Understanding Performance for Two 802.11b Competing Flows

(Revised on Apirl 7, 2007)
Kan Cai, Michael J. Feeley, Sharath J. George
Department of Computer Science
University of British Columbia
Vancouver, BC, Canada, V6T 1Z4
{kcai, feeley, sharathg}@cs.ubc.ca

***Abstract*— It is well known that 802.11 suffers from both inefficiency and unfairness in the face of competition and interference. This paper provides a detailed analysis of the impact of topology and traffic type on network performance when two flows compete with each other for airspace. We consider both TCP and UDP flows and a comprehensive set of node topologies. We vary these topologies to consider all combinations of the following four node-to-node interactions: (1) nodes unable to read or sense each other, (2) nodes able to sense each other but not able to read each other's packets and nodes able to communicate with (3) weak and with (4) strong signal. We evaluate all possible cases through simulation and show that, for 802.11b competing flows, the cases can be reduced to 11 UDP and 10 TCP models with similar efficiency/fairness characteristics. We also validate our simulation results with extensive experiments conducted in a laboratory testbed.**

I. Introduction

In IEEE 802.11, nodes regulate access to the airspace they share in a decentralized fashion using a CSMA/CA and random backoff protocol. Nodes with packets to send engage in an uncoordinated competition for channel access by delaying transmission until sender and receiver see clear air and by backing-off and re-transmitting when collisions occur. The goal of this approach is to share the common airspace fairly and efficiently without requiring centralized channel administration or direct coordination among peer nodes.

Unfortunately, in congested environments things often do not go according to plan. It has been shown that the protocol often exhibits unpredicted performance degradation [1, 2] and unfair channel allocation [1, 3, 4] due to the node topology and other environmental factors. This behavior is not currently well understood and the complexity of the environment and the decentralized nature of the protocol make understanding elusive. Nevertheless, as 802.11 popularity grows, congestion increases and emerging applications such as media streaming place new demands on network performance predictability, there is a growing need for a deeper understanding of how 802.11 deals with congested traffic in practice [5].

The early understanding of 802.11 was that sending nodes are confronted with two types of potentially competing nodes: *hidden-* and *exposed-terminals* [1]. The issue for the protocol is that the sender decides when to send, but it is the channel conditions at the receiver that determine successful delivery. Node hidden from the sender can cause corruption and nodes exposed to the sender, but hidden from the receiver may not. Recently, Chen et al. [4] extended this basic model to observe that sender-receiver node pairs can have *incomplete* or *inconsistent* views of network topology. They argue that incomplete information leads to network inefficiency while inconsistent information leads to unfairness. They do not show, however, what network conditions lead to one or the other of these problems.

The best attempt we know of to describe the conditions that lead to inefficiency and unfairness is by Garetto et. al. [6]. They model the behavior of a set of four nodes consisting of two competing UDP flows. They model node interaction as a binary condition on each node pair indicating whether the pair is within transmission range of each other. Their approach leads to 16 topologies, which they classify into one of the four categories based on similar performance characteristics.

This paper provides a new model of two-flow competition that extends this earlier work in three ways. First, we model traffic that can be sensed but not read. We show that typically at least 42% of the traffic a node senses is too weak to be read. The difference between readable and unreadable competing traffic is the amount of time the sender waits before attempting to send again [1]. The importance of this difference can be seen in the example in Figure 1, in which two 802.11b flows are either fair or unbalanced toward one or the other depending only on whether the two senders can read, sense or not sense each other's packets.

Second, we model both UDP and TCP traffic. The key difference between UDP and TCP is that TCP has a counter flow of transport-level ACK messages. We show that the presence of this counter flow is important to understanding the behavior of wireless congestion.

Finally, we consider flows where sender and receiver are close enough that the flow is resilient to noise generated by the competing flow.

