



Contents lists available at ScienceDirect

Computer Networks

journal homepage: www.elsevier.com/locate/comnet

Improved modeling of IEEE 802.11a PHY through fine-grained measurements [☆]

Jeongkeun Lee ^a, Jiho Ryu ^b, Sung-Ju Lee ^a, Ted Taekyoung Kwon ^{b,*}

^a Hewlett-Packard Laboratories, 1501 Page Mill Road, Palo Alto, CA 94304, USA

^b Seoul National University, Building 301, Room 553-1, 599 Gwanak-ro, Gwanak-gu, Seoul 151-742, Republic of Korea

ARTICLE INFO

Article history:

Available online xxx

Keywords:

IEEE 802.11
Physical layer capture
Interference
Carrier sense
Simulation model

ABSTRACT

In wireless networks, modeling of the physical layer behavior is an important yet difficult task. Modeling and estimating wireless interference is receiving great attention, and is crucial in a wireless network performance study. The physical layer capture, preamble detection, and carrier sense threshold are three key components that play important roles in successful frame reception in the presence of interference. Using our IEEE 802.11a wireless network testbed, we carry out a measurement study that reveals the detailed operation of each component and in particular we show the terms and conditions (interference timing, signal power difference, bitrate) under which a frame survives interference according to the preamble detection and capture logic. Based on the measurement study, we show that the operations of the three components in real IEEE 802.11a systems differ from those of popular simulators and present our modifications of the IEEE 802.11a PHY models to the NS-2 and QualNet network simulators. The modifications can be summarized as follows. (i) The current simulators' frame reception is based only on the received signal strength. However, real 802.11 systems can start frame reception only when the Signal-to-Interference Ratio (SIR) is high enough to detect the preamble. (ii) Different chipset vendors implement the frame reception and capture algorithms differently, resulting in different operations for the same event. We provide different simulation models for several popular chipset vendors and show the performance differences between the models. (iii) Based on the 802.11a standard setting and our testbed observation, we revise the simulator to set the carrier sense threshold higher than the receiver sensitivity rather than equal to the receiver sensitivity. We implement our modifications to the QualNet simulator and evaluate the impact of PHY model implementations on the wireless network performance; these result in an up to six times increase of net throughput.

© 2009 Elsevier B.V. All rights reserved.

1. Introduction

There have been extensive research efforts analyzing physical layer issues such as interference in wireless communications. The impact of interference on frame reception and throughput, however, still needs further investigation. One important problem is modeling the reception process at the physical layer in the presence of interference. A receiver can start receiving and decoding a transmitted frame only when it successfully identifies a predetermined signal pattern, which is preamble

[☆] This work was supported in part by NAP of Korea Research Council of Fundamental Science and Technology (KRCF) and the Ministry of Knowledge Economy, Korea, under the Information Technology Research Center support program supervised by the Institute of Information Technology Advancement. (Grant No. IITA-2009-C1090-0902-0006) The ICT at Seoul National University provides research facilities for this study.

* Corresponding author.

E-mail addresses: jklee@hp.com (J. Lee), jhyu@mmlab.snu.ac.kr (J. Ryu), sjlee@hp.com (S.-J. Lee), tkkwon@snu.ac.kr (T.T. Kwon).

detection. When the preamble is missed, the receiver can still sense there is an ongoing transmission if the received signal power is above a preconfigured level. This is energy detection. An IEEE 802.11 system can perform carrier sensing by using preamble detection and energy detection.

Recently, several studies examine physical layer capture and its ramifications [1–3]; the capture effect is a physical layer mechanism to deal with interference or collisions. In wireless networks, the medium is shared by all the nodes and a collision or interference takes place if two or more nodes in the vicinity transmit frames simultaneously. In most of the literature, collided frames are typically assumed to be garbled. With the capture effect however, one frame can survive the collision and be successfully received by the receiver depending on the relative signal power and the arrival timing of the involved frames. Thus, the throughput of the flows in wireless networks, which are subject to concurrent transmissions, is substantially influenced by the capture logic implementation. Hence we should take it into account in modeling the IEEE 802.11a PHY (physical layer).

From the previous measurement studies on IEEE 802.11 capture [1,2], we learned that the capture effect works differently depending on the 802.11 chipsets. Kochut et al. [2], from their experiments on the wireless cards with the Prism chipset [4], discovered that a stronger frame that arrives during the reception of a weaker frame can be captured if the stronger frame arrives *within* the weaker frame's preamble time. By contrast, with the Atheros chipset [5], the stronger frame can be captured even if it arrives *after* the weaker frame's preamble time [1]. In this paper, through our new testbed that enables more precise topology control, we observe more accurate conditions of the capture effect than our earlier work [1]. The thorough investigation of capture conditions in this paper can provide more precise and comprehensive 802.11 PHY models for analytic and simulation studies.

From measurement studies, we present IEEE 802.11a PHY models in which both carrier sensing and physical layer capture mechanisms are key building blocks. We observe notable discrepancy among the IEEE 802.11 standard, the widely used simulators (e.g., NS-2 [6] and QualNet [7]), and real 802.11 chipsets. We modify the QualNet simulator (version 3.9.5) to investigate how the different 802.11a PHY models affect the wireless network performance in various wireless network scenarios.

The contributions of this paper are summarized as follows. First, we present a detailed model of IEEE 802.11a PHY reception with a focus on the capture process and carrier sense. Second, we use a comprehensive set of experiments that reveals the conditions (i.e., collision timing, signal-to-interference-ratio, PHY bitrate) under which the capture effect takes place and identify two distinct capture models of different 802.11a chipsets. Third, through the testbed experiments, we discover and modify the simulators that do not correctly reflect the behavior of real 802.11a systems. Fourth, through the QualNet simulation, we show that different models of 802.11a capture and carrier sense yield substantially different network performance, up to more than six times increase of aggregate TCP throughput.

The rest of this paper is organized as follows: in Section 2, we describe the frame reception and capture procedures of the IEEE 802.11a PHY. Section 3 describes how we set up a testbed to experiment with the capture effect. From a comprehensive set of experiments, Sections 4 and 5 detail the factors that decide the capture effect and the carrier sense, respectively. Section 6 describes the current simulators' reception model and how the QualNet simulator should be modified. Section 7 presents a performance evaluation using the modified QualNet simulator. A survey of related work is given in Section 8. Section 9 concludes this paper.

2. 802.11a Frame reception, capture and carrier sense

In this section, we first overview the IEEE 802.11a PHY frame reception procedure and the carrier sense mechanism therein. We then elaborate on all the cases of the physical layer capture based on the IEEE 802.11 standard [8] and our testbed results.

2.1. 802.11a PHY receive procedure

As shown in Fig. 1, an 802.11a frame begins with a PLCP (Physical Layer Convergence Protocol) preamble that consists of OFDM training symbols. For a successful reception, while a receiver hears a transmitted PLCP preamble, the receiver should be able to (i) detect and measure the signal power which is equal to or stronger than the receiver's *minimum modulation and coding rate sensitivity* (RXsens)¹, (ii) perform Automatic Gain Control (AGC) to adapt to the received signal power, whose value falls in a wide range, and (iii) synchronize its timing with the training symbols. This is called the *preamble detection* process.

If the preamble detection is successful, the receiver now recognizes the start of a valid 802.11 frame transmission and searches for a PLCP header that follows the preamble. The PLCP header contains the information about the modulation/coding bitrate, the frame length and the parity bit. If the PLCP header reception is successful without any error detected by the parity, the receiver goes into a *receiving* state. Note that all these PHY operations are independent of the MAC header information. Thus, the PHY can lock onto the frame and go into a *receiving* state even when the received frame is a unicast frame destined to a different node.

After the PLCP header, the MAC header and data follow, and then a CRC is piggybacked after the MAC data for frame error checksum. The receiver generates a MAC CRC error if the MAC frame is corrupted. To summarize, in order to receive a frame successfully, the receiver must go through

¹ According to the 802.11 standard [8], the *receiver minimum input sensitivity* is defined as the received signal strength level at which the frame error rate of a 1000-octet frame is less than 10% and it is a bitrate-dependent value. The *minimum modulation and coding rate sensitivity* (RXsens) indicates the minimum sensitivity for the lowest modulation and coding bitrate, 6 Mbps. The required minimum sensitivity for 6 Mbps is -82 dBm by the standard but many 802.11 chipset vendors provide lower sensitivities for increased communication range and higher throughput. For example, the RXSens of Atheros 802.11a chipsets is known to be about -91 dBm [9] and it is also consistent with our measurements.




Fig. 1. 802.11 PHY frame format.

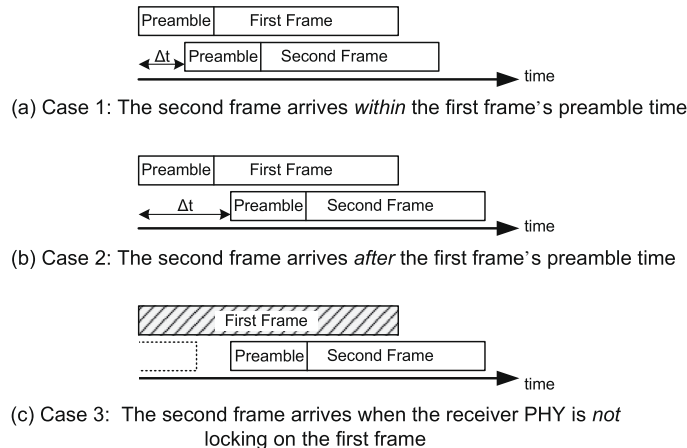


Fig. 2. Three capture cases.

three steps without error: preamble detection, PLCP header reception, and MAC CRC check.

In 802.11a, the PLCP header is always encoded and transmitted at the lowest bitrate, 6 Mbps, regardless of the bitrate of the MAC frame. The preamble training symbols are always the same for all frames. That is, the success of PLCP preamble detection and PLCP header reception are independent of the MAC frame bitrate; thus a higher SINR (Signal-to-Interference-and-Noise-Ratio) is required for a higher bitrate transmission of the MAC frame to successfully pass the MAC CRC check.

2.2. 802.11a Carrier sense

As described above, the 802.11a PHY goes into the *receiving* state after successfully detecting the preamble and receiving the PLCP header. In this case, the PHY also holds the carrier sense busy for the duration of the ongoing frame. If the receiver detects the signal energy but the preamble portion is missed, according to the IEEE 802.11a standard [8], the receiver holds the carrier sense busy for any signal 20 dB above the receiver's RXSens.

To summarize, (i) the 802.11a PHY determines the busy medium by the PLCP preamble (and header) detection (PD) or the energy detection (ED), and (ii) the 802.11a standard defines the ED threshold (EDThres) to be 20 dB higher than the RXSens.²

2.3. Capture effect in 802.11a PHY

In this section, based on our measurements and the prior studies on capture timing and SINR threshold [1–3],

² In the IEEE 802.11 standard [8], the terminology 'ED threshold' of carrier sense is used only in the sections for 11b and 11g and no specific terminology is used to indicate the threshold in the section for 11a. However, we use the same terminology 'ED threshold' in this paper for ease of presentation.

we classify the capture effect into three cases and describe how the 802.11 PHY processes frame capture in each case.

