Revamping the IEEE 802.11a PHY Simulation Models

Jiho Ryu[†], Jeongkeun Lee[‡], Sung-Ju Lee[‡] and Taekyoung Kwon[†]
[†]School of Computer Science and Engineering
Seoul National University, Korea
{jhryu, tk}@mmlab.snu.ac.kr
[‡]Multimedia Communications & Networking Lab
Hewlett-Packard Laboratories, Palo Alto, USA
{jklee, sjlee}@hp.com

ABSTRACT

In simulating wireless networks, modeling of the physical layer behavior is an important yet difficult task. Modeling and estimating wireless interference is receiving great research attention, and is crucial in a wireless network performance study. The implementation of physical layer capture, preamble detection, and carrier sense threshold plays an important role in successful frame reception in the presence of interference. We showed in our previous testbed study that the operations of the frame reception and the capture effect in real IEEE 802.11a systems differ from those of popular research simulators. We present our modifications of the IEEE 802.11a PHY models to the current simulators. The modifications can be summarized as follows. (i) The current simulators' frame reception is based only on the received signal strength. However, the real 802.11 systems can start the 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) The current simulators set the carrier sense threshold equal to the receiver sensitivity. The standard however states that it should be 20 dB higher than the receiver sensitivity. We implement our modifications to the QualNet simulator and conduct a wireless network performance study to evaluate the impact of PHY model implementation.

Categories and Subject Descriptors

I.6.4 [**Simulation and Modeling**]: Model Validation and Analysis

General Terms

Design, Verification

Keywords

Simulations, Physical Layer Capture, Interference, Carrier sensing, IEEE 802.11a

MSWiM'08, October 27–31, 2008, Vancouver, BC, Canada.

1. INTRODUCTION

There have been extensive research efforts on analyzing the physical layer issues such as interference in wireless communications. The impact of interference on frame reception and throughput, however still needs further investigation. When we analyze the throughput of flows in wireless networks, the crucial problem is modeling the reception process at the physical (PHY) 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 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. Although these mechanisms are simple, the IEEE 802.11 standard, the widely used simulators (e.g., NS-2 [1] and QualNet [2]), and the real 802.11 chipsets do not operate the same way.

Recently, several studies reported the ramifications of the physical layer capture [3, 4, 5]; the capture effect is a physical layer mechanism by which the interference or collisions are dealt with. In wireless networks, the medium is shared by all the nodes. Thus, 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, the stronger frame at the receiver can be successfully received when the Signal-to-Interference-and-Noise Ratio (SINR) of the interested frame to the other frame(s) is higher than a *capture threshold.* 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 802.11 reception process.

From the previous measurement studies on IEEE 802.11 [3, 4], we learned that the capture effect works differently depending on the 802.11 chipset models. Kochut *et al.* [4], from their experiments with the wireless cards with the Prism chipset [6], discovered that the stronger frame that arrives during the reception of the weaker frame can be captured if the stronger frame arrives *within* the weaker frame's preamble time. With the Atheros chipset [7] however, the stronger frame can be captured even if it arrives *after* the weaker frame's preamble time [3].

Although the testbed study helped us to learn the different capture operations of the different chipsets, it is difficult to experiment all possible scenarios and topologies on

Figure 1: 802.11 PHY frame format.

a testbed. Thus, we modify the QualNet simulator (version 3.9.5) to investigate how the refined capture models affect the wireless network performance with various wireless network scenarios. Throughout this paper, we do not consider RTS/CTS since the capture of RTS frames are not different from the capture of MAC data frames [4].

The contributions of this paper are summarized as follows. First, we present the detailed model of IEEE 802.11a PHY reception and capture process. Second, different 802.11a chipsets exhibit different reception behaviors, depending on the implemented capture logic. We identify the two distinct capture models. Third, through the testbed experiments, we discover and modify the parts of the simulators that do not correctly reflect the behaviors of real 802.11a systems. Fourth, through the QualNet simulation, we show that different models of 802.11a reception process yield substantially different network performances.

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 the current simulators' reception model and how we modify the QualNet simulator. Section 4 presents performance evaluation using the modified QualNet simulator. In Section 5, we present the carries sense mechanism of IEEE 802.11a with simulation results. A survey of related work is given in Section 6. Finally, Section 7 concludes this paper.

2. 802.11A RECEIVE AND CAPTURE

In this section, we describe the 802.11a PHY receive procedure [8] and present three capture cases.

2.1 802.11a PHY Receive Procedure

As shown in Figure 1, an 802.11a frame begins with PLCP (Physical Layer Convergence Protocol) preamble that consists of OFDM training symbols. Upon receiving a transmitted PLCP preamble, the receiver (i) detects and measures the signal energy and (ii) synchronizes its timing with the training symbols: we call this the *preamble detection* process. If the preamble detection is successful, the receiver recognizes it as the start of a valid 802.11 frame transmission and searches for a PLCP header that follows the preamble. The PLCP header contains modulation/coding bitrate, frame length information and a parity bit. If the PLCP header reception is successful without any error detected by a parity, the receiver goes into a *receiving* state.

After the PLCP header, the MAC data follows and 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 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 bit-rate, 6Mbps, regardless of the bitrates of the MAC frame. The preamble training symbols are always the same for all frames. Thus, the successes of PLCP preamble detection and header reception are independent of the MAC frame bitrate while a higher SINR is required for a higher MAC bitrate to successfully pass the MAC CRC check.