In all, we characterize the two-competing flow scenario using 21 802.11b models (11 UDP and 10 TCP) and 19 802.11g models (9 UDP and 10 TCP) that predict network

[1]Note that 802.11b and 802.11g devices behave differently in the sensing state due to their PLCP header difference. Details will be provided in Section III.

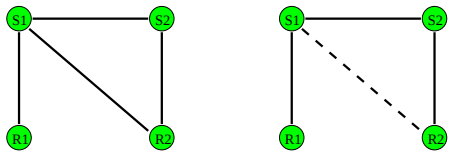

(a) Fair Allocation (b) Flow 2 Wins

Fig. 1. Example of difference between decoding (solid line) and only sensing (dashed line) packets.

performance based only on topology. We validate our models using simulation and experimentation in a lab testbed.

II. Related Work

The performance woes of 802.11 competing flows, i.e., inefficiency and unfairness, are well-known facts. Prior studies have succeed to identify one main culprit: problematic topologies. One attempt is to distinguish between hidden-terminal and exposed-terminal topologies [1–3]. Recent work from Chen et. al. further points out the importance of incomplete channel status assessment and inconsistent channel status [4]. Incomplete channel information leads to packet collisions; inconsistent channel information leads to unfair channel sharing. These categorizations are both correct, however, they are not specific enough to help wireless devices to adapt to continuously changing environments.

Self-adaptation in 802.11 networks has drawn a lot of attention recently to improve performance for 802.11 devices. Some have investigated physical carrier sensing in detail [7–11]. Their intention is to adjust or disable carrier sensing function so as to maximize spatial reuse and avoid packet collisions. In contrast, our goal is to understand the impact of sensing range on network performance compared to the other two link states, which provide reference models which indeed facilitate this sensing-range adaptation.

The other adaptation mechanisms include altering the MAC backoff durations [12, 13], enabling RTS/CTS virtual carrier sensing [4, 14], switching from sender-initiate mode to receiver-initiate mode [4, 15], and adapting the transmission rate and time scheduling [16, 17]. Our work is not to compete with any of these approaches but to suggest to them when and what adjustment they should take if they are aware of the model they belong to.

There have been several analysis works [18–21] on 802.11 MAC DCF protocol performance including throughput, delay, queue performance, etc. These models are, however, mostly analytical rooted at the Markov Chain model proposed by Bianchi [18]; recent work from Kim et al. [21] is based on a fluid model. Even though it is helpful to understand the theoretical limits of 802.11 networks, they are not instructive when severe performance degradation and unfairness problems occur. Our models, on the other hand, follow a more experimental methodology to examine the fairness and performance issues when competition occurs.

Our modelling study is inspired by the recent work from Garetto et al. [6]. To our best knowledge, their work is the only attempt so far to analyze two-competing-flow topologies, "the building block of any complicated scenarios" [6]. They investigate 16 topologies and divide them into four classes for analysis. They provide useful insights on long- and short-term unfairness through analytical analyses and simulations.

Our work substantially improves Garetto's work on a few aspects. First, we include more real parameters such as sensing range and TCP traffic in our analysis; we investigate 1296 scenarios rather than 16. Second, these added parameters result in many more models: 9 UDP models and 10 TCP models for 802.11g flows; 10 UDP models and 11 TCP models for 802.11b flows. We outline and explain, in full detail, the performance differences amid these models. Last, we simulate all possible topologies using Glomosim [22] and the results closely match our models. A real four-node testbed has been used to confirm the correctness of our analysis by examining a few of the interesting two-flow topologies.

III. Modeling Performance by Topology

Our work takes an empirical approach based on simulation and experimentation. We first devised a set of node-topology parameters and node-performance characteristics. We then simulated every combination of topologies characterized by these parameters and grouped them by common performance characteristics. Finally, we used testbed experiments to validate topologies whose performance differed from previous work or which constituted interesting inflection points in the performance-parameter space. This section describes these parameters, characteristics and our assumptions.