2.3.1. Case 1: the second frame arrives *within* the first frame's preamble time

As shown in the example of Fig. 2a, suppose the first frame arrives at a receiver and subsequently, the second frame arrives while the receiver is still receiving the first frame's preamble. If the second frame's signal power is strong enough for the receiver to detect an energy increase (over the first frame's signal power) above a predefined threshold, which we call *preamble capture threshold* or simply *capture threshold*, the receiver drops the first frame's preamble and tries to detect the second frame's preamble. We call it *SC (Second frame Capture)*. If the receiver has not completely detected the first frame yet (i.e., PLCP preamble is not fully processed), the SC can happen even when the signal increase due to the second frame over the first frame is below the *capture threshold* (Section 4.1.2). For ease of presentation, we define *capture* as the PHY operation of selecting one frame to receive between the frames in a collision (the first or the second) regardless of the success of the selected (or captured) frame reception (i.e., regardless of the MAC CRC error). In other words, the capture decision is made at the time of the second frame arrival, but the success of the captured frame reception is determined later.³

If the second frame's SINR (i.e., the ratio of the second frame's signal power to the first frame's signal power plus noise) is high enough to decode the second frame's preamble, the PLCP header and the MAC data without error, the second frame is successfully captured. Kochut et al. [2] showed that the Second frame Capture within preamble time occurs with Prism [4] chipset 802.11 cards and we

³ For example, if another (third) frame arrives during the reception of the captured second frame and the receiver decides to capture the third frame, the second frame reception fails.

will also show this capture occurs with Atheros chipset [5] 802.11 cards in Section 4.

If the power increase due to the second frame's preamble is too small to detect or is below the capture threshold, the receiver retains its lock onto the first frame's preamble and tries to synchronize with it. We call it *FC (First frame Capture)*. However, as the second frame's signal power increases (i.e., interference power increases from the standpoint of the first frame), the captured first frame's PLCP header and/or MAC frame can be corrupted. The First frame Capture results in a successful reception when the first frame ends without PLCP header error or the MAC CRC error.

If transmission times of three or more frames overlap, the 'first' and 'second' frames could be a misnomer. To address those situations, a frame that the receiver has currently locked on to is called "the first frame." The next frame that arrives during the reception of "the first frame" is called "the second frame" for ease of presentation.

2.3.2. Case 2: the second frame arrives after the first frame's preamble time

In Case 2, the second frame arrives after the first frame's preamble time as shown in Fig. 2b. The receiver has already synchronized its timing with the first frame and the first frame's PLCP header has passed; thus, the receiver is in the *receiving* state. In this case, in order to capture the second frame, Message-In-Message (MIM) mode should be implemented in 802.11 PHY [10,11]. In the MIM mode, if the power increase due to the second frame is above the capture threshold, the receiver drops the first frame and tries to synchronize its timing with the second frame preamble; i.e., SC occurs even when the second frame arrives after the first frame preamble. If the MIM mode is *not* implemented or if the power increase is below the *capture threshold*, the receiver retains its reception of the first frame; i.e., FC occurs. In either case (SC or FC), the uncaptured (or unselected) frame signal increases the interference power and can affect the MAC CRC check of the captured frame.

Note that the PHY's locking onto a signal and going into a *receiving* state are distinct operations. The receiver locks onto a signal in the early stage of preamble detection process and it goes into a *receiving* state when it completes the preamble detection and the PLCP header reception without error. The *capture threshold* comes into a play when the receiver has locked onto a previous frame and the MIM mode is required to capture the second frame when the receiver is in a *receiving* state due to the first frame.

The implementation of MIM mode is chipset-dependent. In Section 4, we will show that the second frame that arrives later than the first frame preamble time can be captured with Atheros chipset cards that are believed to implement the MIM mode. However, when we test with the Prism chipset cards, the Second frame Capture does not happen when the second frame arrives after the first frame preamble time even when the second frame is much stronger than the first frame. Kochut et al. [2] carry out their measurement study with Prism chipset cards and report that the Second frame Capture can happen only when it arrives *within* the preamble of the first frame.

2.3.3. Case 3: the second frame arrives when the receiver is *not* receiving the first frame

In Cases 1 and 2, if the receiver has already locked onto the first frame, the energy increase at the receiver due to the newly arriving second frame must be above the *capture threshold* for the receiver to drop the first frame and switch to the second frame. In Case 3 however, the receiver's PHY is not occupied by the first frame's signal, and it can start the reception of the second frame regardless of the *capture threshold*, as long as the SINR for the second frame is high enough to detect the second frame preamble in the presence of the first frame's interference.

There are two scenarios in Case 3, where the receiver fails to lock onto the first frame. First, the first frame is not captured due to another frame as illustrated by the dotted frame in Fig. 2c. Second, the transmitter of the first frame is located outside the communication range of the receiver and the receiver cannot detect and/or synchronize with the first frame preamble. We call this capture in Case 3 *SC-GI (Second frame Capture with Garbled Interferer)* because the first frame is garbled from the receiver's perspective.

In the following sections, we will show measurement results that support the capture and carrier sense mechanisms described in this section and will provide the actual model of off-the-shelf 802.11 wireless cards.

3. Experimental methodology

In this section, we explain the testbed settings for our capture experiments.⁴ There are two methods to construct the testbed topology to carry out the capture experiments. One way is to use wires (e.g. coaxial cables) and the other way is to use an air interface; we need to control the node placement, antenna attachment/detachment, and transmission power control (e.g. refer to our previous testbed study [1] for details.). In this paper, we use wires to simplify the process of creating a desired topology. The precision and reliability of this approach has been verified by previous work [12,13].

To thoroughly investigate the 802.11 PHY capture effect and carrier sense, we built a testbed that enables us to observe (almost) every possible capture (or collision) scenario. The output of each capture depends on the combination of timing relation, signal strength, and PHY bitrate. We consider two transmitters whose transmissions create collisions (their transmission duration times overlap) and one receiver that exhibits the capture effect. 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 with every possible 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 Section 3.2. We need a global view on transmission/reception times between the three nodes (the two transmitters and the receiver) to

⁴ The setting for the carrier sense experiments is relatively simple, and will be described in Section 7.2.

accurately analyze the timing relation. Similar to prior studies [2,14], we use two sniffer nodes, each of which is dedicated to monitor each transmitter. We construct a global timeline of the frame transmissions/receptions among the three nodes by combining the observed transmission time logs from the sniffers and the receiver's reception time log. We use the Time Synchronization Function Timer (TSFT) in 802.11 as a hardware timestamp and achieve microsecond-level accuracy in timing analysis [1].

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 [15] using the Atheros AR5112 chipset is installed in each node. Each 802.11 radio card in each node has two antenna ports; we artificially dedicate one antenna for transmission (TX) and the other for reception (RX). We disabled RX antenna diversity⁵ to fix the RX antenna and consistently measure the received signal strength (RSS). Using separate TX and RX antennas also helps us build a sophisticated topology in our testbed.

By modifying the NetBSD kernel and Atheros device driver, we implemented a user-level utility that allows us to set various parameters such as the MAC level retry limit and the minimum/maximum contention window size. We performed per-packet transmission power control and PHY bitrate control in the application program through *setsockopt* socket application programming interface (API). We also enabled our measurement application to obtain the RSS, noise power level, PHY bitrate, and hardware timestamp of each received packet.

We use an 802.11a channel, which is verified to be free from external interference. The channel status is monitored and analyzed by AiropEEK [16] for verification. Our RSS measurement data exhibit negligible variation of noise power over time (mostly less than 1 dB).

3.2. Topology design

In most experiments, we use five nodes: two transmitters, one receiver, and two sniffers. One transmitter is assigned a role of an interferer and two sniffers are used to monitor the exact timing of the frame transmission from the transmitters. The nodes are connected to each other through cables attached to either the TX or RX antenna port of their 802.11 cards. Fig. 3 depicts the topology of the five nodes on the testbed. The TX/RX sign indicates which antenna port is connected to a cable. Specifically, the following conditions should be taken into account to create a desired topology.

- (1) The condition of the links from the two transmitters⁶ to the receiver should be good enough to allow the receiver to successfully receive a frame from any

⁵ This refers to the antenna switching capability by which a radio dynamically selects the better antenna for frame reception.

⁶ Depending on context, we use “the sender and the interferer” interchangeably with “the two transmitters”.

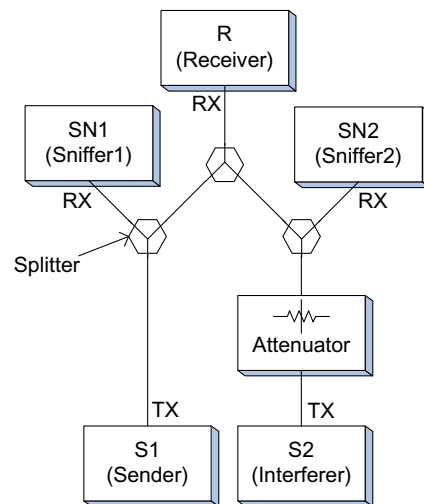


Fig. 3. Testbed topology for the capture effect.

transmitter even with the minimum transmission power. The reason is that we need to obtain the RSS of the interferer's frame at the receiver.

- (2) In Fig. 3, 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 (in our experiments, $[-5 \text{ dB}, 25 \text{ dB}]$ is sufficient) by controlling the transmission powers of the sender and the interferer. We also use a signal attenuator to easily generate the wide SIR range.
- (3) The sender and the interferer should not sense each other in any case to generate every timing relation (i.e., the difference between transmission times). To confirm that they do not sense each other, we check that the sender (or the interferer) alone can broadcast the same traffic rate (say, frames per second) compared to when both broadcast simultaneously.
- (4) The sniffers SN_1 and SN_2 must always receive its corresponding transmitters S_1 's and S_2 's frames, respectively. For instance, SN_1 must receive all S_1 's frames transmitted at S_1 's minimum transmission power despite S_2 's simultaneous transmission at S_2 's maximum transmission power.

3.3. Driver settings and traffic generation

We carry out the experimental study with both broadcast and unicast 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. Throughout this paper, we do not consider RTS/CTS since the capture of RTS frames is not different from that of MAC data frames [2].

There is an algorithm called ANI (Ambient Noise Immunity) that resides in the HAL (Hardware Abstraction Layer)

of the Atheros driver. When operating in ad hoc mode with OFDM PHY, ANI operates in a faulty manner resulting in reduced receive sensitivity [17,18], which affects preamble detection. We observed that ANI reduces carrier sense sensitivity as the preamble detection is a part of the carrier sense mechanism. This sensitivity issue has been actively discussed in the MadWifi community [17] and its bug patch was released [19]. We have upgraded our testbed with the latest Atheros HAL and driver and turn off ANI.