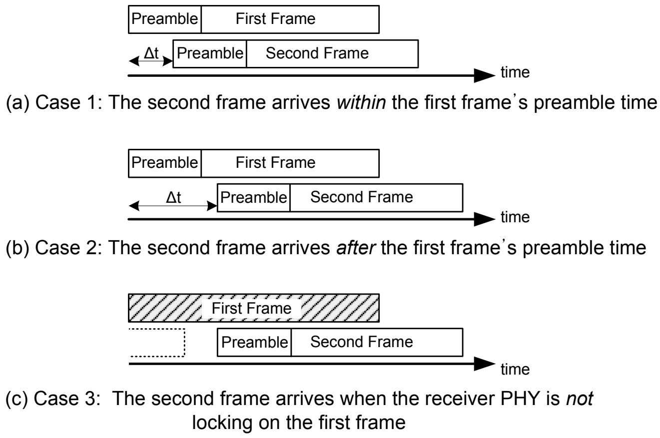

Figure 2: Three capture cases.

2.2 Capture Effect in 802.11a PHY

In this section, based on the measurement studies on capture timing and SINR threshold [3, 4, 5], we categorize the capture cases and describe how the 802.11 PHY processes frame capture in each case.

2.2.1 *Case 1: The Second Frame Arrives* within *the First Frame's Preamble Time*

As shown in the example of Figure 2(a), 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. In this case, the receiver has not yet completely locked onto the first frame. If the second frame's signal power is strong enough for the receiver to detect an energy increase above a certain threshold, which we call *capture threshold*, the receiver drops the first frame's preamble and tries to detect the second frame's preamble. We call it as SFC (Second Frame Capture). For ease of presentation, we define the *capture* as the PHY operation of selecting one frame to receive between the frames in a collision (the first vs. the second) regardless of the success of the selected (captured) frame reception (i.e., with or without 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. [1]

If the second frame's SINR (with the first frame signal as interference power) is high enough to receive second frame's preamble, PLCP header and MAC data without error, the second frame is successfully captured. Kochut *et al.* [4] and Lee *et al.* [3] showed that the capture within preamble time occurs with Prism [6] and Atheros [7] chipset wireless cards, respectively.

If the energy 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 as FFC (First Frame Capture). However as the second frame's signal energy in-

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

creases interference power, 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 with no PLCP header error or MAC CRC error.

The 'first' and 'second' frames do not necessarily indicate the exact frame arrival order. We address the situation when the receiver has locked on and is currently receiving a 'prior' frame; and subsequent frame(s) arrive during the prior frame reception. The newly arrived frame can be the second, third, or any other order as long as it arrives before the end of the prior frame reception. We just call it a '*second*' frame for ease of presentation.

2.2.2 Case 2: The Second Frame Arrives **after** the First Frame's Preamble Time

In this scenario, the second frame arrives after the first frame's preamble time as shown in Figure 2(b). The receiver has already synchronized its timing with the first frame and if the first frame's PLCP header has passed, the receiver is in the *receiving* state. In order to capture the second frame, message-in-message (MIM) mode should be implemented in 802.11 PHY [9, 10]. In the MIM mode, if the energy increase due to the second frame is above the capture threshold, the receiver drops the first frame and begins to synchronize its timing with the second frame preamble, i.e., SFC. If the MIM mode is not implemented or if the energy increase is below the capture threshold, the receiver retains its reception of the first frame, i.e., FFC. In either the first or second frame capture case, the un-captured frame signal increases the interference power and can affect the MAC CRC check of the captured frame.

The implementation of MIM mode is chipset-dependent. In our measurement study [3], we showed that the second frame that arrives later than the first frame preamble time can be captured with Atheros chipsets that are believed to implement the MIM mode. However, when we tested with the Prism chipset wireless 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.* [4], in their measurement study with Prism chipset cards, also report that the second frame capture can happen only when it arrives *within* the preamble of the first frame.

2.2.3 Case 3: The Second Frame Arrives when the Receiver is **not** Receiving the First Frame

In Case 2, the receiver has already synchronized with (and may have received) the first frame and need to drop the first frame to capture the newly arrived second frame: chipset must implement the MIM mode to allow the second frame capture. In Case 3 however, the receiver is in *idle* or *sensing* state, and it can start the reception of the second frame without the MIM mode support, as long as the SINR for the second frame is high enough to detect the second frame preamble.

There are two scenarios 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 Figure 2(c). 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.

3. SIMULATOR MODIFICATIONS

NS-2 [1] and QualNet [2] are two of the most popular network simulators but their 802.11 PHY and MAC implementations do not precisely model the 802.11 frame reception process, especially the capture.

We first describe the reception models of NS-2 and QualNet. As the QualNet model is more precise than NS-2, we explain how to improve the QualNet model based on our observations of capture timing and capture threshold [3].

3.1 Current Simulator Models

The flow chart of the current NS-2 and QualNet RX models is illustrated in Figure 3(a). In the flow chart, we denote the frame that newly arrives at the receiver 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, we begin the flow chart from the moment at which a new frame arrives when the PHY state is *idle* or *sensing*[2] instead of the *receiving* state. The other case when a new frame arrives in the middle of the PHY *receiving* state will be dealt later in the flow chart.