A. Model Parameters

We consider three parameters in our model. The first is link state. A two-flow scenario consists of four nodes with six links between them. Two of these links are between sender and receiver of the flows, which are assumed to be in transmission range. The other four links — between senders, between receivers, and between sender and receiver of different flows — can have one the three link states: out of range, in sense-range only, and in transmission range. The other two parameters are traffic type and flow robustness. Taken together these three parameters and their possible values yield a total of 1296 [2] distinct network/traffic-type topologies. The remainder of this section describes these three parameters in more detail.

a) Tri-State Link: There are three states between any pair of nodes depending on the transmission power, carrier sensing threshold and background noise: (1) *Transmission Range* (TR), in which a node can clearly receive a packet from the other node; (2) *Sensing Range* (SR), in which a node can only sense the signal from the other node, but is not able to capture its packet correctly; (3) *Out of Range* (OR), in which a node cannot sense any signal from the other node at all.

Receiving a packet or sensing a packet has different impacts on the length of delay before a node sends its next packet. When a node is in transmission range of another node, it is able to set its NAV (Network Allocation Vector) correctly and then use DIFS to contend for the airspace with the others. When it senses a packet whose *payload* it can not decode, however, it follows a different approach.

[2] 2 protocols (802.11b, 802.11g); 3 link states; 4 inter-flow links; 2 traffic type; 2 interference levels. The number of scenarios is $2 * 3^4 * 2 * 2^2$.

If the packet is sent by a 802.11b node and the sensing node is able to decode the preamble and the *whole* PLCP (Physical Layer Convergence Protocol) header that uses a more reliable 1Mbit/s or 2Mbit/s code modulation than the payload, then it will use EIFS, a longer period than DIFS, to hold back its transmission in order to avoid interfering with the MAC-ACK packet of the other flow. However, when the sensed traffic is actually a MAC ACK, this extended wait is unnecessary.

If the packet is sent by a 802.11g (or 802.11a) node that uses the standard ERP-OFDM (or OFDM) modulation scheme, then the sensing node is unlikely to decode the PLCP header if it cannot decode the payload. This is because *part* of the PLCP head, the SERVICE field, is also encoded by the higher-rate modulation that the payload uses. When a node fails to decode the PLCP header, it still uses DIFS *not* EIFS to schedule its next packet. Thus, if the sensed traffic is actually a MAC data packet, this sensing node might not back off long enough to avoid interference to the returning MAC ACK packet.

We have analyzed both scenarios: in one, the sensing node can decode the PLCP header successfully; in the other, the sensing node cannot. We conclude 21 models for 802.11b flows and 19 models for 802.11g flows. *Due to the space limitation, we can only present the 802.11b models in this technical report.* The 802.11g models are reported in the our submission to MASS 2007 [23].

b) Traffic Type: We investigate the performance of two traffic types in all the topologies: (1) *UDP traffic* (2) *TCP traffic*. The key difference between UDP and TCP is that TCP flows consist of two sub flows, i.e., the *TCP-DATA subflow* and the *TCP-ACK subflow*.

c) Flow Robustness: The distance between a flow's sender and receiver also plays an important role in a noisy environment. If the nodes are close enough, the signal strength at the receiver is strong enough that the flow is resilient to most noise, while a distant sender provides a weak signal that is vulnerable to noise. We examine two points along this signal-strength continuum using the *interference-level* parameter which takes on two possible values: (1) Interference-Susceptible and (2) Interference-Immune.

B. Performance Characteristics

We classified the simulation results according to two qualitative performance metrics: *fairness* and *communication efficiency*. The first metric has three values: fair or unfair with one or the other of the flows dominating. The other metric classifies interference between the flows by indicating whether there is interference and, if so, which packets conflict: data packets send by *flow senders*, ACK packets sent by *flow receivers* or one type from each flow.

C. Assumptions

Finally, we make three important simplifying assumptions. First, we assume that link conditions are symmetric; that is node A has the same view of B's traffic and B has of A's traffic.

