maximum power for 6~24 Mbps rates, respectively. This made our experimental setting more difficult and the 54 Mbps plots include some spikes. Nonetheless, the tendency of SIR range change over PHY rates is easy to notice.

Before going into the detailed discussion with various PHY bit rates, we summarize our key findings for better understanding. First, the two capture stages (preamble detection stage and FCS check stage) have their own and different SIR thresholds for successful stage clear. Second, SIR threshold for preamble detection is not affected by the sender frame's encoding rate (preamble is always encoded at 6 Mbps) but is affected by the interferer frame's state (SLC or SLG). Third, SIR threshold for the FCS check increases as the bit rate increases. Fourth, the interferer's frame encoding rate does not show direct relation to the capture SIR threshold.

5.1 SIR Thresholds for Multiple Bit Rates

SIR in the SF case (Fig. 7 (a)) moves from 0 dB to more than 20 dB. Remember that in the SF case, only the FCS check stage matters because the sender's frame (thus its preamble) arrives first without interference. Interferer's frame arrives when the receiver is already locked to the sender's frame and is considered as white noise. As it hinders the decoding of sender's frame body, a higher SIR is required to decode a sender frame encoded at a higher bit rate. In interference in (a) but hardly in (b).

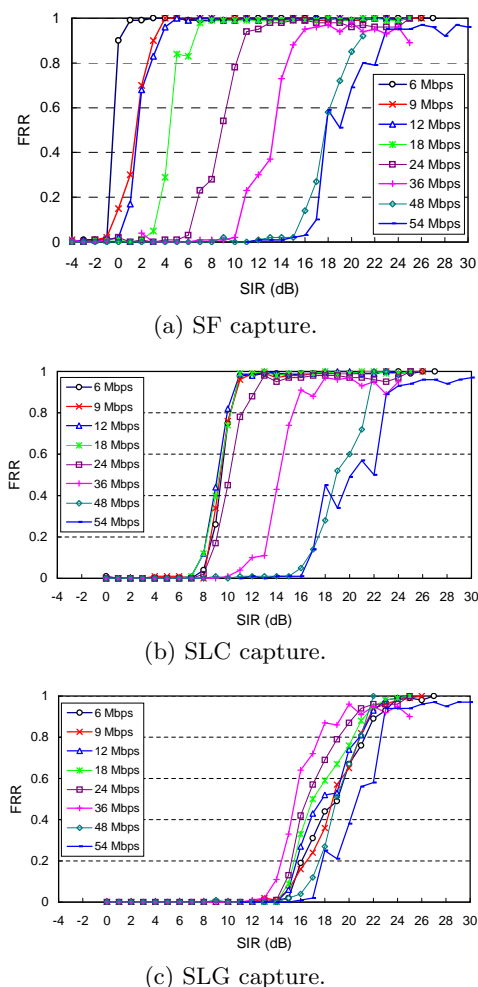


Figure 7: Sender’s PHY rate changes while the interferer’s PHY rate is fixed at 6Mbps.

summary, we can interpret the SIR thresholds of the SF case as the SIR thresholds for successful FCS check: for example, 0 dB for 6 Mbps and 23 dB for 54 Mbps.

As shown in Fig. 7 (b), SLC’s SIR threshold tends to hold at 10 dB until 24 Mbps and higher bit rates. Note that this bit rate is also when the SF’s SIR threshold gets larger than 10 dB. The 10 dB SLC threshold seems to be an intentionally designed hardware setting to protect the previous frame before deciding to capture the frame that arrives later, which is in line with the energy increase threshold in switching to the MIM mode (as mentioned in Section 2). It is also possible that the required SIR to detect (retrain to) the sender’s preamble in the presence of SLC interference is 10 dB. Because the preamble of an 802.11a frame is always encoded and sent at 6 Mbps regardless of the bit rates for frame body encoding, the required SIR for the preamble detection is expected to be the same whatever bit rate is used for the frame body. In this theory, the required SIRs for a successful FCS check of 6~18 Mbps frame body are smaller than or equal to 10 dB; if the preamble is successfully decoded, which means SIR is higher than 10 dB, FCS check stage does not matter for 6~18 Mbps. When the bit rate is lower than 24 Mbps, SIR threshold for a successful frame capture is determined

by the preamble detection stage. In contrast, when the bit rate is 24 Mbps or higher, SIR threshold for a successful frame capture is determined by the FCS check stage.

