

Sniffing out the correct Physical Layer Capture model in 802.11b

Andrzej Kochut, Arunchandar Vasam, A. Udaya Shankar, Ashok Agrawala *
Department of Computer Science, University of Maryland, College Park, Maryland 20742
{kochut, arun, shankar, agrawala}@cs.umd.edu

Abstract

Physical layer capture (PLC) in 802.11b refers to the successful reception of the stronger (higher signal strength at receiver) frame in a collision. PLC causes significant imbalance in the throughputs of sources. Existing 802.11b simulators, including ns2 and Qualnet, assume that PLC occurs only if the stronger frame arrives first at the receiver. We show empirically that in reality PLC occurs even if the stronger frame arrives later (but within the physical layer preamble of the first frame). Consequently, throughput unfairness in reality can be significantly (up to 15%) higher than with the former PLC model. We have modified the ns2 simulator to account for this and Qualnet will be incorporating a fix in their next release.

To identify which frames were involved in collisions, when their transmissions started, and which of them were retrieved, we have devised a novel technique using multiple sniffers and instrumented device drivers to reconstruct from the air interface all tx/rx events in a WLAN to within 4 μ s accuracy. This allows us to quantify the causal links from the PHY layer through the MAC layer to the observed application layer imbalance. It also shows that the arrival times of colliding frames routinely differ by as much as 20 μ s due to inherent uncertainties of 802.11b firmware clock synchronization and rx/tx turnaround delays, and that the frame to arrive first can be either the stronger or the weaker with equal likelihood.

1. Introduction

Consider two 802.11b stations, A and B, transmitting data to a wired sink S through access point C, in the DCF (Distributed Coordination Function) mode under the fol-

* This work was supported in part by the Maryland Information and Network Dynamics (MIND) Laboratory, Fujitsu Laboratories of America, and by the Department of Defense through a University of Maryland Institute for Advanced Computer Studies (UMIACS) contract. Access to the Qualnet network simulator was provided by Scalable Networks Inc., under the Qualnet University Program.

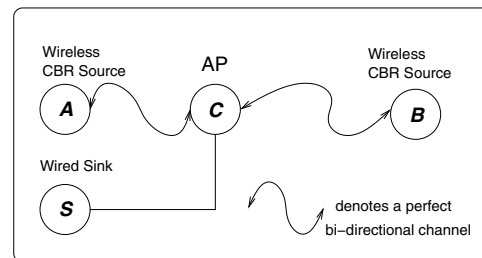


Figure 1. Scenario showing imbalance in throughputs obtained by sources though their channel to the AP is perfect.

lowing natural conditions: (1) each of A and B have a “perfect” channel to C when the other is not transmitting, i.e., gets the maximum throughput at the highest bitrate; (2) A’s signal at C is stronger than B’s signal at C by a few dB; and (3) A, B, and C are all in range of each other, i.e., no hidden terminal. The workload considered is a UDP source on each station with a constant rate sufficient to saturate the available capacity. The scenario is shown in Figure 1.

Our experiments, done with a variety of 802.11 cards, show that A’s goodput *as well as throughput* are consistently higher than B’s in the long term (over intervals longer than, say, 10 seconds). Here, goodput refers to the number of successful transmissions per second and throughput to the number of (re)transmissions per second, including lost tries.

Because A and B achieve equal performance when the other is not transmitting, the imbalance can only be because of interaction between A and B when both are transmitting simultaneously. Indeed, conditions 1-3 above allow us to conclude that the imbalance can only be because collisions are resolved in favor of A, rather than, say, both frames being lost. The phenomenon of the stronger frame in a collision being received successfully is known as *physical layer capture (PLC)* [1].

Note that the imbalance cannot be caused by B’s frames having a poor signal-to-noise ratio, because B experiences a perfect channel when A is not transmitting. The imbalance also cannot be caused by B switching to a lower trans-

mission bitrate, as that would only reduce each station's throughput by same amount. This is because, as explained in detail in [2], the 802.11 MAC guarantees fair access to the media (in terms of the number of transmissions) in the long term and does not abort a transmission in progress (no matter what the bitrate).

Here are possible models for what happens when frames of different signal strength collide at a receiver:

- I Both frames are lost.
- II The stronger frame is received correctly provided it is sufficiently stronger and it is the first to arrive.
- III The stronger frame is received correctly provided it is sufficiently stronger, regardless of whether it arrives before or after the weaker frame.