We generate UDP application traffic at various rates (between 2.5 and 15 Mbps) and with payload size of 1000 bytes at both transmitters. The UDP traffic generation rate is determined based on the PHY bitrate used for each experiment run. Recall that we must produce a wide range of arrival time difference. In each experiment run, we schedule the frequency of transmissions so that transmission times occupy between a quarter and a half of the total air time. This makes collisions occur frequently and generates various 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.

4. Measurement-based modeling of preamble detection and capture

In order to conduct the PHY modeling of an off-the-shelf 802.11a testbed thoroughly, we examined more than 30 million packets transmitted at various 802.11a bitrates and classified them into the three capture cases based on the collision timing relations between the two transmitted frames. Throughout this section, both the sender and the interferer transmit to the broadcast address unless a different setting is specified. We call a frame from the sender a FoI (Frame of Interest) and will explain the experimental results from the viewpoint of FoI.

4.1. Effect of SIR and collision timing on capture performance

We analyze measurement logs of transmissions with the bitrate fixed to 6 Mbps and classify all the frame collisions into one of the three collision cases: (1) FoI arrives at the receiver ahead of the interferer's frame, (2) FoI arrives later than the interferer's frame, and (3) FoI arrives later than the interferer's frame, which was already garbled as explained in Section 2.3.3. That is, the receiver could not synchronize timing with the interferer's frame. Thus, the three cases correspond to FC (First frame Capture), SC (Second frame Capture), and SC-GI (Second frame Capture with Garbled Interferer) from the FoI's perspective, respectively. We analyze the required SIR (Signal-to-Interference Ratio) of each capture case. As reference data, we also plot the FRR (Frame Reception Ratio) vs. SNR (Signal-to-Noise-Ratio) curve when there is no interference. That is, we increase the transmission power of the sender in the presence of the ambient noise power only in order to compare SIR and SNR. The case of no interference is called Clear Channel Reception (CCR).

FRR is defined as the ratio of the number of successfully received packets at the receiver to the total number of

transmitted packets from the sender including the packets dropped due to collision with the concurrent transmissions from the other transmitter. Given RSSs (in dBm scale) of the FoI (from the sender) and the interference frame at the receiver, the SIR is defined in dB scale as "RSS of FoI minus RSS of the interference frame". Because the reported noise level by the driver is -96 dBm or smaller and the signal power from the sender and interferer used in our experiments ranges -80 to -50 dBm, the noise power is negligible (40–40,000 times smaller than signal power) and we can deem SIR to be the same as SINR (Signal-to-Interference-and-Noise-Ratio) throughout this paper.

4.1.1. First frame Capture (FC)

We first look at the FC case where a FoI arrives first at the receiver ahead of the interferer's frame. Fig. 4 plots the FRR of the FoI at the receiver with various SIR. To calculate FRR, we consider all the FC collisions that have SIR in the range of $[X - 0.5, X + 0.5]$, where X is an integer of SIR, and the FRR at the SIR X is the ratio between the number of successfully received frames from the sender at the receiver and the total transmissions from the sender.

Fig. 4 shows that the FRR of FC reaches 0.9 near 1 dB and FC requires lower SIR than the other capture cases and the receiver sensitivity of CCR. The fact that the receiver has already locked onto the FoI at the time of interference frame's arrival seems to help the receiver to overcome the interference signal with relatively low SIR.

Timing relations (i.e., the difference of the arrival times of the frames) of the FC cases are shown in Fig. 5a. The x -axis shows the arrival time difference $\Delta t =$ (the arrival time of the second frame) $-$ (the arrival time of the first frame (FoI, in this case)) in the interval unit of $3.2 \mu\text{s}$ which is the 802.11a OFDM symbol duration [8], and the y -axis shows the FoI's SIR against the interference frame.

Fig. 5a shows three distinct performance regions of the FC capture with respect to Δt : $[0, 6.4] \mu\text{s}$, $(6.4, 16] \mu\text{s}$, and larger than $16 \mu\text{s}$. When the interference frame arrives within $6.4 \mu\text{s}$ after the FoI's arrival, about 4 dB SIR is required for the receiver to capture FoI with an 80% or more success ratio. (Note that the 4 dB SIR requirement is similar to that of SC-GI and CCR.) As Δt increases up to $16 \mu\text{s}$, the SIR requirement gradually decreases to 1 dB. The SIR requirement value for a given FRR converges after $16 \mu\text{s}$. This phenomenon is explained by the 802.11a OFDM

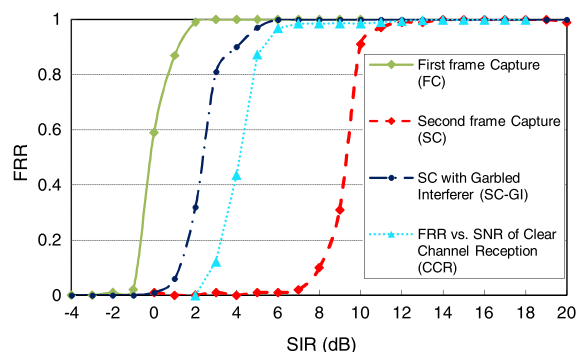


Fig. 4. FRR vs. SIR with 6 Mbps sender/interferer's transmissions.

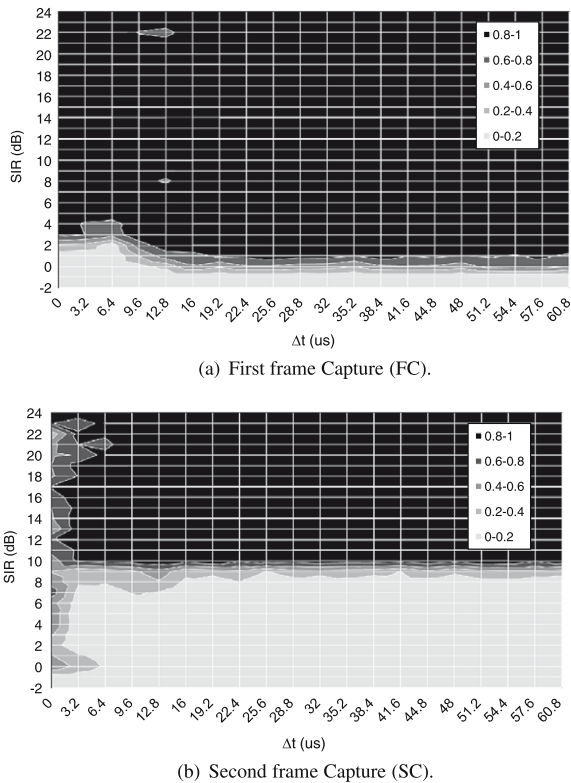


Fig. 5. FRR as a function of the SIR and the arrival time difference Δt for 6 Mbps FC and SC cases. Δt is “the arrival time of the second frame minus the arrival time of the first frame”.

preamble structure [8], which consists of three parts: in the entire $16 \mu s$ preamble time, the first $5.6 \mu s$ is used for signal detection and Automatic Gain Control (AGC), the second part from $5.6 \mu s$ to $8 \mu s$ is used for coarse frequency offset estimation and timing synchronization, and the remaining part from $8 \mu s$ to $16 \mu s$ is for channel and fine frequency offset estimation. Thus, as the interference frame arrives later, the receiver requires less SIR threshold to detect and synchronize to the FoI signal. If the interfering frame arrives after preamble detection of the FoI is completed ($16 \mu s$), the receiver is able to capture the FoI as long as its signal is only 1 dB (1.26 times) stronger than the interfering frame. Our measurement results for different PHY bitrates show a similar tendency for FC capture; as the interfering frame arrives within the preamble time (zero to $16 \mu s$), the required SIR threshold for successful capture is around 1–2 dB.

Although the transmission time of a 1000-byte payload frame is $1444 \mu s$, we plot FRR only up to $60.8 \mu s$ to observe the early phase of frame reception in detail. We also observed that the FRR vs. SIR pattern remains the same after $60.8 \mu s$: the time duration that the FoI’s transmission overlaps with the interfering frame does not change the SIR threshold to survive the collision. From this, we can also infer that the SIR capture threshold is independent of packet length although the longer packet will have more chance of collisions.

4.1.2. Second frame Capture (SC)

As mentioned in Section 2.3.2, the Atheros chipset we use for measurements implements MIM (Message-In-Message) mode and it supports the Second frame Capture (SC) irrespective of the arrival timing of the FoI – before or after the preamble time of the interfering frame. Fig. 5b verifies it by showing the constant 10 dB SIR threshold and no noticeable change in FRR until the end of preamble time, $16 \mu s$.

As shown in Fig. 4, the receiver requires a higher SIR value to success the SC capture than the other capture cases. The SIR threshold to achieve FRR = 0.9 is approximately 10 dB. We believe that this higher SIR threshold for SC is related to the *capture threshold*, which will be detailed in Section 4.2.

4.1.3. Second frame Capture with Garbled Interferer (SC-GI)

In the SC-GI case, the FoI arrives at the receiver while the garbled frame’s transmission is ongoing. Because the receiver does not lock onto the first (garbled) frame at the time of the FoI’s arrival, the receiver can start receiving the FoI regardless of the *capture threshold*, as long as the SINR for the FoI is high enough. The interferer’s frame (or the first frame) is regarded as noise power by the receiver; thus, the SIR curve in SC-GI is similar to the SNR curve in CCR.

As we discussed in Section 2.3, the interferer’s frame (first frame) can be garbled in the following two cases: (i), the first frame can be un-captured due to another frame as illustrated by the dotted frame in Fig. 2c, and (ii) the transmitter of the first frame is located outside the communication range of the receiver and hence the receiver cannot detect the first frame preamble. Case (ii) can frequently occur in wireless networks. To test case (ii), we arrange the nodes to form a topology sketched in Fig. 6. 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 an 802.11a radio. Because the senders are hidden from each other, their transmissions can take place concurrently and collide. The capture experiments in this topology setting produce the same capture threshold as case (i) in Fig. 4 (SC-GI).

4.2. Multiple PHY bitrates

In all scenarios in our experiments, we observe that the PHY bitrate of the interferer’s frame does not have any meaningful impact on capture effect. We thus focus on the impact of the bitrate of the sender’s frame only. Fig. 7




Fig. 6. Interferer is out of receiver’s communication range.