The SLG SIR pattern in Fig. 7 (c) does not show a major difference over multiple PHY rates. This contradicts the conventional belief that “the higher the PHY rate, the more vulnerable the communication.” Using the two stage capture model again, the preamble detection of the sender’s frame requires high SIR (up to 24 dB) because the receiver is already busy attempting to lock onto the interferer’s garbled frame when the sender’s preamble arrives at the receiver. This SIR of 24 dB is even higher than the SIR threshold for the highest bit rate’s successful FCS check, which is 23 dB for 54 Mbps. Thus, in the SLG capture, the preamble detection stage seems to determine the capture SIR threshold and not the FCS check stage.

Interestingly, as the PHY rate increases, the SIR graph shape of three capture cases become similar and converges to [15 dB, 20 dB]. From 36 Mbps and higher, SF and SLC patterns become similar and at 54 Mbps all three capture cases show similar curves.

5.2 Implication and Discussion

One of the major implications of the above SIR threshold graphs that when we construct a SIR-based network conflict graph we should use different SIR thresholds for different frame arrival timings (SF, SLC and SLG) even when the other parameters (power, PHY rate) are the same. For example, the SIR threshold for 6 Mbps frame’s complete capture (in other words, conflict-free SIR threshold) varies from 0 dB to 24 dB as the timing relation varies from SF to SLG. To our best knowledge, the frame arrival timing has not been considered in any of the conflict graph based interference model.

We discuss how we seek to ensure the confidence of the test results. We have examined the hardware (Atheros card) dependency. The observed SIR (or capture) thresholds are tested and validated over more than 10 Atheros radio cards. The SIR thresholds for the SF and the SLC cases and the SIR ranges for the SLG case are not dependent on individual Atheros cards. The measured SIR threshold difference is mostly less than 1dB.

6. RELATED WORK

There have been studies on the capture effect since analog FM modulation schemes, mostly in random access networks (e.g. [12]). The existing capture models for the 802.11 networks are described in [18] where overlapping transmission time and/or signal power difference (or SIR) are taken into account to enable the capture. Its SIR-based capture model is analyzed in [13]. The MIM mode and the triggered condition for the retraining process are proposed in [4]. According to our experiments, we believe this mechanism is embodied in the Atheros chipsets.

Recent experimental studies on the capture effect are made on 802.11 networks and sensor networks. The experimental results on the capture effect on 802.11b testbeds are presented in [10, 8, 9]. An experimental work on different capture cases (SF and SLIC) in 802.11 networks is presented in [10]. In [8], after providing empirical studies on how the capture effect causes unfairness, the authors propose to remedy the unfairness by adjusting the MAC parameters such as

the retransmission limit and power control (minimum contention window size, TxOp (Transmission Opportunity) and AIFS (Arbitration Inter-Frame Space) in 802.11e). However, as the carrier sense range is large in the experiments in [10, 8], all the nodes on the testbed can sense the other node's transmission. The capture effect is investigated only for the cases where the nodes select the same backoff slot. Although the nodes are supposed to start the transmissions exactly at the same moment, there is up to 20 microsecond difference between the arrival times at the receiver due to the RX/TX turnaround time delay and inherent uncertainties in the 802.11 firmware clock synchronization [10]. Hence, the timing relation between the sender and the interferer is not thoroughly investigated in the those studies. Another interesting study on the 802.11b testbed with the Prism chipsets is reported in [9], where a wireless emulator is devised to disable carrier sense to observe and insert controllable propagation loss. A wide range of the arrival time difference in collisions therefore can be investigated. However, only up to 100 μ s difference is investigated which is less than the preamble time of an 802.11b frame. Moreover, SLG capture case is not discussed.

Capture effects in sensor network testbeds that consist of Mica2 Motes with Chipcon CC1000 radios [1] are studied in [19, 17]. The capture behavior under various settings is tested in [17]. It is observed that the SIR threshold to trigger the capture may change over 6 dB range depending on the transmission powers. The capture experiments with multiple interferers is also conducted. The frames are however transmitted simultaneously and hence the timing relation is not investigated. Similar to our work, [19] reports that the capture can happen regardless of the timing relation between the two frames from the two transmitters. However, their measurement is conducted in Chipcon CC1000 radios and diverse capture scenarios are not considered.

The capture-aware interference models and estimation mechanism is discussed in [11], which also studies the relation of carrier sense and interference and their impact on the performance of two contending links.