Clearly, we can rule out model I because it does not generate the observed imbalance. Both models II and III generate imbalance in favor of A, model III more so than model II. To the best of our knowledge, existing analytical models for 802.11 MAC throughput [3, 4, 5] assume model I. Simulators such as ns2 [6] and Qualnet [7] use (deterministic or probabilistic versions of) model II. *We show that in reality, collisions are resolved according to model III.*

The 802.11b standard [8] does not preclude this kind of capture (see PHY receive process in pages 202-203). Each frame starts with 192 bits of Physical Link Control Protocol (PLCP) preamble and header, consisting of 128 sync bits, 16-bit Start Frame Delimiter (SFD), and 48 bits indicating, among other things, the bitrate and length of the MPDU (data) that follows. These 192 bits are always transmitted at 1 Mbps. If a stronger frame arrives while a receiver is receiving a frame, the stronger frame can be captured provided it starts before the SFD of the first frame.

Naturally, the difference between models II and III with respect to throughput imbalance is most significant when the stronger frame almost always starts second; we demonstrate such a simulation scenario in Section 3, showing throughput imbalance of 15% with model III as opposed to 0% with model II. The PLC imbalance is typically magnified by protocol control mechanisms as one goes higher up in the protocol stack: the MAC's Binary Exponential Backoff (BEB) algorithm magnifies the raw PLC imbalance by increasing the backoff window of the losing station; the TCP congestion control magnifies the imbalance at the MAC layer by throttling the send rate; and so on.

We also show that the time instants of start of reception of colliding frames at the receiver routinely differ by as much as 20 microseconds, and that the frame to start reception first can be either the stronger or the weaker frame with equal likelihood. This happens even when all stations are in range of each other, due to the rx/tx turnaround delay at stations and inherent uncertainties in 802.11 firmware

clock synchronization. The time differences can be more if hidden terminals exist.

To demonstrate these behaviors in real-life 802.11 networks, we have developed a technique for reconstructing from the air interface the timeline of all transmission (including retransmission) and reception events in an 802.11 WLAN to an accuracy of $4\mu\text{s}$. Hence we can identify *all* the frames that collided, their arrival times, and whether any of them were received (due to physical layer capture). Combining this with an instrumentation of the WLAN device driver and the TCP/UDP/IP subsystem, we are able to trace the impact of physical layer capture through the protocol stack, and thereby explain the observed imbalance in the application layer throughputs.

1.1. Impact of our findings on simulators

ns2, probably the most detailed network simulator in the open source domain, uses a deterministic version of model-II; if the earlier of two colliding frames is stronger than the later by 1 dB, it is captured. Qualnet, a state-of-the-art commercial network simulator, uses a probabilistic version of model-II; if the earlier frame is stronger, it treats the later frame's signal as noise for the earlier frame, computes the resulting bit error rate (this depends on the earlier frame's bitrate and modulation), and appropriately flips bits of the earlier frame.

Neither simulator allows the stronger frame to be captured if it starts later than the weaker frame, and it seems reasonable to expect the same of other 802.11-capable network simulators. Consequently, the throughput imbalance predicted by simulation can differ significantly from the imbalance observed in reality, for example, the imbalances can differ by 15% with ns2. The difference can be greatly reduced by changing the simulator's collision handling to that of model III, and we have shown this for ns2. We discuss this further in Section 3. We have informed the maintainers of ns2 and the Qualnet support staff about this issue, and Qualnet will incorporate a fix in their next release [9].

1.2. Identifying collisions and their resolution

How does one determine what happens in a collision when one or both frames may be lost? To identify the correct collision model (I, II, or III), we need to identify in each collision precisely which frames were involved in the collision, which frame arrived first, which frame was stronger, and how the collision was resolved (were both frames lost or was one of them received).

Clearly, we need to generate a global timeline of all transmissions (including retransmissions) and all receptions in the WLAN. This is not trivial for several reasons. The information needed for this is available only at the firmware

level or the air interface (e.g., the number of retransmissions a frame underwent is not available at the device driver). Access to the firmware code is not feasible (in our case, because of the stiff license fee). Hence, we need to *sniff* packets in the air.

However, a single sniffer cannot capture everything because it too is subjected to PLC (so the frame retrieved by it may, in fact, be the one lost by the intended receiver). So we need many sniffers, which then leads to the problem of merging the logs of different sniffers on a common timeline, i.e., synchronizing the timestamps from the various sniffers to an accuracy adequate to resolve collisions. We solve this problem, achieving timing accuracy of better than $4\mu\text{s}$, by using timestamps generated by the firmware and synchronizing them via receptions of beacon frames. For the details of our method please see [10].