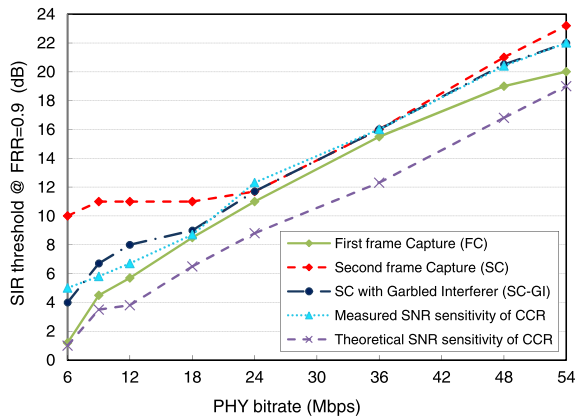


Fig. 7. SIR thresholds at FRR = 0.9 for 802.11a bitrates. Clear Channel Reception (CCR) shows the receiver sensitivity levels (@FRR = 0.9) in terms of SNR without any interference.

provides the SIR threshold for successful capture of the sender's frame (FOI) as the bitrate increases in the three capture cases. The FRR of each plot in Fig. 7 corresponds to one experiment which lasts 30–50 s. The number of sender's attempted transmissions in one session varies depending on the different bitrates but is at least 12,000. For comparison purposes, we also plot the measured SNR thresholds in CCR and the theoretical minimum SNR thresholds in CCR. The latter SNR thresholds are calculated by assuming perfect synchronization and channel estimation in [20]. Our measured SNR thresholds are approximately 3–4 dB higher than the theoretical minimum SNR thresholds required by the Atheros chipset.

Before elaborating on the effect of various PHY bitrates, we summarize our key findings. First, the two capture stages (preamble detection and MAC CRC check) have their own different SIR thresholds for successful stage fulfillment. Second, the *capture threshold* for SC is not affected by the FOI's encoding rate and the threshold is around 11 dB. Third, the SIR threshold for the MAC FCS check increases as the bitrate increases. Fourth, the interferer's frame encoding rate does not reveal any observable relation to the capture SIR threshold.

The SIR for FC (Fig. 7) increases from 1 dB to 20 dB as the bitrate of the FOI increases. Remember that in the FC case, only the MAC CRC check stage matters because the FOI arrives first without interference. The interferer's frame arrives when the receiver is already locked onto the FOI and hence is dealt with as noise by the receiver. As it hinders the decoding of sender's MAC frame, the higher SIR is required to decode the FOI transmitted at the higher bitrate. In summary, we can interpret the SIR thresholds of the FC case as the SIR thresholds for the successful MAC CRC check. Again, FC requires lower SIR than the other capture cases over all the 802.11a bitrates and the FC curve is close to the theoretical minimum SNR thresholds at the bitrate of 6, 9 and 54 Mbps. The fact that the receiver has already locked onto the FOI when the interference frame starts seems to improve the capture performance.

In Fig. 7, the curves of FC, SC-GI, and SNR of CCR show a similar trend while SC's SIR threshold sustains around

11 dB for 6–18 Mbps. It starts increasing when the bitrate exceeds 24 Mbps, and joins the trend of other curves. The 11 dB SIR threshold for SC seems to be an artificially set parameter by chipset designers to protect the previous frame before deciding to capture the next frame that arrives later, which is also called the *capture threshold*. The capture threshold only determines whether to switch from the first frame to the second frame's preamble; the threshold is not related to the encoding and modulation rate for the MAC frame. From the curves of FC and SC-GI, the required SIR thresholds for a successful MAC FCS check of 6–18 Mbps frame are less than or equal to 11 dB; if the preamble is successfully decoded, which means the SIR is higher than 11 dB, the FCS check stage is successfully passed for 6–18 Mbps. Overall, when the bitrate is lower than 24 Mbps, the SIR threshold for a successful frame capture is determined by the preamble detection stage. In contrast, when the bitrate is 24 Mbps or higher, the SIR threshold for a successful frame capture is determined by the MAC CRC check stage.

The 11 dB capture threshold also explains the irregular FRR vs. SIR pattern in the SC case in Fig. 5b when $\Delta t < 3.2 \mu\text{s}$. When the FOI arrives at the very early stage of the first frame's preamble detection, the receiver can capture the newly arriving FOI even at SIRs lower than 11 dB. The results for 9, 12, and 18 Mbps also show a similar pattern but the lowest SIR that exhibits FRR > 0.2 (0 dB SIR for 6 Mbps in Fig. 5b) increases as the bitrate increases. From 24 Mbps, there is no such irregular pattern in the region $\Delta t < 3.2 \mu\text{s}$. Since SIRs higher than 11 dB are required to decode 24 Mbps or higher transmissions without error, switching to the second frame with SIRs below 11 dB will fail eventually due to the MAC CRC check.

The curve of SC-GI in Fig. 7 is very close to that of the measured CCR sensitivity over all bitrates and this verifies that the interferer's signal is deemed as noise power by the receiver in the case of SC-GI.

4.3. Unicast vs. broadcast

In all of the previous experiments, both the sender and the interferer transmit broadcast traffic for the sake of convenience (e.g., no retransmission, no ACK). For comparison purposes, we also tested the case when the interferer transmits a unicast frame to a node other than the receiver (e.g., Sniffer 2). This test aims to check whether the receiver still locks on and keeps receiving the interferer's frame even when it is not destined to the receiver, i.e., the SC or the SC-GI happens with unicast frames destined to another receiver. Our unicast experiments with 802.11a 6 Mbps PHY rate for both the sender and the interferer resulted in the same SIR threshold as the broadcast test in Fig. 4.

4.4. Implication on interference modeling

One of the major implications of our SIR threshold results is that when we construct a SIR-based network conflict graph we should use different SIR thresholds for different frame arrival timings (FC and SC) even when the other parameters (i.e., 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 2 dB to 11 dB as the timing relation varies from FC to SC. To our best knowledge, the frame arrival timing has not been considered in any of conflict graph-based interference models.

5. Measurement-based modeling of carrier sense

As described in Section 2.2, the 802.11a PHY determines the busy medium by the PLCP preamble (and header) detection (PD) or the energy detection (ED) and the 802.11a standard defines the ED threshold (EDThres) to be 20 dB higher than the RX sensitivity (RXSens). In order to observe the PD and ED behavior from our Atheros-based testbed, we measure the medium access probability (MAP) of two concurrently transmitting broadcast senders, S1 and S2, where S2 does not sense S1's transmission: thus, S2 is able to access/use the medium whenever it wants to transmit. We control S2's transmission power so that we can observe S1's varying MAP as a function of S1's received SNR of S2's transmission. We use SNR because Atheros chipsets implement the ED carrier sense by comparing the signal power against the noise level [21]. S1 transmits saturated traffic and it always has a packet to send in its TX queue.

We first consider the case when S2 transmits unsaturated traffic. We control S2's transmission timing to create the setting where S1 misses the preamble of S2's frame in most cases so that S1 depends more on the energy detection than preamble detection. The solid line in Fig. 8 plots the MAP of S1 as a function of S1's received SNR of S2's transmission. It clearly shows the distinct operations of PD and ED. As SNR increases from 2 dB, S1's MAP decreases from 1 because S1 begins to detect the preamble of S2's frame. To compare, we also plot the FRR of S2's frame at S1 that was measured separately. The FRR rises when SNR is at 2 dB and it verifies that the PD happens. The FRR curve shows that S1's PHY is able to detect S2's preamble almost perfectly at 7 dB SNR but S1 does not sense many of S2's frames as the solid line (S1's MAP) still shows high value of 0.9. The S1's MAP drops down to 0.6 at around SNR 11 ~ 12 dB and stabilizes when SNR reaches 13 dB and we believe the energy detection plays its role in this region.

In Fig. 8, we also show the S1's MAP when S2 transmits saturated traffic. When the medium is saturated, S1 has

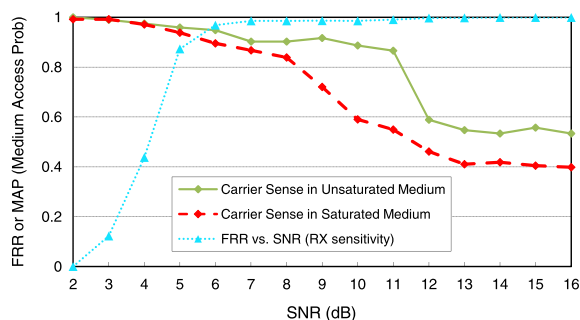


Fig. 8. Carrier sense performance as a function of SNR. FRR vs. SNR plot is put together to better understand preamble detection behavior.

more chances to sense S2's preamble and the MAP curve exhibits a smoother transition from 1 to 0.4 than the case with S2's unsaturated traffic. However, the MAP also stabilizes when SNR is 13 dB. Thus, we conclude EDThres is 13 dB for the Atheros chipset. If we choose 5 dB for RXSens (@FRR = 0.9), EDThres is 8 dB higher than RXSens. Compared to the 2 dB where PD starts to work, EDThres is 11 dB higher. In Section 7, we will set EDThres to be 10 dB higher than RXSens in conjunction with the report from [9] and our earlier work [22].

6. Simulator modifications

NS-2 [6] and QualNet [7] are two of the most popular network simulators but their 802.11 PHY and MAC implementations do not precisely model off-the-shelf 802.11 card behavior. We first describe the reception models of NS-2 and QualNet. As the QualNet model is more precise than NS-2, we choose to improve the QualNet model based on our observations in Section 4.

6.1. Current simulator models

The flow chart of the original NS-2 and QualNet RX models is illustrated in Fig. 9a. In the flow chart, a frame that newly arrives at the receiver is referred to as a NEW frame. If the NEW frame's preamble is successfully detected and the receiver locks onto the frame, the frame is called a RCV frame. For the ease of presentation, the flow chart begins at the moment when a new frame arrives when the PHY state is *idle* or *sensing*⁷ instead of the *receiving* state. The other case when a new frame arrives in the middle of the PHY *receiving* state will be explained later in the flow chart.

When a NEW frame arrives, its Received Signal Strength (RSS), denoted by $rss(NEW)$, is compared with the RX sensitivity (RXSens). If $rss(NEW)$ is less than the RXSens, the NEW frame is discarded. In QualNet, $rss(NEW)$ is added to the interference power (int_power) until the end of the NEW frame transmission. However, NS-2 does not add up the interference power.⁸ If $rss(NEW)$ is larger than the RXSens, the receiver starts its reception of the NEW frame. The NEW frame becomes the RCV frame and the receiver PHY state becomes *receiving*. When $rss(NEW)$ is smaller than the RXSens, QualNet compares the sum of all currently transmitting frames' signal powers and noise level with the RXSens for carrier sense.

When another NEW frame arrives during the RCV frame reception, NS-2 immediately makes the capture decision of the RCV frame: if the RCV frame is stronger than the NEW frame by CPTThres (capture threshold, 10 dB in NS-2), the RCV frame is captured and the NEW (or the old) frame is ignored. Otherwise, both the RCV and the NEW frames are discarded. QualNet uses a more enhanced model: it treats the NEW frame's signal power as the interference power for the RCV frame

⁷ When the PHY is *sensing*, it can receive but can not transmit.

⁸ The signal power of a discarded frame can interfere with another frame but NS-2 ignores the discarded frame.