7. CONCLUSION

Using our 802.11a testbed, we performed experimental studies on the capture effect and presented the precise terms and conditions (timing, SIR, and PHY bit rate) in which the receiver can successfully decode a packet from a sender despite the simultaneous transmission by another sender. Based on the arriving timing relation between the frames of the sender and the interferer, we classified the capture scenarios into three cases and we presented and analyzed the capture behavior in these different cases. We believe this is the first 802.11 measurement study that presents the capture scenario where the stronger frame arrives during the reception of the *garbled* weaker frame. This case is different from the scenario where the ongoing weaker frame transmission is not garbled and is recognized clearly by the receiver. We observed that the successful capture of a frame involved in a collision is determined through two stages: preamble detection and frame body FCS check. The analysis of our experiments will help us better understand interference/capture and its impact on throughput, which is crucial for the design of network protocols and the improvement of the wireless network capacity.

8. REFERENCES

- [1] Chipcon, <http://focus.ti.com/lit/ds/symlink/cc1000.pdf>.
- [2] Wistron NeWeb CM9, http://www.wnweb.com/wireless/wireless_mini-pci.htm.
- [3] AiroPeek NX wireless LAN analyzer, <http://www.wildpackets.com/>.
- [4] Jan Boer et al., Wireless LAN with Enhanced Capture Provision, US Patent 5987033, Nov. 16, 1999.
- [5] J. Arnbak. Capacity of slotted aloha in rayleigh fading channels. *IEEE J. Select. Areas Commun.*, SAC-5:261–269, Feb. 1987.
- [6] K. Cheun and S. Kim. Joint delay-power capture in spread-spectrum packet radio networks. *IEEE Trans. Commun.*, 46(4):450–453, 1998.
- [7] M. Costa. Writing on dirty paper. *IEEE Trans. Inform. Theory*, 24:374–377, 1978.
- [8] S. Ganu, K. Ramachandran, M. Gruteser, I. Seskar, and J. Deng. Methods for restoring mac layer fairness in ieee 802.11 networks with physical layer capture. In *Proc. ACM REALMAN'06*, Florence, Italy, May 2006.
- [9] G. Judd. Using Physical Layer Emulation to Understand and Improve Wireless Networks. Technical Report CMU-CS-06-164, Ph.D Thesis, School of Computer Science, CMU, Oct. 2006.
- [10] A. Kochut, A. Vasani, A. Shankar, and A. Agrawala. Sniffing out the correct physical layer capture model in 802.11b. In *Proc. IEEE ICNP*, Berlin, Germany, Oct. 2004.
- [11] J. Lee, S. Lee, W. Kim, D. Jo, T. Kwon, and Y. Choi. RSS-based Carrier Sensing and Interference Estimation in 802.11 Wireless Networks. In *Proc. IEEE SECON*, San Diego, CA, 2007.
- [12] J. J. Metzner. On improving utilization in ALOHA networks. *IEEE Trans. Commun.*, COM-24(4):447–448, 1976.
- [13] T. Nadeem, L. Ji, A. Agrawala, and J. Agre. Location Enhancement to IEEE 802.11 DCF. In *Proc. IEEE INFOCOM*, Miami, USA, Mar. 2005.
- [14] K. Pahlavan and A. H. Levesque. *Wireless Information Networks, 2nd Ed., Chapter 11*. Wiley, 2005.
- [15] C. H. Rentel and T. Kunz. A clock-sampling mutual network synchronization algorithm for wireless ad hoc networks. In *Proc. IEEE WCNC*, pages 638–644, New Orleans, USA, Mar. 2005.
- [16] K. Romer. Time Synchronization in Ad Hoc Networks. In *Proc. ACM MOBIHOC*, Long Beach, CA, Oct. 2001.
- [17] D. Son, B. Krishnamachari, and J. Heidemann. Experimental study of concurrent transmission in wireless sensor networks. In *Proc. ACM SenSys*, Nov. 2006.
- [18] C. Ware, J. Chicharo, and T. Wysocki. Simulation of Capture Behaviour in IEEE 802.11 Radio Modems. In *Proc. IEEE VTC'01 Fall*, Oct. 2001.
- [19] K. Whitehouse, A. Woo, F. Jiang, J. Polastre, and D. Culler. Exploiting the capture effect for collision detection and recovery. In *Embedded Networked Sensors, 2005. EmNetS-II. The Second IEEE Workshop on*, May 2005.