1.3. Causal link between PLC and higher-layer throughput imbalance

Several papers in the literature (e.g., [11, 12]) have pointed out that difference in signal strengths at the receiver is accompanied by unfairness in application layer throughput. They correlate application layer throughput to signal strength, but they do not explain *why* it happens. They do not examine or quantify the causal links from the PHY layer through the MAC layer to the application layer.

We instrument wireless stations at the lowest layer available to us – the device drivers of the WLAN cards and the TCP subsystem in the Linux kernel. We first isolate the effect of PLC by itself by examining an ad-hoc network in the broadcast mode, *thereby disabling the MAC backoffs and retransmissions*. Then, we consider how the MAC magnifies the PHY layer imbalance. We explicitly identify the frames which collided and how they were resolved at the receiver, and relate the observed difference in this to the throughputs obtained by different stations.

1.4. Summary of contributions

Our main contributions are as follows:

- We identify from real-life measurements that model III is the correct model for physical layer capture. Because the throughput and goodput imbalance caused by model III is significantly higher than that caused by models II and I, this has a practical consequence to simulators, including ns2 and Qualnet, which use model II. Incorporating collision model III in ns2 makes it more realistic. Qualnet will incorporate this fix in their next release.
- We develop a novel technique for reconstructing from the air interface a time line of all transmission (including

retransmission) and reception events in an 802.11 WLAN to an accuracy of $4\mu\text{s}$. This allows one to identify all the frames that collided, their arrival times, and how collisions were resolved at a receiver. For details see [10].

- We present a detailed analysis of the causal link between PLC and application layer throughput imbalance. This explains why a station with a perfectly fine channel to a sink suffers an unfair degradation in service when another station whose frames have stronger signal strength at the sink starts to compete with it.

1.5. Roadmap

Section 2 surveys related work. Section 3 identifies issues with simulator models and suggests fixes in light of our findings. Sections 4, and 5 describe our experiments and our inferences. Section 6 concludes.

2. Related work

Related work falls into two groups. The first group considers models for physical layer capture. A survey of communications literature shows that existing analytical models for obtaining probability of frame capture typically consider the stronger frame starting first (e.g., [13, 14, 15]). However, there are patents, including one on the WaveLan precursor of 802.11 [16], that also consider the stronger frame starting later.

To the best of our knowledge, reference [17] is the only work which tries to identify what capture model is used in 802.11. This work simulates different capture models [13, 14, 15, 16] to identify which of these explain the extent of unfairness observed. Thus, it indirectly infers the appropriate capture model, but does not provide any real-life evidence. Our work, on the other hand, provides concrete evidence of the correct capture model by the use of a novel technique that reconstructs all collisions and their resolutions.

The second group of related work is real-life measurement based studies which consider unfairness in 802.11 throughput. References [11] and [12] do not consider the impact of PLC, but focus on MAC-layer performance. Reference [18] analyzes location-dependent throughput performance, but does not explain the mechanism leading to this. Reference [19] is closest to our work. It correlates the TCP throughput to signal strength at the macroscopic level in presence of a hidden terminal effect only. Specifically, the authors neither identify collisions and PLC nor provide evidence of how collisions are resolved and how this impacts higher-layer protocols. We are able to precisely identify and quantify the impact of PLC by itself and how it is magnified by higher layers.

3. A better PLC model for simulators

In this section we first make some observations from our empirical measurements. Then we use these to point out the scenario where there is a significant difference between predicted throughput imbalances with capture models II and III.

Consider a collision of RTS frames transmitted by two stations A and B at a receiver. Let T_A denote the time instant when the reception of the RTS frame transmitted by source A starts, and T_B denote that of source B. Figure 2 shows a histogram of $T_A - T_B$ for all collisions at the receiver obtained from a typical real-life experiment (described in Section 5). Source A is the stronger, and almost all collisions (99 %) are resolved in its favour. From these results, we make the following observations:

- The stronger frame is almost always retrieved even if it starts second.
- The arrival times of colliding frames routinely differ by as much as $20\mu s$, and the frame to arrive first can be either the stronger or the weaker with approximately equal probability.
- RTS frames are also successfully captured.




Figure 2. Histogram of the difference between starting time instants of two colliding RTS frames in a typical experiment.