When a NEW frame arrives, its Received Signal Strength (RSS), denoted by $rss(NEW)$, is compared with the RXSens (RX sensitivity, the minimum signal strength for a frame to be received). If $rss(NEW)$ is smaller 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.[3] 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 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 CPThres (capture threshold, 10 dB in NS-2), the RCV frame is captured and the NEW 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 (*int_power*+=$rss(NEW)$), computes the bit error rate of the RCV frame based on the RCV frame's bit rate and SINR, and appropriately generates 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 FFC (First Frame Capture) but not SFC (Second Frame Capture).

3.2 Modified Simulator Model

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

- **PHY0**: The current QualNet model. Only FFC (first

[2]When the PHY is in the *sensing* state, it can receive but can not transmit.

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

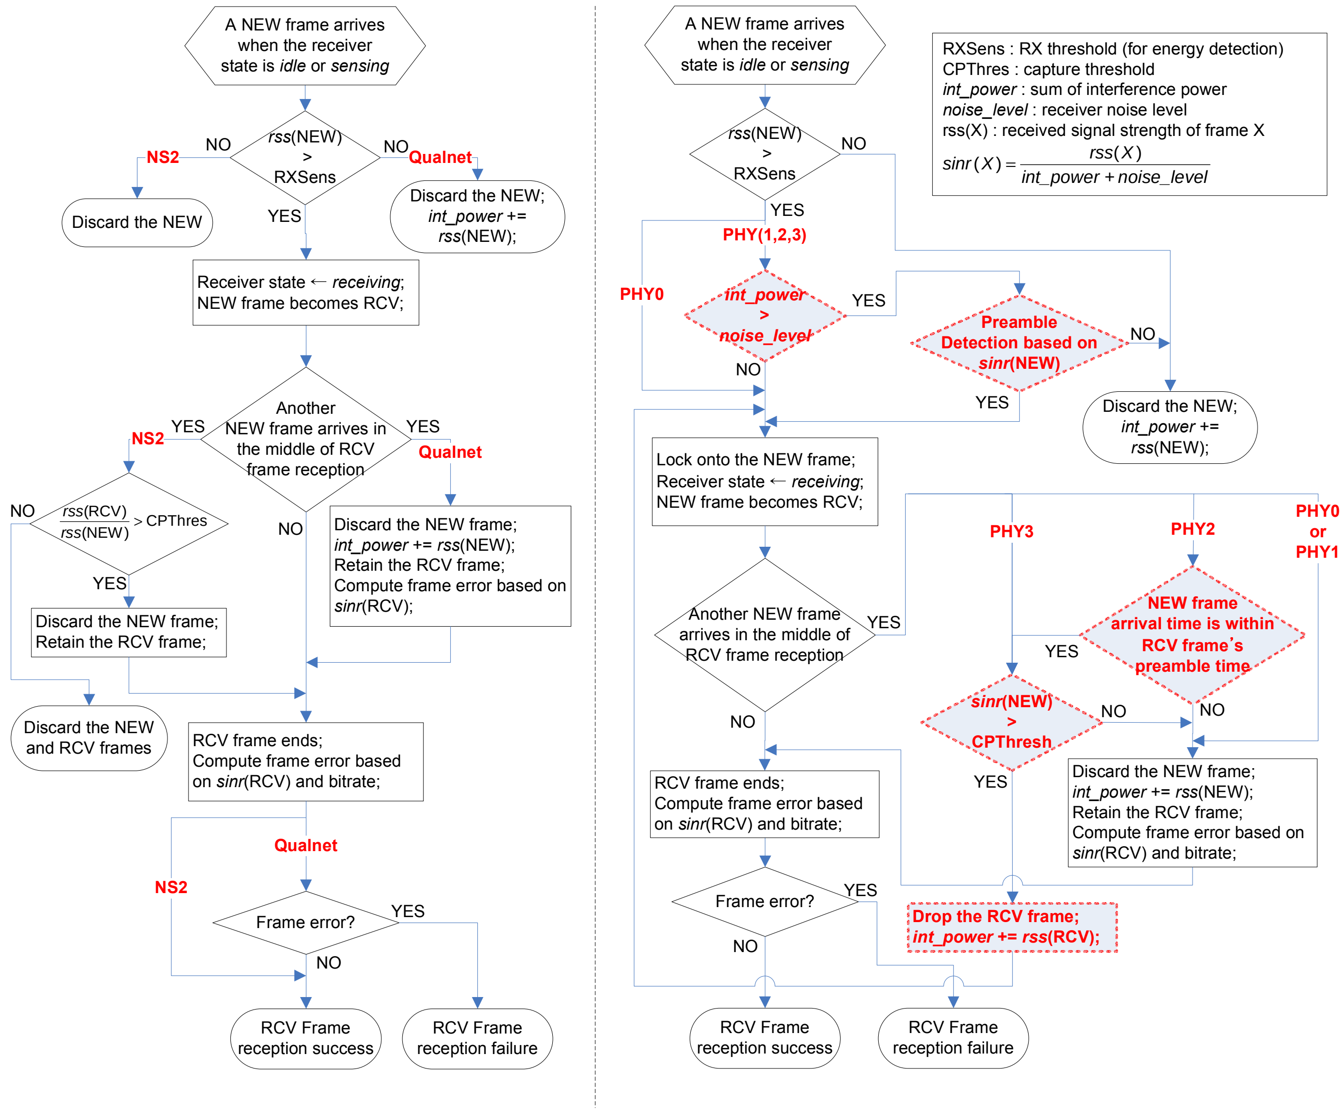

(a) Original Simulator Model (b) Revised Simulator Model

Figure 3: Simulator model flow charts.

frame capture) is implemented.