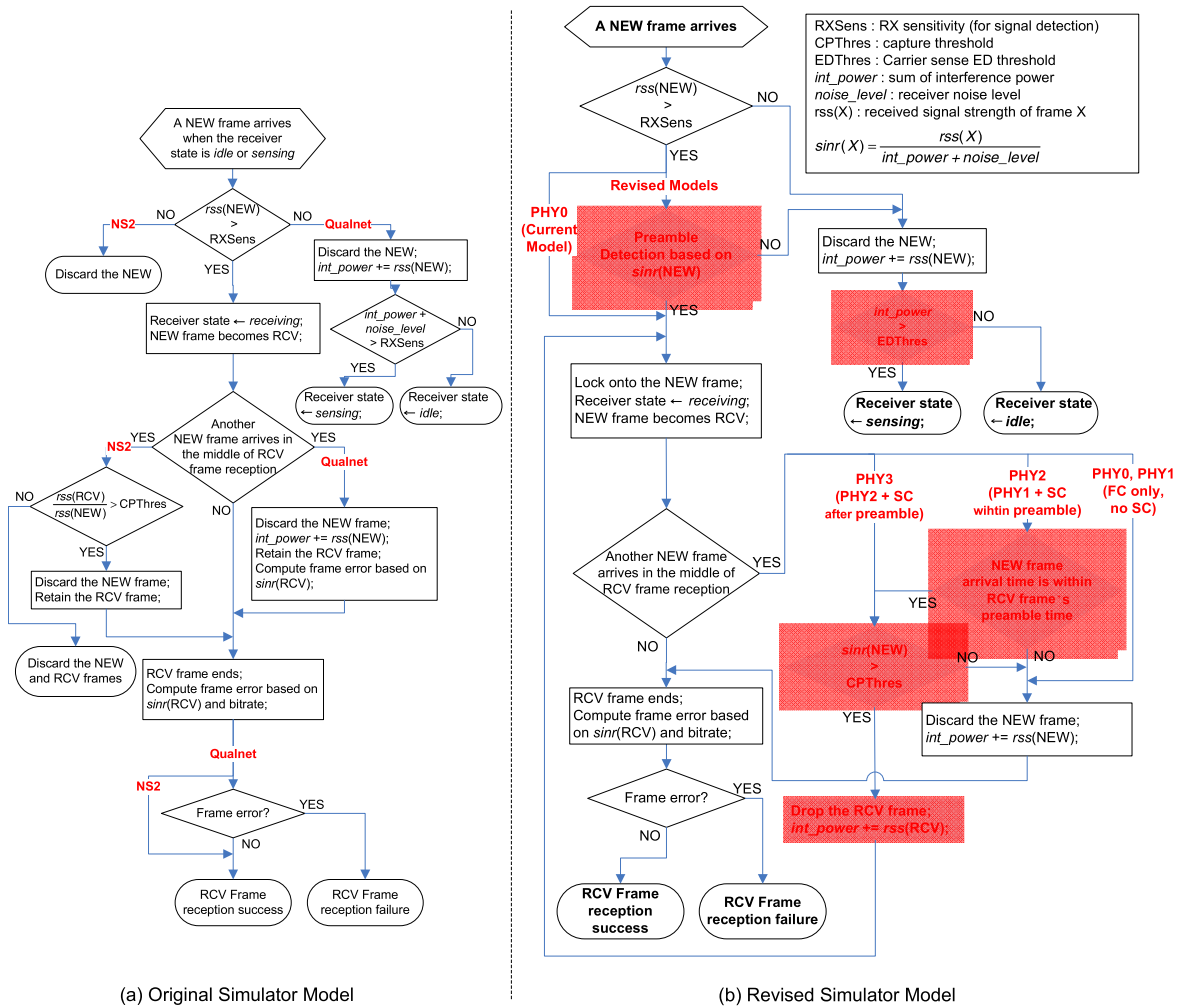


Fig. 9. Simulator model flow charts.

($int_power += rss(NEW)$), computes the bit error rate of the RCV frame based on the RCV frame's bitrate and SINR, and appropriately decides the frame error. If there is no error at the end of the RCV frame reception, the frame is delivered to the MAC layer.

To summarize, the current NS-2 and QualNet implements FC (First frame Capture) but not SC (Second frame Capture). And they do not use a separate ED threshold, which is higher than the RXSens.

6.2. Modified simulator model

Because the QualNet simulator models the RX process in presence of interference better than NS-2, our RX model featuring the capture effect is described by explaining how to modify QualNet (version 3.9.5). We enhance QualNet by augmenting two components in the RX model: the SINR-based preamble detection and the capture algorithm. To effectively show the impact of each component in the enhanced simulator model, we define four PHY models as follows.

- **PHY0:** The QualNet version 3.9.5 model. Only FC (first frame capture) is possible.
- **PHY1:** PHY0 + SINR based preamble detection.
- **PHY2:** PHY1 + SC (Second frame Capture) *within* first (RCV) frame's preamble time.
- **PHY3:** PHY2 + SC *after* first (RCV) frame's preamble time (MIM mode is supported).

The flow chart for the revised simulator model is presented in Fig. 9b. Our revision is highlighted by the dotted boxes.

6.2.1. SINR-based preamble detection

As discussed in Section 2.1, the frame RX process requires not only energy detection but also PLCP preamble detection and PLCP header reception, which are vulnerable to interference. However, QualNet does not consider interference when deciding whether to receive a newly arriving frame. Thus, we enhance the RX model to check $sinr(NEW)$ as well as $rss(NEW)$ before locking on the newly arriving frame and moving into the *receiving* state.

In order to determine the SINR threshold for the preamble detection, we refer to the FRR curves of 6 Mbps bitrate in SC-GI as shown in Fig. 4. Fig. 4 shows that the SIR threshold (@FRR = 0.9) in SC-GI is 4 dB while the SIR threshold of FC is about 1 dB. Because the preamble, the PLCP header, and the MAC frame body are all encoded at 6 Mbps bitrate, the only difference between SC-GI and FC is whether the interference signal is present (SC-GI) or not (FC) when the receiver detects the NEW frame's preamble. Thus, we should consider the SC-GI case to determine the SINR threshold for the preamble detection. According to the FRR curve of SC-GI in Fig. 4, we determine the success of preamble detection by a probability function linearly increasing from zero to one as the SINR increases from 1 dB to 5 dB.⁹

If $rss(NEW)$ is greater than the $RXSens$ and the preamble detection is successful, the receiver goes into the *receiving* state and the NEW frame becomes the RCV frame. Note that the SINR-based preamble detection logic is not applied to the PHY0 model (i.e., the current QualNet model).

6.2.2. Capture models

If a NEW frame arrives during the reception of the RCV frame, we apply different capture algorithms depending on the PHY models. The PHY0 and PHY1 models follow the QualNet implementation: the NEW frame is discarded and its signal is added to the interference power for the RCV frame. In the PHY2 model, if a NEW frame arrives after the RCV frame's preamble time, the NEW frame is discarded. If a NEW frame arrives within the RCV frame's preamble time and $sinr(NEW)$ is greater than the *capture threshold* (CPT_{thres}), the NEW frame is captured (the RCV frame is dropped), and $rss(RCV)$ is added to the interference power. When we calculate $sinr(NEW)$, $rss(RCV)$ is also considered as the interference power for the NEW frame. The PHY3 model compares $sinr(NEW)$ to the CPT_{thres} regardless of the arrival timing between the two frames. In other words, the PHY3 model supports the Message-In-Message (MIM) mode; even if the NEW frame arrives after the RCV frame's preamble time, the NEW frame can be captured.

Because the QualNet does not have the CPT_{thres} (since it does not support SC), we use the SIR threshold result of SC from Section 4, which reports that at least 11 dB SINR is required to capture a NEW frame that arrives during the reception of a previous frame. In line with the settings in [22], we use 10 dB CPT_{thres} instead of 11 dB as a simulation parameter since this value also slightly varies over different chipsets. Because the capture decision is made at the time of preamble detection, the 10 dB CPT_{thres} is independent of the MAC data bitrate. We conclude that Prism chipsets follow the PHY2 model and Atheros chipsets follow the PHY3 model according to [2] and our observations in the previous section.

6.2.3. Carrier sense models

QualNet compares the sum of all currently transmitting frames' signal powers plus the noise level to the $RXSens$ for

carrier sensing. First, we revised QualNet to use a separate parameter (EDThres), which is configurable. Second, we revised QualNet not to compare the sum of the frame signal power of the noise power with the threshold because (i) the 802.11a standard defines the thresholds ($RXSens$ and EDThres) in dBm that is independent of the noise level and (ii) Atheros chipsets implements EDThres in terms of SNR. In either case of the standard and the Atheros implementation, the noise level should not be added to the signal power when we determine busy medium.

If the EDThres is the same as the $RXSens$ (for example, in the QualNet model or in 802.11b), the preamble detection does not play an important role in the carrier sense mechanism because the energy detection module will eventually sense the channel busy even when the preamble portion was missed. However, when the EDThres is (much) higher than the $RXSens$, the success or failure of the preamble detection greatly affects the carrier sensing. As discussed in Section 6.2.1, we revised QualNet to have the SINR-preamble detection logic as it is an important part of 802.11a carrier sense.

6.3. Additional qualnet modifications

6.3.1. Desynchronization

According to the IEEE 802.11 standard, when a node detects the medium is busy, it freezes its backoff timer and suspends its backoff procedure. However, as earlier studies (e.g., [23]) pointed out, in QualNet, the backoff procedure is suspended after the backoff timer is decreased by propagation delay. Hence, we fix the problem to prevent nodes from having unrealistically low collision rate due to the desynchronization [23].

6.3.2. Timing jitter

Because frame transmissions in QualNet are slotted and scheduled based on one global clock, all the frames in a collision from the mutually carrier sensing senders start transmissions exactly at the same time. Thus, the arrival time difference between two frames in a collision is determined solely by the propagation delay. In other words, when two frames collide because the backoff counters reach zero at the same time, the frame from the sender closer to the receiver always arrives at the receiver prior to the frame from the farther sender. However, in real wireless communications, senders are not perfectly slot-synchronized. Clock drift, missed beacon and hardware jitter contribute to timing jitter. We observed that transmission time difference can be up to 4 μ s but in most cases fell into $[-2 \mu$ s, 2 μ s] window from our measurements. Thus, we randomly select jitter from $[-2 \mu$ s, 2 μ s] and apply it before the start of each frame's transmission.

7. Performance evaluation

To evaluate the four PHY models (PHY0–PHY3) and observe the effect of different ED thresholds using QualNet, we consider two radio environments: indoor and outdoor radio propagation models. For the indoor propagation model, we use the log-distance model [24,25]. Here, we

⁹ We observe that different chipset models and driver settings (e.g., ANI on/of) exhibit slightly different shape of curves of SC-GI. In our earlier work [22], we used the 4 dB–10 dB model.