In light of these observations, we identify issues with the PLC model commonly used in simulators. We restrict citing source code to the ns2 simulator (version 2.26) due to licensing concerns, and briefly discuss Qualnet at the end of this section. Our experimental results have helped uncover two issues in the current modeling of physical layer capture in ns2.

3.1. PLC II vs PLC III

In a collision of data frames, ns2 does not display capture of the stronger frame if it starts second, contrary to ex-




Figure 3. Topology used for ns Experiment

```
// pktRx_ is the first frame
// p is the second frame
if (pktRx_->rxPower/p->rxPower > CPTresh-
old)
    capture(p);
else
    if (p->rxPower/pktRx_->rxPower > CPTresh-
old)
        capture(pktRx_);
else
    collision(p);
```

Figure 4. Code deciding capture in ns2 with our suggested improvement (in box).

perimental results. The pseudo-code for ns2 capture (paraphrased from the file `mac/mac-802_11.cc`) is outlined in Figure 4. Here, `pktRx_` refers to the first frame in the collision and `p` refers to the second frame. If the ratio of `pktRx_`'s signal strength to `p`'s signal strength is greater than `CPTresh` (which is set to 10, or equivalently, a difference of 1 dB), then the routine `capture` is executed.

We now identify a scenario where using the existing model causes significant deviation from reality. We simulate in ns2 the topology shown in Figure 3. We disabled RTS/CTS to avoid another problem described below. The distances A C and BC are 50 meters and 200 meters, respectively. The workload consists of two equal rate CBR sources, one on each station, with a total rate enough to saturate the available capacity.

Because transmissions are slotted and ns2 uses one global clock, all the frames in a collision start transmission at the same time. Hence, the time difference between the receptions of the two frames is determined solely by the propagation delay. In our scenario, A's frame arrives at C before B's frame.

Case 1: Suppose A and B transmit with the same power. Then A's signal is stronger at C than B's signal at C by more than 1 dB (in ns2, the signal strength is inversely proportional to a power of the distance, either $1/d^2$ or $1/d^4$). So the stronger signal arrives before the weaker signal and the existing ns2 code resolves all collisions in favor of A. The resulting difference in their long-term throughputs is 14.45%.

Case 2: Suppose B transmits with high enough power such that the ratio of signal strength of B at C to that of A at C is higher than the capture threshold of 1dB. Now the stronger signal starts second at C, and therefore, the existing ns2 code loses both signals in a collision. *NS-2 shows no unfairness in this scenario, contrary to what is observed in reality. It would be more realistic to observe the same unfairness as in case 1, but in B's favor.*

Further, the histogram in Figure 2 shows that it is not the propagation delay (which is of the order of nanoseconds) which decides which frame arrives first at the receiver, but rather the inherent uncertainties of clock synchronization (drift, missed beacons, etc.) and rx/tx turnaround time. Therefore, if one were simulating a wireless technology that has capture model II, then our results show that it is crucial to accurately model these timing jitters.

3.2. Issue II : lack of RTS capture

In a collision of RTS frames, ns2 does not display capture behavior at all (i.e., even if stronger starts first), contrary to experimental results. Suppose two RTS frames collide in ns2. If the stronger frame is second, then both frames are lost (as discussed in Issue I). If the stronger frame is first, then it is captured (as it should be). However, the ns2 capture sub-routine (paraphrased in Figure 5) wrongly marks the medium as busy for the duration of the *second* frame *plus an additional EIFS duration*. Consequently, the CTS response to the first RTS is never sent, effectively resulting in the first RTS frame being dropped. *However, our experiments described in the Section 5 show that when RTS/CTS frames are used, the stronger RTS frame always receives a CTS frame in response.*

```
set_nav(txtime(p) + eifs.);
free(p);
```

Figure 5. Capture sub-routine in ns2.

Issue I can be fixed by minor modifications to the existing code and this results in an unfairness of 15%, which reflects reality better. Issue II has been taken up with ns2 maintainers for further debugging.

As mentioned before, Qualnet uses a probabilistic capture model if the stronger frame starts first, and therefore is affected by Issue I. Subsequent to our communication, their support staff have scheduled a fix to their simulator in a future release.

4. Isolating the unfairness caused by PLC

The goal is to analyze the impact of physical-layer capture effect, all by itself, on 802.11b performance. For this, we consider WLAN networks in adhoc mode in which stations *broadcast* data, i.e., with destination MAC address FF:FF:FF:FF:FF:FF. Because there is no unicast, there are no ACKs and hence no timeouts, retransmissions, or binary exponential backoff, and the transmission bit rate is fixed at 2Mbps (the highest bit rate for broadcast). Thus all MAC-layer mechanisms are eliminated except for carrier sensing (and the initial backoff). (Adhoc mode is needed because in infrastructure mode, the AP retransmits broadcasts from clients.)