- **PHY1**: PHY0 + SINR based preamble detection.
- **PHY2**: PHY1 + SFC (second frame capture) *within* first (RCV) frame's preamble time.
- **PHY3**: PHY2 + SFC *after* first (RCV) frame's preamble time (MIM mode is supported).

The flow chart for the revised simulator model is presented in Figure 3(b). Our revision is highlighted by the dotted boxes.

3.2.1 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 header reception, that are vulnerable to interference. However, QualNet does not consider interference when deciding whether to receive a newly arrived frame or not. Thus, we revise the RX model to check $sinr(NEW)$ as well as $rss(NEW)$ before going into the *receiving* state.

In order to determine the SINR threshold for the preamble detection, we refer to [3] where the FRR (Frame Reception Rate) of 6Mbps transmission is observed to be a function of SINR in the Case 3 in Section 2.2. In this case, an ongoing transmission already exists when the NEW frame arrives and the success of the NEW frame capture is affected by whether the PLCP preamble detection and header reception of the previous frame have been successful or not. We observed that the FRR begins to rise above zero when SINR is 4dB. FRR reaches almost 100% when SINR is 10dB or higher while the FFC threshold for 6Mbps transmissions is about 0dB. Because the PLCP header and MAC data are encoded at the same 6Mbps bit rate, we reason that the relatively high (4∼10dB) SINR is required to detect (synchronize with) PLCP preamble in the presence of interference. Thus we determine the success of preamble detection based on a probability function linearly increasing from zero to one as the SINR increase from 4dB to 10dB.

Another observation from [3] is that as interference signal power becomes weak and close to the noise level, the interference power does not affect the frame capture.[4] Thus, we compare the interference power with the noise level and apply the preamble detection logic only if the interference power is greater than the noise level.

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).

3.2.2 Capture Models

If another NEW frame arrives during the reception of the RCV frame, we apply different capture algorithms based on the different PHY models. The PHY0 and PHY1 models follow the current QualNet implementation: always discard the NEW frame and treat its signal as interference power for the RCV frame. In the PHY2 model, if the NEW frame arrives after the RCV frame's preamble time, the NEW frame is discarded. If the NEW frame arrives within the RCV frame's preamble time and $sinr(NEW)$ is greater than the CPThres, 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 CPThres 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 original QualNet does not have the CPThres (as it does not support SFC), we use the measurement data of the Atheros wireless cards in [3], which reports that at least 10dB SINR is required to capture the NEW frame that arrives during the reception of a previous frame. Because the capture decision is made at the time of preamble detection, the 10dB CPThres is independent of the MAC data bitrate [3]. We conclude that Prism chipsets that we test follow the PHY2 model and Atheros chipsets follow the PHY3 model based on the reports from [3, 4] and our observations in the previous section.

3.3 Additional QualNet Modifications

3.3.1 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 previous studies (e.g., [11]) 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 [11].

3.3.2 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 of the tie in the backoff countdown race, 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 jitters contribute to timing skews. We observed that transmission time difference could be up to $4\mu s$ but in most cases fell into [-$2us$,$2us$] window. Thus, we randomly select skew time from [-$2us$,$2us$] and apply it before the start of each frame's transmission.

[4]FRR vs. SIR curve becomes similar to a clear channel reception curve (FRR vs. SNR) that is measured when there is no interference.

Table 1: Parameters used in simulation.

Parameter	Value
Path loss model	Two-ray (n=2) Log-distance (n=4)
Shadowing model	Constant
Shadowing mean	4dB
TX Power	16dBm
RX Sensitivity	-88dB
CS Threshold	-88dB
Rate adaptation mechanism	Auto Rate Fallback (ARF)
RTS/CTS	Disabled
TCP payload size	1400bytes

4. PERFORMANCE EVALUATION

4.1 Simulation Parameters

To evaluate the four PHY models (PHY0$\sim$PHY3) using QualNet, we consider two radio environments: indoor and outdoor radio propagation models. For the indoor propagation model, we use the log-distance model [12, 13]. Here, we set the path-loss exponent (n) to 4. For outdoor propagation model, we use the two-ray path-loss model with $n = 2$ [12]. Multi-path fading is not considered for either model. With TX power and RX sensitivity in Table 1, the maximum transmission ranges at 6Mbps in outdoor and indoor propagation models are $328.6m$ and $48.5m$, 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 increases. To remove the effect of routing, we intentionally arrange the sender and the receiver of each flow to be only one hop away.