Table 1
Parameters used in simulation.

| Parameter | Value |
|-----------------------------|--|
| Path loss model | Two-ray ($n = 2$) Log-distance ($n = 4$) |
| Shadowing model | Constant |
| Shadowing mean | 4 dB |
| TX power | 16 dBm |
| RX sensitivity (RXSens) | -88 dBm |
| ED threshold (EDThres) | RXSens + (0, 10, 20 dB) |
| Capture threshold (CPThres) | 10 dB |
| Preamble Detection | Linearly increasing prob. function: [1 dB, 5 dB] \rightarrow [0, 1] |
| Rate adaptation mechanism | Auto Rate Fallback (ARF) |
| RTS/CTS | Disabled |
| TCP payload size | 1400 bytes |

set the path-loss exponent (n) to 4. For outdoor propagation model, we use the two-ray path-loss model with $n = 2$ [24]. Multi-path fading is not considered for either model. With TX power and RX sensitivity in Table 1, the maximum transmission ranges at 6 Mbps in outdoor propagation and indoor propagation models are 328.6 m and 48.5 m, respectively.

Each node uses the 802.11 DCF without RTS/CTS in ad hoc mode. Other simulation parameters are listed in Table 1. We evaluate the performance of the four PHY models as the number of flows increase. To remove the effect of routing and multi-path relaying, we intentionally arrange the sender and the receiver of each flow to be only one hop away.

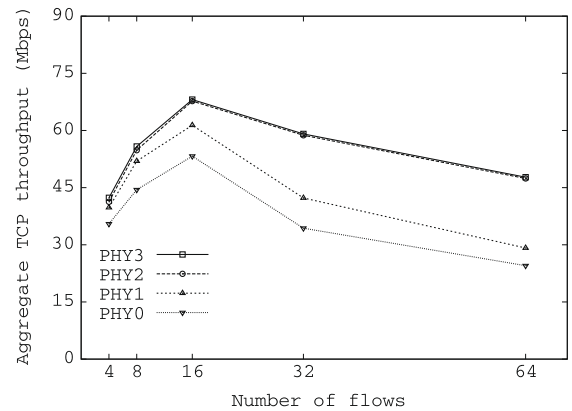
We first evaluate the performance of the revised models under a TCP traffic model and the results with UDP traffic will be presented later in this section.

To generate TCP traffic, each sender transmits a large file using an FTP application, which leverages TCP. The TCP payload size is set to 1400 bytes. We use Auto Rate Fallback (ARF) for the rate adaptation mechanism. We set the terrain size to $1000 \times 1000 \text{ m}^2$ for the outdoor propagation model and $149 \times 149 \text{ m}^2$ for the indoor propagation model.¹⁰ We equally divide the simulation terrain into a number of cells, where the number is equal to the number of sender–receiver flows, and randomly place each sender within each cell. We randomly choose the distance between a sender and its receiver in a range from 10 m to 57 m in outdoors with a uniform distribution. Likewise, the distance between the sender and the receiver ranges from 3.5 m to 20 m uniformly for indoor propagation models.¹¹ To plot the performance metrics for each number of flows, we average results from 30 different network topologies (i.e., 30 different sender–receiver placements) created with different random number seeds.

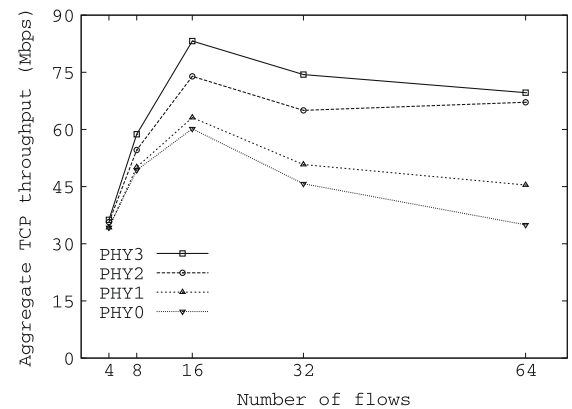
In the following subsection, we discuss the effect of SINR-based preamble detection and different capture behavior by observing the performance of the four PHY models while EDThres is set to be the same as RXSens

¹⁰ To simulate the equivalent network area in terms of hop count, we scale down the indoor area by considering the ratio of indoor and outdoor TX ranges. $149 \text{ m} \approx 48.538 \text{ m} \times 1000 \text{ m} \div 328.602 \text{ m}$.

¹¹ Both 57 m and 20 m are the largest distance that packets are received at 54 Mbps in absence of interference in outdoor and indoor propagation models, respectively.



(a) Outdoor propagation ($n=2$)



(b) Indoor propagation ($n=4$)

Fig. 10. Average aggregate TCP throughput.

according to the default QualNet setting. The effect of different EDThres settings on the performance will be discussed later. Lastly, we will show the performance result with UDP traffic.

7.1. Performance of preamble detection and capture models

In Fig. 10a, we show the aggregate TCP throughput of the four PHY models as the number of sender–receiver pairs varies in the outdoor propagation model. The aggregate throughput of the real chipset models (PHY2 and PHY3) are greater than that of the QualNet model (PHY0). The advantage of the capture logic implementation over non-capture models is substantial; however, the difference between PHY2 and PHY3 is almost negligible. As shown in Fig. 11, the number of SC (Second frame Capture) occurrence difference between PHY2 and PHY3 is small in the outdoor propagation model. Let us consider a topology where the distance between a sender and a receiver (d_s) is shorter than the distance between an interferer and a receiver (d_i): $d_s < d_i$. The SIR at the receiver is expressed as $(d_i/d_s)^n$ where n is a path loss exponent. As n increases, a signal attenuates more rapidly over distance and the SIR gets larger. Thus, the SIR in the indoor propagation model

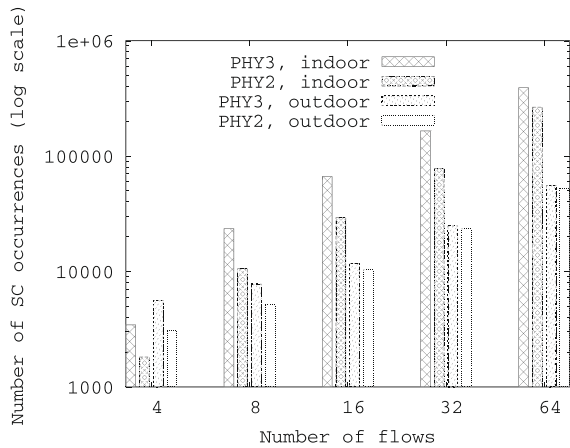


Fig. 11. Number of SC occurrences of PHY3 and PHY2.

is more likely to exceed the CPT_{thres} and the sender's frame can be successfully captured at the receiver with a higher probability. Fig. 11 verifies the large difference of the SC occurrences between PHY2 and PHY3 in the indoor model. Note that the y-axis of Fig. 11 is in log-scale. We plot the aggregate TCP throughput in the indoor propagation model in Fig. 10b, where PHY3 achieves a noticeable throughput gain over PHY2.

Simply adding the SINR-based preamble detection logic (PHY1) yields a notable gain over PHY0. The preamble detection logic prevents the PHY layer from locking on and going into *receiving* state upon a useless frame and allows the PHY layer to transmit its own frame. The performance gain of PHY2 over PHY1 shows the impact of FC (first frame capture). Another interesting observation is that the aggregate throughput of every PHY model reaches the highest value at the point when the number of flows is 16 and decreases beyond that point. The throughput reduction is more notable with PHY0 and PHY1 because they suffer from collisions more severely due to the lack of capture capability in contrast to PHY2 and PHY3. The other cause of the reduction is the carrier sense mechanisms: the reduction disappears as we use higher carrier sense thresholds (ED_{thres}). More details will be given in Section 7.2.

Fig. 12 shows the Jain's fairness index [26] among the flows in the indoor propagation model. The fairness index of 1 means that all the flows obtain the same throughput. As the PHY model evolves from PHY0 to PHY3, the fairness improves. The SINR-based preamble detection logic and the capture capabilities not only increase the throughput performance but also improve the fairness between flows. The fairness in the outdoor propagation model shows a similar pattern.

To evaluate the impact of the different PHY models on the MAC layer performance, we measure *MAC efficiency*, defined as the ratio of the number of successful MAC frame transmissions to the total number of MAC frame transmissions including the retransmissions. If there are no retransmissions (or no transmission failures), MAC efficiency is 1. In Fig. 13, we plot the MAC efficiency as well as the aggregate TCP throughput to analyze the correlation between

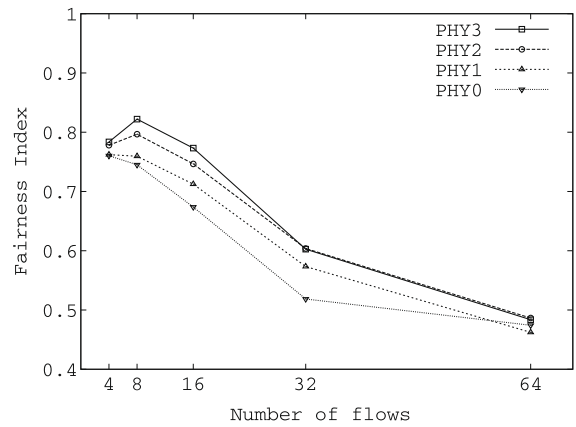


Fig. 12. TCP flow fairness (indoor).

two metrics. We measure the metrics for 30 topologies of the two radio propagation models, each of which consists of 16 flows. As before, we create 30 different topologies (sender–receiver placements) with different random number seeds. The seed number is used as a topology index number in Fig. 13a and b that plot both aggregate TCP throughput and MAC efficiency metrics in outdoor and indoor propagation models, respectively. Since we use ARF for the rate adaptation mechanism, MAC efficiency of each PHY model is relatively high, around 0.9. Overall, there are marginal differences in both metrics among the PHY models in outdoor environments. On the other hand, in indoor environments, somewhat higher variation in MAC efficiency among the PHY models results in large TCP throughput difference. Recall that all flows span only one hop to remove the effect of routing. If each flow spans mul-

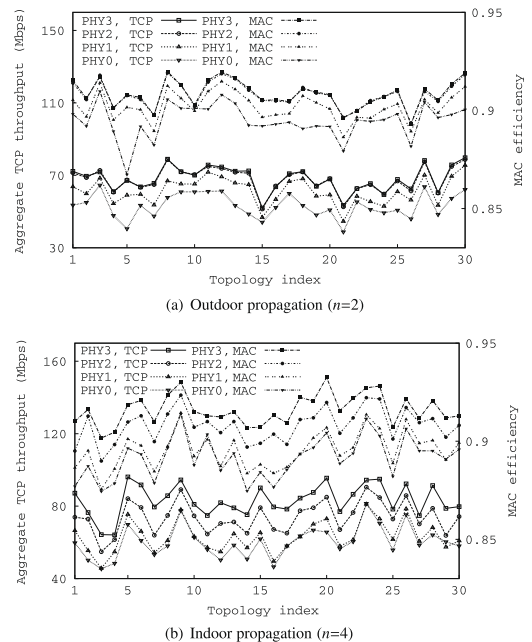


Fig. 13. TCP throughput and MAC efficiency (16 flows). The 30 topologies are created with 30 different random number seeds.

tiple hops, MAC layer retransmissions may generate TCP timeouts, which will invoke even larger TCP throughput fluctuation. Across all the settings, PHY3 achieves the highest MAC efficiency and the largest aggregate TCP throughput. This result demonstrates that the capture effect enhances the MAC frame delivery ratio and the ARF algorithm can leverage the capture effect.