Figure 6. Approximate locations of nodes A, A', B, B', and C. The drawing is not to scale.

We use two stations, A and B, broadcasting in the adhoc mode, and two sniffers, A' and B', one placed close to each source. A third sniffer, C, acts as a sink. For all experiments reported in this paper we use identical IBM Thinkpad R-32 laptops running customized RedHat Linux and equipped with Prism2 [20] chipset wireless cards. For the details of hardware and software configurations used as well as our modifications to the kernel software please see [10].


Figure 6 shows positions of sources and sniffers in an indoor setting for the first scenario. We use UDP sources at A and B (because TCP sources cannot broadcast). Each UDP source broadcasts 10,000 data packets at a constant sending rate chosen such that the firmware queue is always non-empty. Each UDP packet is 1472 bytes, this includes a sequence number. This results in a 802.11 frame payload of 1500 bytes (after adding 20-byte IP header and 8-byte UDP header), which is the Ethernet-based MTU used by current WLAN device drivers (although the WLAN standard allows 2304 bytes). Sniffer C logs all of the packets that were successfully received.

To verify that both sources are in range of the sink, we performed experiments in which only one source was sending at a time. In all cases, the sink received all 10,000 packets sent by the source correctly, thus verifying that each source has a perfect channel to the sink. Likewise, we ver-

ified that the sources are in range of each other by pinging one from another. Experiments were repeated with several cards, and we observed consistent results.



(a)



(b)

Figure 7. Received signal strength (in dBm) vs. time (in sec.) (a), and goodput (in KBps) vs. time (in sec.) (b) for a typical ad-hoc mode experiment.

Metrics: We use two sets of metrics, one quantifying the extent of the physical-layer capture, and the other quantifying the impact of the effect on higher-layer performance. For each source, let N denote the total number of frames transmitted by the source, let S denote the number of these frames that were not received at the sniffer, and let L_i , for $i = 1, \dots, N$, be an indicator variable that is 1 if frame i was involved in a collision (according to our collision analysis).

Our metrics that quantify *extent of capture* for each source are:

- **Fraction of transmitted packets that collided:**

$$F_C = (\sum_i L_i) / N$$




Figure 8. Non-overlapping distributions of signal strengths of A and B at C.

- **Fraction of transmitted packets that were lost:**

$$F_L = S/N$$

Our metrics for estimating the *impact on higher-layer performance* are:

- **Throughput:** the number of “transmission-complete” interrupts over time.
(These interrupts are generated by the WLAN card and logged by the source driver. Because there are no ACKs, a “transmission-complete” interrupt does not imply successful transmission. Hence the throughput metric really quantifies the source’s share of the media access.)
- **Goodput:** the number of frames of the source received by the sink over time.
- **MAC Service Time:** The average time between successive transmission-complete interrupts.

(Because the firmware queue is never empty, the MAC service time of a frame is the time from the previous frame’s transmission-complete interrupt to this frame’s transmission-complete interrupt.)

Results of the experiments: Figure 7a represents the time evolution of the signal strength of the two source stations as perceived by the sink sniffer in a typical experiment. Each point in the figure is the averaged value of signal strengths of 50 frames around the frame received at that time instant. Figure 8 shows the corresponding histogram.

Figure 7b depicts the goodputs achieved by both sources during one of the experiments over time. The stronger source, A, has higher goodput than the weaker source, B. The overall goodput (computed over the entire experiment) for the stronger source was typically **12% higher** than that of the weaker source.

Station	Signal (dBm)	Num. of packets that			Fraction of pkts that			Goodput (KBps)	Service time (msec.)	Throughput (KBps)
		were sent	were lost	collided	collided (F_C)	were lost (F_L)	collided and lost			
A	8.64	10,000	6.23	862.12	8.62	0.06	0.64	116.68	12.54	117.73
B	-19.51	10,000	870.67	862.12	8.62	8.70	100.00	104.20	12.80	117.91

Table 1. Results of ad-hoc mode experiments. Each number is an average computed over several experiments. Fractions have been multiplied by 100. Observe that the average goodput of station A is larger even though its average throughput is nearly the same as that of station B.