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) [14] for the rate adaptation mechanism. We set a terrain size to $1000 \times 1000m^2$ for the outdoor propagation model and $149 \times 149m^2$ for the indoor propagation model.[5] We equally divide the simulation terrain into the cells (whose number is the same with that of sender-receiver flows) and randomly place each sender within each cell. We choose the distance between a sender and its receiver in a range from $10m$ to $57m$ in outdoors with a uniform distribution. Likewise, the distance between the sender and the receiver ranges from $3.5m$ to $20m$

[5]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 meters $\simeq$ 48.538 meters $\times$ 1000 meters $\div$ 328.602 meters.

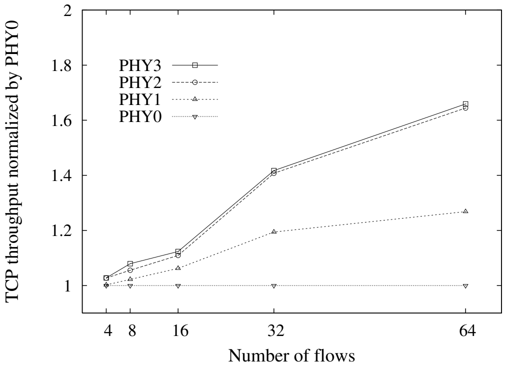

(a) Outdoor propagation (n=2)

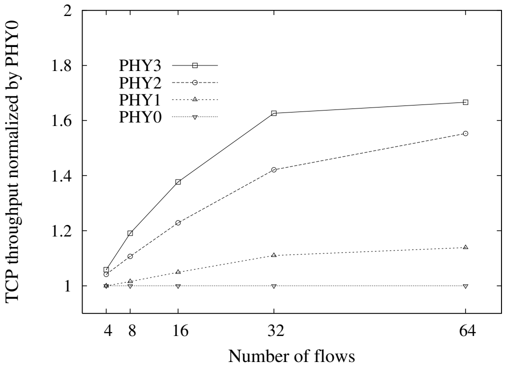

(b) Indoor propagation (n=4)

Figure 4: Average aggregate TCP throughput normalized by PHY0.

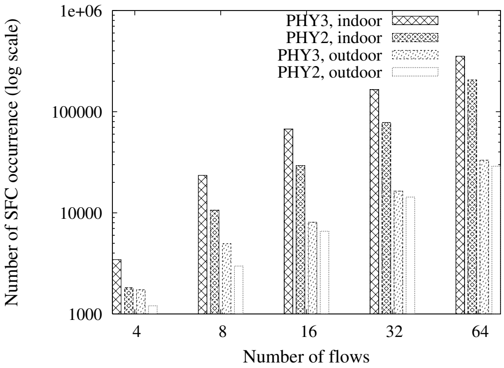

Figure 5: Number of SFC occurrences of PHY3 and PHY2.

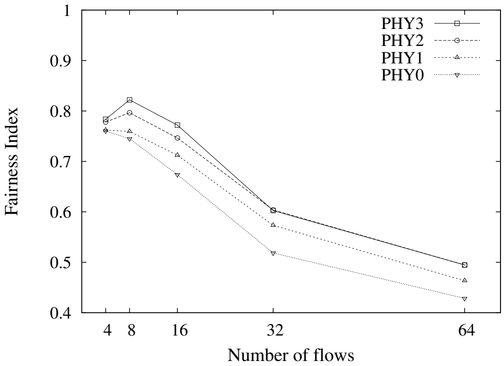

Figure 6: TCP flow fairness (indoor).

uniformly for the indoor propagation models.[6] To plot the performance metric for each number of flows, we average 30 different network topologies (or 30 different sender-receiver placements).

4.2 Performance Comparison

In Figure 4(a), we normalize the aggregate TCP throughputs of the four PHY models with reference to PHY0 as the number of sender-receiver pairs varies in the outdoor propagation model. The aggregate throughputs of the real chipset models (PHY2 and PHY3) are higher than that of the current model of the simulator (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 Figure 5, the number of SFC occurrence difference between PHY2 and PHY3 is also small in the outdoor propagation model. Since the path loss exponent is only 2, it is not easy for an interferer to be outside of the sender's CS (Carrier Sense) range and yet have the SINR at the receiver to be higher than the 10dB CPThres. Simply adding the SINR-based preamble detection logic (PHY1) yields a notable gain over PHY0. The preamble detection logic prevents the PHY layer from going into the *receiving* state upon receiving a useless frame and allows the PHY layer to transmit its own. More detailed explanations will be given in Section 5.2). Overall, the performance gain increases as the number of flows increases. In Figure 4(b), we plot the normalized TCP throughput in the indoor propagation model. Since the path loss exponent is higher than in the outdoor model, it is more likely that an interferer is out of the sender's CS range and the SINR at the receiver is higher than the CPThres. Figure 5 shows the large difference of the SFC occurrences between PHY2 and PHY3 and it verifies our conjecture. Note that the y-axis of Figure 5 is in log-scale. In the indoor propagation model, a signal attenuates more rapidly over distance than in the outdoor propagation model. Hence, even though there are more hidden interferers, the sender's frame can be successfully captured at the receiver with a higher probability. Therefore, PHY3 achieves a noticeable throughput gain over PHY2 in Figure 4(b).

Figure 6 shows the Jain's fairness index [15] among the flows in the indoor propagation model. 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 see the impact of the 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 Figure 7, we plot the MAC efficiency as well as the aggregate TCP throughput to analyze the correlation between two metrics.