Second, we restrict our analysis to *two* flows under the belief that pair-wise interference is common enough to warrant isolated study and under the hope that these results will provide a building block for analysis of more complex topologies such as mesh networks [6].

Third, we assume that nodes out of sensing range do not interfere with each other's traffic. This assumption is built into common network simulators and approximates expected behavior. However, at higher sending rates, the assumption does not hold, though it is likely rate-adaption schemes incorporated on most 802.11 adaptors would lower sending rate if faced with significant interference, even from an otherwise invisible node. Experiments conducted in our testbed indicate that at a low sending rates of 6 Mbps (802.11g), any signal above -70 dBm would not be vulnerable to noise from a node that could not be sensed.

IV. 802.11B MODELS

We simulated each of the 648 802.11b scenarios characterized by our model; details of the simulation are presented in Section V-B. As explained in Section III-A.0.a, we assume in these scenarios a sensing node is able to decode the PLCP header of the sensed packet and thus will use EIFS to schedule its next transmission. By grouping performance-similar topologies together, we derive 11 UDP and 10 TCP models.

A. UDP Models

The first 11 models are for competing UDP flows. Figure 2 provides a graphical representation of each model and its legend is shown in Figure 3. Note that, out of the 7 weak-signal models, 4 of them, (i.e., UM_4, UM_7, UM_8 and UM_9) are missed out in Garetto's models.

1) Base cases – Senders TR/SR: We begin with four initial models that two UDP senders can at least sense each other, with the addition of the trivial *independent* model. In these models, it is very unlikely that packet collisions will happen unless the backoff counters of both senders reach zero simultaneously. At times when one sender starts transmitting a data packet, the other sender will wait until the MAC-ACK transmission is over by setting its NAV to either the duration in the MAC header if it can decode that data packet or EIFS if it cannot.

a) UM_1: Independent: For completeness we begin with the trivial model in which two flows that are sufficiently distant from each other that neither flow affects the other.

b) UM_2: Symmetric, Senders TR/SR: In this model, the topology is *symmetric* (i.e., the two links that connect the two senders to the other receivers are of the same link state.) and thus bandwidth is evenly divided.

c) UM_3: Asymmetric (TR, OR), Senders TR/SR: Two models are needed to capture the behaviour when the topology is asymmetric in senders TR/SR scenarios. The first model covers the case where one inter-flow sender-receiver link is out of range while the other is in transmission range. In this case, even though the senders have different understanding of the network topology, bandwidth is still evenly distributed to the two flows. This is because, when two senders can at least sense each other, there is actually no difference between two scenarios where a sender can decode the MAC-ACK packet of the other flow or cannot read it at all. In either case, a sender