Figure 9. Throughput (KB/s) vs. time (sec.). Note, that throughput of both stations is nearly equal.

The time evolution of the instantaneous throughputs is shown in Figure 9. The difference in the instantaneous throughputs varies on the short-term between the two sources, while the overall throughputs differ by **less than 1%**. The service times for the two sources (Table 1) differ by **less than 2%**.

Because there are no ACKs and retransmissions, the throughput seen by the source represents the rate of media-access. Because there is PLC and subsequent loss for the weaker station, the goodputs seen at the sink sniffer differ.

Table 1 presents the results of our collision and loss analysis of the experiments. F_C and F_L are very close (differ by 0.9%) for the weaker source showing that most of the frames that collided were lost by the weaker source, while F_L of the stronger source is close to zero. Nearly all of the collisions (99%) were resolved in favor of the stronger source. The average goodput of station A is larger, even though its average throughput is nearly the same as that of station B.

About 44% of the cases the stronger source started second. The packets of the weaker source that were lost are almost always (99%) the ones for which our global-timeline analysis detected collision.

Summary:

- In about 44% of the collisions, the stronger frame started second but was still captured by the receiver.
- Nearly 100% of collisions were resolved in favor of the station with the stronger signal.
- The resulting long-term goodput imbalance between the two sources was as high as 12%.
- The resulting long-term throughput imbalance between the two sources was practically zero.

5. How higher layers magnify the unfairness

The second round of experiments was designed to quantify the amount by which higher layers magnify unfairness caused by PLC. We present detailed results and analysis for two sources. Results for multiple sources are available in [10]. Experiments were conducted in configuration identical to the one shown on Figure 6, except that in place of node C we placed access point (AP) and a sniffer close to it. We use laptop with Prism2 card that supports *HostAP* mode as an access point. We verified that there is a perfect communication channel between each pair of stations.

Metrics for UDP experiments: For a source, let N be the total number of RTS frames transmitted by the source, and let M be the total number of CTS frames received by the source. Each CTS frame corresponds to an RTS frame that was received. Thus $N \geq M$, and $N - M$ is the number of RTS frames that were transmitted and lost. N and M can be obtained from the log of the sniffer associated with the source.

For a source, our metrics for measuring the *extent of capture* are:

- **Fraction of RTS losses:**

$$F_L = (N - M) / N$$

- **Fraction of RTS collisions:** Let L_i , for $i = 1, \dots, N$, be an indicator variable, which is 1 if the RTS frame i is detected as having collided by our global collision analysis and 0 otherwise.

Protocol	Station	Signal (dBm)	Fraction of packets that				RTT (msec.)	Goodput (KBps)				Goodput difference (%)
			collided (F_C)		were lost (F_L)			link layer		appl. layer		
			DATA	RTS	DATA	RTS		mean	variance	mean	variance	
UDP	A	8.72	0.01	4.78	0.01	0.48	N/A	347.23	3.88	330.45	3.72	18.70
	B	-19.53	0.01	4.82	0.02	4.92	N/A	292.45	1.67	278.56	1.57	
TCP	A	8.82	0.01	4.69	0.01	0.38	106.76	245.52	3.26	234.89	3.12	25.13
	B	-19.44	0.02	4.78	0.02	4.79	136.65	195.37	2.85	187.53	2.73	

Table 2. Analysis of collisions, losses and achieved goodputs during UDP and TCP experiments in the infrastructure mode with RTS/CTS enabled. Each value is an average over several experiments and fractions have been multiplied by 100. Signal strength is measured at the access point.

Protocol	Transmission 1	Transmission 2	% lost by both	% won by transmission 1	% won by transmission 2
UDP	A → AP	B → AP	4.54	95.46	0.00
TCP	AP → A	A → AP	100.00	0.00	0.00
	AP → A	B → AP	5.16	94.84	0.00
	AP → B	A → AP	15.59	84.41	0.00
	AP → B	B → AP	100.00	0.00	0.00
	A → AP	B → AP	5.68	93.82	0.50

Table 3. Analysis of collision resolution for UDP and TCP experiments in infrastructure mode with RTS/CTS. Each value is an average over several experiments. A→B denotes A sends to B.

$$F_C = \sum_i L_i / N$$

Our metrics for measuring the *impact on higher layer performance* are:

- **Link Layer Goodput:** The number of unique IP-level frames successfully transmitted per unit time. This is obtained from the device-driver log of the source by counting the number of “successful transmission” interrupts over time. Because ACKs are used, an interrupt showing “successful transmission” indicates that the frame got through.
- **MAC Service Time:** This is obtained from the device-driver in the same manner as that for the ad-hoc mode experiments. In this case however, this includes the time durations due to retransmissions and binary exponential backoff.