[6]The largest distance from which the packets are received at 54Mbps in the absence of interference is 57 meters for the outdoor and 20 meters for the indoor propagation model.

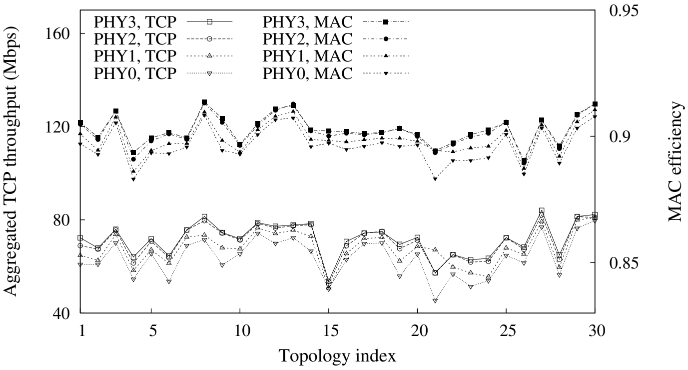

(a) Outdoor propagation (n=2)

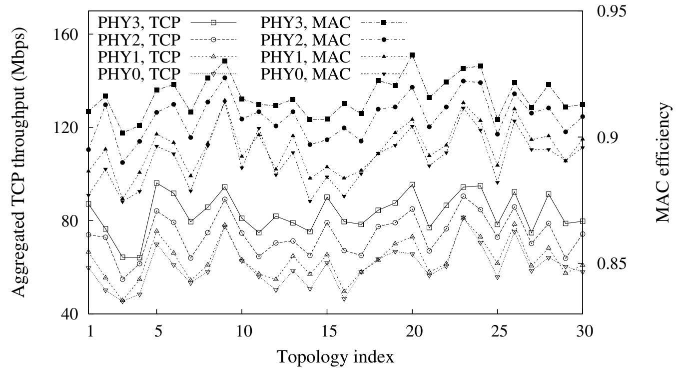

(b) Indoor propagation (n=4)

Figure 7: TCP throughput and MAC efficiency (16 flows).

We measure the metrics for 30 instances of the two radio propagation models, each of which consists of 16 flows. Figures 7(a) and 7(b) plot both 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 multiple hops, MAC layer retransmissions may generate TCP timeouts, which will invoke even larger TCP throughput fluctuation. (Similar observations are made in [4].) 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.

5. 802.11A CARRIER SENSE

As described in Section 2.1, 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 ensures that it holds the carrier sense (CS) busy for the duration of the transmitted frame as indicated in the PLCP header. If the receiver detects the signal energy but the preamble portion was missed, the receiver holds the CS busy for any signal 20dB above the minimum receive sensitivity (RXSens). To summarize, (i) the 802.11a PHY determines the CS busy by the PLCP preamble (and header) detection and the energy detection, and (ii) the 802.11a standard defines the carrier sense threshold (CSThres) to be 20dB higher than the RXSens.

The implication of the 20dB-higher CSThres in 802.11a CS mechanism can be illustrated by comparing the two collision cases (Case 2 and Case 3) in Figure 2(b) and (c). In Section 2.2, we analyzed the difference between the two cases from the viewpoint of the receiver. Now let us consider the two cases from the viewpoint of the sender that wants to transmit its own frame during the transmission of the first frame (from another sender). Let us focus only on the first frame (and ignore the second frame) in Figure 2(b) and (c) as the capture effect is not our concern here. We assume the first frame is strong enough to be detected and received when there is no other interference ($rss(first) >= RXSens$). In Case 2, the sender correctly receives the first frame's preamble and PLCP header and holds the CS busy until the end of the first frame transmission, i.e., the sender defers its own transmission. In Case 3, we assume the transmission of the dotted frame (that can be from either the sender or another sender) has just ended. We have two different consequences depending on $rss(first)$. If $rss(first)$ is smaller than the CSThres, the sender determines the channel is idle and transmits its frame when its backoff counter becomes zero. If $rss(first)$ is larger than the CSThres, the sender determines the channel is busy and defers its transmission.

5.1 Simulator Revision for Carrier Sense

The current QualNet compares the sum of all currently transmitting frames' signal power and noise level to the RXSens for carrier sensing. First, we revised QualNet to use a separate parameter (CSThres), which is configurable in the configuration file. Second, we revised QualNet to consider only the frame signal power without the noise level because the 802.11a standard defines the thresholds (RXSens and CSThres) in dBm scale that is independent of the noise level.

If the CSThres is the same with the RXSens (for example, in the current 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 CSThres is (much) higher than the RXSens, the success or failure of the preamble detection greatly affects the carrier sensing. As discussed in section 3.2.1, we revised QualNet to have the SINR-preamble detection logic.