7.2. Effect of carrier sense ED thresholds

To evaluate the impact of the revised carrier sensing model on the network performance, we run simulations with three difference EDThres settings: (i) the EDThres is equal to the RXSens (default QualNet setting), (ii) the EDThres is 20 dB higher than the RXSens (802.11a standard setting), and (iii) the EDThres is 10 dB higher than the RXSens (Atheros chipset setting). The Atheros chipset setting is based on the observation in Section 5.

In Fig. 14, we show the aggregate TCP throughput in the indoor propagation model for each EDThres setting. The throughput improvement of PHY1–PHY3 over PHY0 increases as the number of flows increases and the EDThres increases. Not only does the standard setting result in a large performance improvement up to 662% (Fig. 14c), but the Atheros setting (with the 10 dB increase) also shows ample improvement up to 421% (Fig. 14b).

In particular, PHY1's throughput gain over PHY0 in the Atheros setting (up to 276%) and in the standard setting (up to 459%) are much greater than that of the QualNet setting (up to 13%). Recall that the PHY1 model implements the SINR-based preamble detection logic in addition to the RX process of PHY0. Suppose an 18 Mbps frame (denoted by NEW) arrives at the receiver when the receiver is in an *idle* state and $\text{RXSens} < \text{rss}(\text{NEW}) < \text{EDThres} = \text{RXSens} + 20 \text{ dB}$ and $\text{sinr}(\text{NEW}) < 5 \text{ dB}$. As shown in the flow chart of Fig. 9a, in the original model, the receiver locks on the frame and goes into the *receiving* state because $\text{rss}(\text{NEW})$ is larger than RXSens ; however, this frame will be corrupted because at least 10 dB SINR is required to decode the 18 Mbps frame without error (Fig. 7). If another frame (denoted by NEW') arrives at the receiver with a high SINR ($> 10 \text{ dB}$) during the reception of NEW, the receiver cannot receive NEW' because it has locked onto NEW. Moreover, if there is a frame to send from the receiver node, the node wastes its chance to send its own frame

due to NEW. However, in the PHY1–PHY3 models of Fig. 9b, the receiver may fail both the preamble detection (due to the low $\text{sinr}(\text{NEW})$) and the energy detection (because $\text{rss}(\text{NEW}) < \text{EDThres}$) and stay in the *idle* state; thus, it has a better chance to receive NEW' or send its own frame. Hence, (1) the default simulator model (PHY0)'s lack of SINR-based preamble detection logic causes inefficient PHY operations and (2) the combination of SINR-based preamble detection and high EDThres prevents the PHY from being occupied by a useless frame (NEW) and enables the PHY to receive a stronger useful frame (NEW') or to transmit its own frame.

The performance gain of PHY3 over PHY2 and of PHY2 over PHY1 are also increased in the Atheros and standard settings. In those settings, senders transmit more aggressively due to the higher EDThres and the chance of SC also increases.

Based on the simulation results, we believe that the 802.11a carrier sense mechanism can benefit from better channel utilization than the 802.11b/g carrier sense mechanism that sets the EDThres more conservatively (equal to or lower than the RXSens). Since the EDThres is 20 dB higher (100 times stronger) in the 802.11a standard, 802.11a stations are allowed to transmit more aggressively than 802.11b stations. If the PHY capture was not supported, the increased transmission attempts would only aggravate interference and decrease MAC efficiency. However, from our simulation results, although not included in this paper, we observed that the MAC efficiency of the increased EDThres settings does not decrease compared with that of the QualNet setting ($\text{EDThres} = \text{RXSens}$); the increased transmission attempts with the unchanged MAC efficiency means the increase of the successfully received frames. Hence, thanks to the capture effect (both FC and SC), the increased transmission attempts in 802.11a networks result in better spatial reuse especially when there are a large number of flows.

7.3. Experiments with UDP traffic

As TCP performs poorly over wireless links due to the congestion control mechanism, we also conduct simulation experiments with UDP traffic. The same topologies and simulation settings are used as described at the beginning of this section except senders generate UDP flows using




Fig. 14. Aggregate TCP throughput of the PHY models in the indoor propagation model with different EDThres settings.




Fig. 15. Aggregate UDP throughput of the PHY models in the indoor propagation model with different EDThres settings.

CBR applications. The CBR rate of each flow is set high enough to make each sender's transmission buffer always filled with packets to send.

Fig. 15 shows the aggregate UDP throughput in the indoor propagation model for the three EDThres settings. The overall tendency remains similar to TCP in comparing Fig. 15 with Fig. 14. The aggregate UDP throughput is larger than the TCP throughput at all points; due to the lack of congestion control, UDP senders transmit more aggressively than TCP senders. However, the relative throughput gain of PHY1–3 over PHY0 is similar to those in TCP. For example, with 64 flows, the throughput gain of PHY3 over PHY0 is 643% with UDP flows in Fig. 15c and is 662% with TCP flows in Fig. 14c. The fairness index and MAC efficiency with UDP flows are also similar to TCP and hence are omitted. In general, the impact of the revised PHY models is substantial for both TCP and UDP traffic models.

8. Related work

8.1. Capture effect

There have been studies on the capture effect mostly in random access networks (e.g. [27]). One of the early 802.11 capture studies is in [11], which identifies overlapping transmission time and/or signal power difference as the factors that decide the capture. Its SIR-based capture model is analyzed in [28]. The MIM mode and the triggering condition for the retraining process (i.e. synchronizing with the next incoming frame) are proposed in [10]. According to our experiments, we believe this mechanism is embodied in the Atheros chipsets.

Recent experimental studies on the capture effect are carried out in 802.11 networks [2,14,3] and sensor networks [29,30]. The first experimental study that discusses the Second frame Capture (SC) in 802.11b networks is in [2], which shows that the SC within preamble can happen with Prism chipsets. In [14], after providing empirical studies on how the capture effect causes unfairness, the authors propose to remedy the unfairness problem by adjusting the MAC parameters such as the retransmission limit, minimum contention window size and transmission power. However, as the carrier sense range is large in the experiments of [2,14], all the nodes on their testbed can sense the other nodes' transmissions. The capture effect is inves-

tigated only for the cases when the backoff counters reach zero at the same time. Although the nodes are supposed to start the transmissions exactly at the same moment, there is up to 20 μ s 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 [2]. Hence, the timing relation between the sender and the interferer is not thoroughly investigated in those studies. Another interesting study is reported in [3], where a wireless emulator is devised to disable carrier sense to insert and observe controllable propagation loss. The author of [3] quantifies the effects of timing and signal power on the performance of the First frame Capture (FC) using their "wired" 802.11b testbed with Prism wireless cards. However, the Second frame Capture (SC) after preamble and SC-GI capture cases are not discussed in the literature on 802.11 capture [2,14,3].

Capture effects in sensor network testbeds that consist of Mica2 Motes with Chipcon CC1000 radios [31] are studied. The capture behavior under various network settings are tested in [30]. It is observed that the SIR threshold to trigger the capture may change over a 6 dB range depending on the transmission powers. Capture experiments with multiple interferers are also conducted. The frames are however transmitted simultaneously and hence the timing relation is not investigated. Similar to our work, [29] reports that 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 are discussed in [32], which also studies the relation of carrier sense and interference, and their impact on the performance of two contending links.

8.1.1. 802.11 Capture simulation models

There have been several simulation studies on the capture effect. However, their simulation models do not precisely model the capture process in real 802.11 chipsets. To our best knowledge, the work of [2] is the first attempt to revise the capture simulation model based on the observations from off-the-shelf 802.11 hardware. From the Prism chipset-embedded 802.11b testbed measurement, the authors modified NS-2 to account for the SC within the first frame's preamble time. Chang et al. [23]

analyze the aggregated network throughput when there are concurrent multiple flows. In their simulation, the authors modify QualNet to reflect the SC within the first frame's preamble time [2]. Both [2] and [23] considered only the single hop carrier sense range; capture with hidden interferers was not considered. In this paper, we consider capture with hidden interference and show the performance of the capture model that supports the SC after preamble (MIM mode). Along with the SINR-based preamble detection logic, our simulator modifications reflects the actual systems more precisely.

Two recent studies [33,34] revise 802.11 MAC, PHY, and channel models for NS-2 and augment incomplete interference and capture models. The revised model in [33] supports only the SC within preamble time and considers the RSS in evaluating successful preamble reception while [34] presents more accurate capture models that include SINR-based frame reception and MAC frame capture (i.e., the SC after preamble). As the main focus in [33,34] is not on the capture model revision, they do not evaluate the impact of revised PHY models (of preamble detection and capture) on the network performance.

Both Ware et al. [11,35] and Ganu et al. [14] illustrate the throughput unfairness problem caused by the capture effect in the network topology where all senders carrier sense each other. In our simulation, we compare the fairness of different PHY models when multiple senders are hidden from one another. In addition, [11,14] consider only the case of two senders and one common receiver; it is hard to generalize the throughput unfairness problem for multi-hop wireless networks.

8.2. Carrier sense

There are a number of studies that evaluate the effect of the carrier sense (energy detection) threshold on the network throughput performance [36–38]. They suggest to tune the carrier sense threshold to maximize the network performance considering various factors such as transmission power, bitrate, MAC overhead, interference, etc. The recent work in [39] proposes an algorithm to detect selfish 802.11 devices that increase their carrier sense thresholds to obtain a higher, unfair share of available medium resource. Similar to our observations, [36,38] show that the use of a small carrier sense range (i.e., high carrier sense ED threshold) can allow stations to attempt more concurrent transmissions. However, none of the previous research efforts consider the preamble detection as an important carrier sense mechanism; instead, they focus only on the tuning of the carrier sense threshold.

Measurement reports on the carrier sense ED threshold are also available in [40,9]. The Intel wireless card and its proprietary firmware used in [40] set the default carrier sense ED threshold to be equal to the receiver sensitivity level for both 802.11g and 802.11a while the authors in [9] show about 10 dB higher carrier sense threshold above the RX sensitivity from their Atheros based 802.11a testbed, which is in line with our observations on the Atheros cards. Neither of them separately evaluated the effects of preamble detection and energy detection on the carrier sense performance.