Results of UDP experiments: The time evolution of the signal strength remains the same as the Figure 7a for ad-hoc experiments. Figure 10 shows the link-layer goodput of the stations over time.

The stronger source finishes first for the same number of packets. The overall goodput of the stronger source is **18.7%** higher. The average MAC service time of the weaker




Figure 10. Link-layer goodput (in KBps) vs time (sec) of two sources in a typical UDP experiment in infrastructure mode with RTS/CTS.

source is 4.96 ms and that of the stronger is 4.19ms, a difference of **18.37%**.

Tables 2 and 3 present the detailed results of the experiments. The values in the tables were obtained by averaging results of multiple experiments. The variation of each metric is very small. Nearly 5% of the RTS frames sent by each source collided at the AP. 95.46% of the collisions were resolved in favor of station A. In about 53% of the collisions,

the stronger station, *A*, started sending later than *B*. The data from one of these experiments was presented as a histogram on Figure 2 in Section 3. F_C was close to F_L for the weaker station, showing that it lost most of the collisions.

Metrics for TCP experiments: We use the following to quantify the impact of PLC on TCP performance:

- Congestion Window
- RTT as observed by TCP
- TCP level loss
- TCP level throughput, as reported by *net-perf*.
- Firmware Service Time (instead of MAC Service time used for UDP)

Firmware service time is the time between submission of frame to the firmware and the successful completion indication. This includes the queuing delay involved in the firmware and is more representative of the end-to-end delay that affects TCP performance.




Figure 11. Link-layer goodput vs time of two stations during a typical TCP experiment in the infrastructure mode with RTS/CTS.

For the details of the instrumentation used to obtain metrics involving TCP state please see [10].

Results of TCP experiments: The signal strengths of the sources at the AP is the same as that in Figure 7a. The time evolution of the link-layer goodputs is shown in Figure 11. All numbers are averages over several experiments. The difference in the long-term link-layer goodput is **29%** and the application layer throughput (reported by *netperf*) difference is **28.57%**.

After slow start, the congestion windows of both TCP sources were at 23 frames (corresponding to 32K advertised receiver buffer) throughout the experiments, except for the very few losses (0.02%) for the weaker station. This implies that the throughputs obtained by the application is effectively determined by the RTT.

The firmware service times of the sources are shown in Figure 12. The X-coordinate of a point in the graph represents the transmission completion time instant of a frame and the Y-coordinate represents the firmware service time of that frame.




Figure 12. Time evolution of the firmware service time (including queuing) for a TCP experiment (with RTS/CTS enabled).

The weaker source has higher instantaneous service time (most of the time) and the difference in the average per-frame firmware service time is **27%**.




Figure 13. TCP's RTT of both stations in the experiment with RTS/CTS.

Figure 13 shows the time-variation of the RTT values seen by TCP. A point in the graph denotes the RTT value in the Y-axis and the time instant the information was obtained in the X-axis. The weaker source has higher RTT in general and the average RTT of the weaker source is higher by **27.9%**. Table 2 summarizes the results.

Table 3 presents the detailed collision analysis of all pairs of simultaneous transmissions at various receivers. The results show that collisions are resolved in favor of the stronger source **93.82%** of the time at the AP. The stronger source also collides with transmissions from AP

to the weaker source and affect the weaker source's reception 15.59% of the time, while the weaker affects the AP's transmission to the stronger only 5.16% of the time.

Summary: The results of the experiments presented above confirm that in the infrastructure mode (with binary exponential backoff, ACKs, and retransmissions):

- Collisions are resolved in favor of the stronger station.
- RTS frames are successfully captured; the stronger station gets the CTS frame.
- In approximately 50% of the cases the stronger RTS frame starts second but is still received properly.
- MAC layer magnifies the discrepancy (due to binary exponential backoff) by approximately 10%.
- TCP may additionally magnify the difference. In our case it does not, because there are no TCP-layer losses as link-layer losses are masked by retransmissions. Therefore, unfairness seen at the MAC-layer propagates to the application-layer unchanged.
- Observed throughput difference at the application layer may be as high as 25% for two sources.