5.2 Simulation Results

To evaluate the impact of the revised carrier sensing model on the network performance, we run simulations with three difference CSThres settings: (i) the CSThres is equal to the RXSens (current QualNet setting), (ii) the CSThres is 10dB higher than the RXSens (Atheros setting), and (iii) the CSThres is 20dB higher than the RXSens (802.11a standard setting). The Atheros setting is based on our measurement from the Atheros-chipset-embedded testbed. In the testbed experiments, we controlled the frame transmission timing of the two senders and created the setting where the sender always misses the preamble of the other sender's frame so

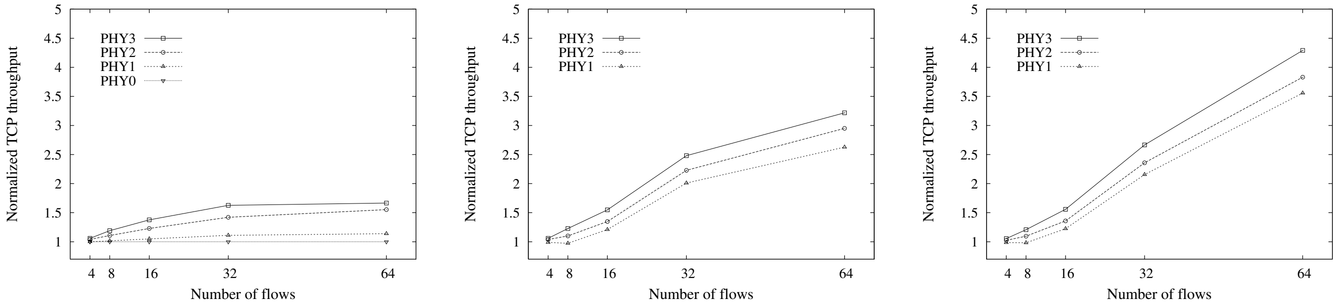

(a) CSThres=RXSens (QualNet) (b) CSThres=RXSens+10dB (Atheros) (c) CSThres=RXSens+20dB (standard)

Figure 8: Aggregate TCP throughput of indoor propagation model with different CSThres settings. Normalized by PHY0(CSThres=RXSens).

that the sender could rely only on the energy detection for carrier sensing. The sender deferred its own transmission when the other sender's frame was 10dB stronger than the minimum RSS for a successful frame reception.

In Figure 8, we show the aggregate TCP throughput of indoor propagation model for each CSThres setting. The TCP throughput is normalized by the TCP throughput of the current QualNet (PHY0, CSThres=RXSens). The throughput improvement of PHY1∼PHY3 over PHY0 increases as the number of flows increases and the CSThres increases. Not only does the standard setting result in substantial performance improvements (Figure 8(c)), but the Atheros setting (with the 10dB increase) also shows considerable improvements (Figure 8(b)).

In particular, the PHY1's normalized throughput gains in the Atheros setting (up to 263%) and in the standard setting (up to 356%) are much greater than that of the QualNet setting (up to 111%). Recall that the PHY1 model implements the SINR-based preamble detection logic on top of PHY0. Suppose an 18Mbps frame (frame A) from node A arrives at node B when node B is in the *idle* state. The RSS of frame A is higher than the RXSens but the SINR at node B is lower than 4dB. In the PHY0 model, node B starts to receive this frame and goes into the *receiving* state; but this frame will be eventually corrupted because at least 10dB SINR is required to decode the 18Mbps frame without error. Thus, the current simulator's lack of SINR-based preamble detection logic can cause an unrealistic and adverse PHY operation. For example, in the current simulator (PHY0 model), if another frame (frame C) arrives at node B with a high SNR (>10dB) during the transmission of frame A, node B can not receive frame C because node B is in the *receiving* state. In the PHY1∼PHY3 models with the preamble detection capability, node B does not detect the preamble of frame A due to the low SINR and will be able to receive the frame C because node B does not go into the *receiving* state. Moreover, if the RSS of frame A is smaller than the CSThres (that is very likely because frame A signal is weak and the CSThres is much higher than the RXSens), node B stays in the *idle* state, i.e., node B can transmit its own frame inspite of frame A. Hence, the preamble detection logic prevents the PHY from going into the *receiving* state upon the arrival of a useless frame (frame A) and enables the PHY to receive a more strong and useful frame (frame C) or transmit its own frame.

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

Based on the simulation results, we believe that the 802.11a CS mechanism can benefit from the better channel utilization than the 802.11b CS mechanism that sets the CSThres more conservatively (same or lower than the RXSens). Since the CSThres is 20dB higher (100 times stronger) in 802.11a, stations are allowed to transmit more aggressively than the 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,[7] we observed that the MAC efficiency in the increased CSThres settings does not decrease from that of the QualNet setting (CSThres=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 FFC and SFC), the increased transmission attempts in the 802.11a networks result in better spatial reuse especially when the number of flows is large.

6. RELATED WORK

There have been several simulation studies on the capture effect. However, their simulation models do not precisely reflect the capture process in real 802.11 systems. The most recent research work on improving the capture simulation model is the paper by Kochut et al. [4]. From the Prism chipset-embedded 802.11 testbed measurement, they observe that a frame with strong signal power can be captured if it arrives (i) before a frame with weak signal power or (ii) after the weak frame's arrival but within the weak frame's preamble time. They have modified NS-2 to account for this SFC within first frame's preamble time. Chang *et al.* [11] analyze the aggregated network throughput when there are concurrent multiple flows. In their simulation, the authors modify QualNet to reflect the SFC within the first frame's preamble time [4]. Both [4] and [11] considered only the single hop carrier sense range; the capture with hidden interferers was not considered. In this paper, we consider the capture with hidden interference and show the performance of the capture model that supports the SFC after preamble time (MIM mode). Along with the SINR-based

[7]The MAC efficiency graph of the improved CS models is not included due to space limit.

preamble detection logic, our simulator modifications better reflect the actual systems.