9. Conclusion

Using our 802.11a testbed, we performed experimental studies on the physical layer capture effect, preamble detection, and carrier sense. We presented the precise terms and conditions (timing, SIR, and PHY bitrate) in which the receiver can successfully decode a frame in the presence of interference. We observed that the successful capture of a frame is determined in two stages, preamble detection and MAC frame CRC check, and we showed the impact of an accurate preamble detection model on capture and carrier sense performance. Based on the experimental results and analysis, we tried bridging the gap between the IEEE 802.11a PHY simulation models and real wireless network systems. We made modifications to the QualNet simulator (version 3.9.5) in the following ways: (i) we implemented the SINR condition in addition to the RSS condition for the preamble reception that determines the start of a frame reception and the frame capture, (ii) we provided vendor specific PHY models (Prism and Atheros) as different chipsets have different implementations, and (iii) we corrected the carrier sense (energy detection) threshold value. In order to evaluate the impact of these changes on the wireless network performance, we conducted comprehensive simulation experiments in various scenarios. Our results show that our modified model increases the aggregated TCP/UDP throughput up to more than six times that of the QualNet model. The analysis of our experiments and the revised simulation models will help us better understand interference/capture and its impact on throughput, which is crucial for the design of network protocols.

Acknowledgement

The authors would like to express appreciation to Mr. Wonho Kim and Dr. Hoon Chang for their help in building a testbed and solving a QualNet timing problem, respectively.

References

- [1] J. Lee, W. Kim, S.J. Lee, D. Jo, J. Ryu, T. Kwon, Y. Choi, An experimental study on the capture effect in 802.11a networks, in: Proceedings of the ACM WINTech, Montreal, Canada, 2007.
- [2] A. Kochut, A. Vasani, A. Shankar, A. Agrawala, Sniffing out the correct physical layer capture model in 802.11b, in: Proceedings of the IEEE ICNP, Berlin, Germany, 2004.
- [3] G. Judd, Using Physical Layer Emulation to Understand and Improve Wireless Networks, Tech. Rep. CMU-CS-06-164, Ph.D Thesis, School of Computer Science, CMU, October 2006.
- [4] PRISM Wireless LAN, Intersil Corporation, <<http://www.intersil.com/globespanvirata/>>.
- [5] Atheros Communications, <<http://www.atheros.com/>>.
- [6] The Network Simulator 2 (NS-2), <<http://www.isi.edu/nsnam/ns/>>.
- [7] QualNet simulator, <<http://www.scalable-networks.com/>>.
- [8] IEEE 802.11-2007, Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications, 2007.
- [9] C. Reis, R. Mahajan, M. Rodrig, D. Wetherall, J. Zahorjan, Measurement-based models of delivery and interference in static wireless networks, in: Proceedings of the ACM SIGCOMM, Pisa, Italy, 2006.
- [10] J. Boer, H.V. Bokhorst, W.J. Diepstraten, A. Kamerman, R. Mud, H.V. Driest, R.J. Kopmeiners, Wireless LAN with Enhanced Capture Provision, US Patent 5987033, November 16, 1999.

- [11] C. Ware, J. Chicharo, T. Wysocki, Simulation of capture behaviour in IEEE 802.11 radio modems, in: Proceedings of the IEEE VTC'01 Fall, 2001.
- [12] G. Judd, P. Steenkiste, Using emulation to understand and improve wireless networks and applications, in: Proceedings of the USENIX NSDI, Boston, MA, USA, 2005.
- [13] R. Chandra, R. Mahajan, T. Moscibroda, R. Raghavendra, P. Bahl, A case for adapting channel width in wireless networks, in: ACM SIGCOMM '08, 2008.
- [14] S. Ganu, K. Ramachandran, M. Gruteser, I. Seskar, J. Deng, Methods for restoring MAC layer fairness in IEEE 802.11 networks with physical layer capture, in: Proceedings of the ACM REALMAN'06, Florence, Italy, 2006.
- [15] Wistron NeWeb CM9, <http://www.wnweb.com/wireless/wireless_mini-pci.htm>.
- [16] AiroPeek NX wireless LAN analyzer, <<http://www.wildpackets.com/>>.
- [17] Madwifi receiver sensitivity problem, <<http://madwifi-project.org/ticket/705>>.
- [18] L. Verma, S. Kim, S. Choi, S.-J. Lee, Reliable, low overhead link quality estimation for 802.11 wireless mesh networks, in: Proceedings of the IEEE WiMesh, San Francisco, CA, USA, 2008.
- [19] Bug patch for Madwifi receiver sensitivity problems in OFDM modulations, <http://madwifi-project.org/attachment/ticket/705/ani_sensitivity_fix.diff>.
- [20] J.T. Bevan, J. Thomson, B. Baas, E.M. Cooper, J.M. Gilbert, G. Hsieh, P. Husted, A. Lokanathan, J.S. Kuskin, D. Mccracken, B. Mcfarl, T.H. Meng, D. Nakahira, S. Ng, M. Rattehalli, J.L. Smith, R. Subramanian, L. Thon, H. Wang, R. Yu, X. Zhang, An integrated 802.11a baseband and MAC processor, in: Proceedings of the IEEE International Solid-State Circuits Conference (ISSCC), San Francisco, CA, 2002.
- [21] K. Jamieson, B. Hull, A. Miu, H. Balakrishnan, Understanding the real-world performance of carrier sense, in: Proceedings of the ACM SIGCOMM E-WIND Workshop, 2005.
- [22] J. Ryu, J. Lee, S.-J. Lee, T. Kwon, Revamping the IEEE 802.11a PHY simulation models, in: ACM MSWiM '08, 2008.
- [23] H. Chang, V. Misra, D. Rubenstein, A general model and analysis of physical layer capture in 802.11 networks, in: Proceedings of the IEEE INFOCOM, Barcelona, Spain, 2006.
- [24] T.S. Rappaport, Wireless Communications: Principles and Practice, second ed., Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 2002.
- [25] D. Xu, J. Zhang, X. Gao, P. Zhang, Y. Wu, Indoor office propagation measurements and path loss models at 5.25 GHz, in: Proceedings of the IEEE VTC'07 Fall, 2007.
- [26] R. Jain, D. Chiu, W. Hawe, A Quantitative Measure of Fairness And Discrimination For Resource Allocation in Shared Computer Systems, Technical Report TR-301, DEC Research Report, September 1984.
- [27] J.J. Metzner, On improving utilization in ALOHA networks, IEEE Trans. Commun. COM-24 (4) (1976) 447–448.
- [28] T. Nadeem, L. Ji, A. Agrawala, J. Agre, Location enhancement to IEEE 802.11 DCF, in: Proceedings of the IEEE INFOCOM, Miami, USA, 2005.
- [29] K. Whitehouse, A. Woo, F. Jiang, J. Polastre, D. Culler, Exploiting the capture effect for collision detection and recovery, in: The Second IEEE Workshop on Embedded Networked Sensors, EmNetS-II, 2005.
- [30] D. Son, B. Krishnamachari, J. Heidemann, Experimental study of concurrent transmission in wireless sensor networks, in: Proceedings of the ACM SenSys, 2006.
- [31] Chipcon, <<http://focus.ti.com/lit/ds/symlink/cc1000.pdf>>.
- [32] J. Lee, S.J. Lee, W. Kim, D. Jo, T. Kwon, Y. Choi, RSS-based Carrier sensing and interference estimation in 802.11 wireless networks, in: Proceedings of the IEEE SECON, San Diego, CA, 2007.
- [33] N. Baldo, F. Maguolo, M. Miozzo, M. Rossi, M. Zorzi, ns2-Miracle: a modular framework for multi-technology and cross-layer support in network simulator 2, in: ICST ValueTools '07, 2007.
- [34] Q. Chen, F. Schmidt-Eisenlohr, D. Jiang, M. Torrent-Moreno, L. Delgrossi, H. Hartenstein, Overhaul of IEEE 802.11 modeling and simulation in ns-2, in: ACM MSWiM '07, 2007.
- [35] C. Ware, J. Judge, J. Chicharo, E. Dutkiewicz, Unfairness and capture behaviour in 802.11 adhoc networks, in: Proceedings of the IEEE ICC, New Orleans, LA, 2000.
- [36] X. Yang, N. Vaidya, On physical carrier sensing in wireless ad hoc networks, in: Proceedings of the IEEE INFOCOM, Miami, FL, USA, 2005.
- [37] H. Zhai, Y. Fang, Physical carrier sensing and spatial reuse in multirate and multipath wireless ad hoc networks, in: Proceedings of the IEEE INFOCOM, Barcelona, Spain, 2006.
- [38] T.-Y. Lin, J.C. Hou, Interplay of spatial reuse and SINR-determined data rates in CSMA/CA-based, multi-hop, multi-rate wireless

networks, in: Proceedings of the IEEE INFOCOM, Anchorage, Alaska, 2007.

- [39] K. Pelechrinis, G. Yan, S. Eidenbenz, S.V. Krishnamurthy, Detecting selfish exploitation of carrier sensing in 802.11 networks, in: Proceedings of the IEEE INFOCOM, Rio De Janeiro, Brazil, 2009.

- [40] V.P. Mhatre, K. Papagiannaki, F. Baccelli, Interference mitigation through power control in high density 802.11 WLANs, in: Proceedings of the IEEE INFOCOM, Anchorage, Alaska, 2007.



Jeongkeun Lee is a research scientist at Hewlett-Packard Laboratories in Palo Alto, CA, USA. He received B.S. and Ph.D. degrees from the School of Computer Science and Engineering, Seoul National University, in 2001 and 2007 respectively. His research interests include wireless LANs, wireless mesh networks, and enterprise and datacenter networking. <http://www.hpl.hp.com/people/jklee/>.



Jiho Ryu received the B.S. degree in computer science (major) and mathematics (minor) from Korea Advanced Institute of Science and Technology (KAIST), Daejeon, in 2005. Currently, he is working towards the Ph.D. degree at the School of Computer Science and Engineering, Seoul National University, Seoul, Korea. His research interests include wireless mesh networks, wireless LANs, sensor networks, and Real-time systems.



Sung-Ju Lee is a Senior Research Scientist at Hewlett-Packard Laboratories in Palo Alto, CA, USA. He received his PhD in Computer Science from University of California, Los Angeles (UCLA) in 2000. His research interests include Wireless LAN, wireless multi-hop networks, network management, and large scale networks. He has published nearly 80 papers in the area of computer networks and systems. He serves as a technical program committee member of various leading networking conferences, and is a standing committee member of IEEE SECON (Conference on Sensor, Mesh and Ad Hoc Communications and Networks). He is a senior member of ACM and IEEE. <http://networking.hpl.hp.com/sjlee>.



Ted Taekyoung Kwon received the B.S., M.S., and Ph.D. degrees from the Department of Computer Science Engineering, Seoul National University (SNU), Seoul, Korea, in 1993, 1995, and 2000, respectively. He is currently an Associate Professor at the School of Computer Science and Engineering, SNU. Before joining SNU, he was a Postdoctoral Research Associate at the University of California at Los Angeles (UCLA) and at City University of New York (CUNY). During his graduate program, he was a Visiting Student at IBM T. J. Watson Research Center and the University of North Texas. His research interest lies in sensor networks, wireless networks, IP mobility, and ubiquitous computing.