6. Conclusions and future work

Existing models for modeling collisions in 802.11 networks typically assume that either both colliding frames are lost or that the stronger frame is received successfully only if it starts first. In this work, we showed experimentally that in reality the stronger frame is retrieved irrespective of whether its reception starts first or second.

This result has significant impact on the modeling of the physical layer in discrete event simulators. Specifically, we have uncovered two issues in the ns2 simulator. The Qualnet simulator will also be modified to account for our results.

Finally, we examined the causal link between physical layer capture and application layer unfairness in great detail. This quantifies the impact of the physical layer capture on the perceived application layer unfairness by itself and examined how it is magnified by the MAC and higher layer control mechanisms.

Directions for future work include the design of control mechanisms at the MAC layer to compensate for the unfairness induced by the physical layer.

7. Acknowledgments

Samrat Bhattacharjee, Arunesh Mishra, Kevin Kamel, Bao Trinh, and Moustafa Youssef, of the Department of Computer Science, helped us greatly in technical issues as well as in conducting some of the experiments. Jonathan

Agre and Lusheng Ji of the Fujitsu Laboratories of America, College Park, provided some of the equipment needed for experiments. Renata Kochut gave valuable support and encouragement. The ICNP referees provided constructive feedback and pointed out some errors. We express our sincere thanks to all of them.

References

- [1] M.Soroushnejad and E. Geraniotis, "Probability of Capture and Rejection of Primary Multiple Access Interference in Spread Spectrum Networks," *IEEE Transactions on Communications*, vol. 39, no. 6, pp 986-994, 1991.
- [2] M. Heusse, F. Rousseau, G. Berger-Sabbatel, and A. Duda, "Performance anomaly of 802.11b," in *Proceedings of IEEE INFOCOM 2003*, San Francisco, USA, March-April 2003.
- [3] J. Bianchi, "Performance analysis of the IEEE 802.11 Distributed Coordination Function.," in *IEEE JSAC*, 2000.
- [4] K.C. Tay, Y.C. and Chua, "A Capacity analysis for the IEEE 802.11 MAC protocol," in *Wireless Networks*, January 2001.
- [5] Cali I., Conti M., and Gregori E., "IEEE 802.11 wireless lan: Capacity analysis and protocol enhancement," in *Proceedings of the IEEE INFOCOM*, 1998.
- [6] "NS2 network simulator," <http://www.isi.edu/nsnam/ns>.
- [7] "Qualnet simulator," <http://www.scalable-networks.com>.
- [8] The Institute of Electrical and Electronic Engineers Inc., "IEEE 802.11, 1999 edition (ISO/IEC 8802-11:1999)," <http://standards.ieee.org/getieee802/802.11.html>.
- [9] *Personal communication with Qualnet support staff*.
- [10] A. Kochut, A. Vasani, A. U. Shankar, and A. Agrawala, "Sniffing out the correct Physical Layer Capture model in 802.11b," *University of Maryland CS-TR-4583.1*, 2004.
- [11] B. Bing, "Measured performance of the IEEE 802.11 wireless lan," in *24th Conf. on Local Comp. Networks*, October 1999.
- [12] G. Xylomenos and G. Polyzos, "Tcp and udp performance over a wireless lan," in *IEEE Infocom*, 1999.
- [13] K. Cheun and S. Kim, "Joint delay-power capture in spread-spectrum packet radio networks," in *IEEE Tr. on Com.*, 1998.
- [14] D.H. Davis and S. Gronemeyer, "Performance of slotted aloha random access with delay capture and randomized time of arrival," in *IEEE Trans. on Communications*, 1980.
- [15] J.C. Arnbak, "Capacity of slotted aloha in rayleigh fading channels," in *IEEE JSAC*, February 1987.
- [16] US Patent 5987033, *Wireless LAN with Enhanced Capture Provision*, <http://www.uspto.gov>, November 1999.
- [17] C.Ware, J.F.Chicharo, and T.Wysocki, "Modelling of capture behaviour in IEEE 802.11 radio modems," in *Invited Paper, IEEE International Conference on Telecommunications*, June 2001.
- [18] B. Rathke, M. Schlager, and A. Wolisz, "Systematic measurement of tcp performance over wireless lans," in *Technical Report, Technical University Berlin TKN01BR98*, 1998.
- [19] C. Ware, J. Judge, J. Chicharo, and E. Dutkiewicz, "Unfairness and capture behaviour in 802.11 adhoc networks," in *IEEE International Conference on Communications*, 2000.
- [20] "Prism chipset documentation," <http://www.intersil.com/design/prism>.