Both Ware *et al.* [10, 16] and Ganu *et al.* [17] 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 each other. In addition, [10, 17] consider only the case of two senders with one common receiver: it is hard to generalize the throughput unfairness problem into multi-hop wireless networks.

There are a number of studies that evaluate the effect of the carrier sense (energy detection) threshold on the network throughput performance [18, 19, 20]. They suggest to tune the carrier sense threshold to maximize the network performance considering various factors such as transmission power, bit rate, MAC overhead, interference, etc. Similar to our observations, [18, 20] show that the use of a small carrier sense range (i.e., high carrier sense threshold) can enable 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.

7. CONCLUSION

In this paper, we tried bridging the gap between the IEEE 802.11a PHY simulation models and the real wireless network systems. We focused on the preamble detection, physical layer capture, and carrier sensing. 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 simulation experiments with four different PHY models in various scenarios. Our results show that our modified model can increase the aggregated TCP throughput of more than 400% compared with the current QualNet model.

8. ACKNOWLEDGMENTS

This work was supported in part by the IT R&D program of MKE/IITA [2007-F-038-02, Fundamental Technologies for the Future Internet] and the KOREN program of National Information Society(NIA) Agency, Korea. The ICT at Seoul National University provided research facilities for this study. We thank Hoon Chang for helping us solve the QualNet timing problem.

9. REFERENCES

[1] The Network Simulator 2 (NS-2), http://www.isi.edu/nsnam/ns/

[2] QualNet simulator, http://www.scalable-networks.com/

[3] Jeongkeun Lee, Wonho Kim, Sung Ju Lee, Daehyung Jo, Jiho Ryu, Taekyoung Kwon and Yanghee Choi, An Experimental Study on the Capture Effect in 802.11a Networks, In Proc. ACM WiNTECH, Montreal, Canada, Sept. 2007.

[4] A. Kochut, A. Vasan, A.U. Shankar and A. Agrawala, Sniffing out the correct Physical Layer Capture model in 802.11b, In Proc. IEEE ICNP, Berlin, Germany, Oct. 2004.

[5] Glenn 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.

[6] PRISM Wireless LAN, Intersil Corporation, http://www.intersil.com/globespanvirata/

[7] Atheros Communications, http://www.atheros.com/

[8] IEEE 802.11a, *Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications: High-speed Physical Layer in the 5 GHz Band.*

[9] Jan Boer et al., Wireless LAN with Enhanced Capture Provision, US Patent 5987033, Nov. 16, 1999.

[10] Christopher Ware, Joe Chicharo and Tadeusz Wysocki, Simulation of Capture Behaviour in IEEE 802.11 Radio Modems, In Proc. IEEE VTC'01 Fall, Oct. 2001.

[11] Hoon Chang, Vishal Misra and Dan Rubenstein, A General Model and Analysis of Physical Layer Capture in 802.11 Networks, In Proc. IEEE INFOCOM, Barcelona, Spain, Apr. 2006.

[12] Theodore S. Rappaport, Wireless Communications: Principles and Practice, 2nd edition, 2002, Prentice-Hall, Inc., Upper Saddle River, NJ, USA.

[13] Ding Xu, Jianhua Zhang, Xinying Gao, Ping Zhang and Yufei Wu, Indoor Office Propagation Measurements and Path Loss Models at 5.25 GHz, In Proc. IEEE VTC'07 Fall, Oct. 2007.

[14] A. Kamerman and L. Monteband, WaveLAN-II: A High-Performance Wireless LAN for the Unlicensed Band, Bell Labs Technical Journal, Vol.2 no.3, pp. 118-133, Summer 1997.

[15] R. Jain, D. Chiu and W. Hawe, A Quantitative Measure Of Fairness And Discrimination For Resource Allocation In Shared Computer Systems, Technical Report TR-301, DEC Research Report, Sept. 1984.

[16] Christopher Ware, John Judge, Joe Chicharo and Eryk Dutkiewicz, Unfairness and Capture Behaviour in 802.11 Adhoc Networks, In Proc. IEEE ICC, New Orleans, LA, June 2000.

[17] Sachin Ganu, Kishore Ramachandran, Marco Gruteser, Ivan Seskar and Jing Deng, Methods for Restoring MAC Layer Fairness in IEEE 802.11 Networks with Physical Layer Capture, In Proc. ACM REALMAN, Florence, Italy, May. 2006.

[18] Xue Yang and Nitin Vaidya, On Physical Carrier Sensing in Wireless Ad Hoc Networks, In Proc. IEEE INFOCOM, Miami, FL, USA, Mar. 2005.

[19] Hongqiang Zhai and Yuguang Fang, Physical Carrier Sensing and Spatial Reuse in Multirate and Multihop Wireless Ad Hoc Networks, In Proc. IEEE INFOCOM, Barcelona, Spain, Apr. 2006.

[20] Ting-Yu Lin and Jennifer C. Hou, Interplay of Spatial Reuse and SINR-Determined Data Rates in CSMA/CA-Based, Multi-Hop, Multi-Rate Wireless Networks, In Proc. IEEE INFOCOM, Anchorage, Alaska, May 2007.