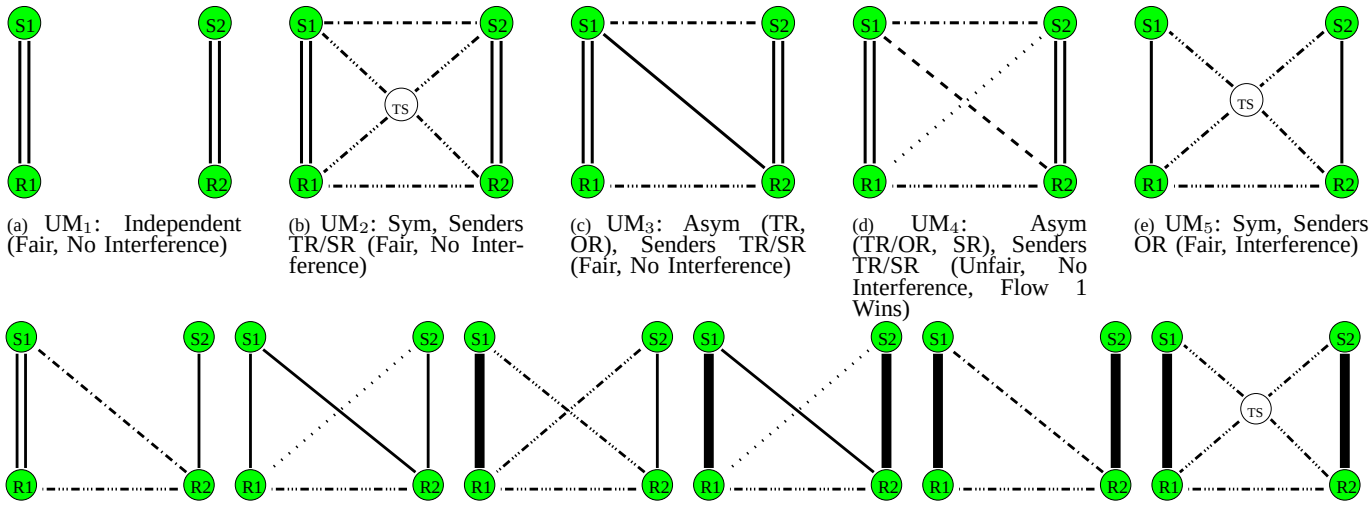

(a) UM_1: Independent (Fair, No Interference)

(b) UM_2: Sym, Senders TR/SR (Fair, No Interference)

(c) UM_3: Asym (TR, OR), Senders TR/SR (Fair, No Interference)

(d) UM_4: Asym (TR/OR, SR), Senders TR/SR (Unfair, No Interference, Flow 1 Wins)

(e) UM_5: Sym, Senders OR (Fair, Interference)

(f) UM_6: Asym (TR/SR, OR), Senders OR (Unfair, Interference, Flow 1 Wins)

(g) UM_7: Asym (TR, SR), Senders OR (Unfair, Interference, Flow 1 Wins)

(h) UM_8: 1 Robust Flow, Senders OR (Unfair, Interference, Flow 1 Wins)

(i) UM_9: 2 Robust Flows, Asym (TR, SR), Senders OR (Unfair, No Interference)

(j) UM_{10}: 2 Robust Flows, Asym (TR/SR, OR), Senders OR (Unfair, Flow 1 loses, No Interference)

(k) UM_{11}: 2 Robust Flows, Sym, Senders OR (Fair, No Interference)

Fig. 2. Models for UDP flows

has to hold back its transmission for a period of EIFS after the other sender finishes sending.

d) UM_4: Asymmetric (TR/OR, SR), Senders TR/SR: The next model covers the remaining asymmetric, senders TR/SR cases. In this model bandwidth allocation unfairly disfavours the flow whose sender is only able to sense the other receiver (i.e. flow 2 in Figure 2(d)). The reason this flow gets less bandwidth is that its sender must wait EIFS after it senses the MAC-ACK packet of the other flow.

2) Senders OR: We now consider three models in which senders are out of range of each other. In these models, since the senders cannot sense each other's packets at all, packet collisions are expected to happen more often due to the hidden terminal problem.

a) UM_5: Symmetric, Senders OR: When the topology is symmetric, similar to UM_2, the two flows receive a fair bandwidth allocation. But unlike the earlier model, the fact that senders can not sense means that DATA packets can be corrupted by the other flow.

b) UM_6: Asymmetric (TR/SR, OR), Senders OR: This model covers the asymmetric, senders OR cases where only one of the two senders is not able to sense the other receiver. This model leads to unfair bandwidth allocation with the flow whose sender senses the other flow dominating. The reason is that both senders send DATA packets at the same time, the packets sent to the receiver of the losing flow will be garbled, but the other receiver sees only the packet from its own flow.

c) UM_7: Asymmetric (TR, OR), Senders OR: This model is similar to UM_4 but with two changes. First, two senders are now out of sensing range. Second, one sender has to be in transmission range of the other flow. In other words, the topology asymmetry (SR, OR) belongs to UM_6 instead. In this model, due to the EIFS effect, the flow with sender only sensing the other receiver loses.

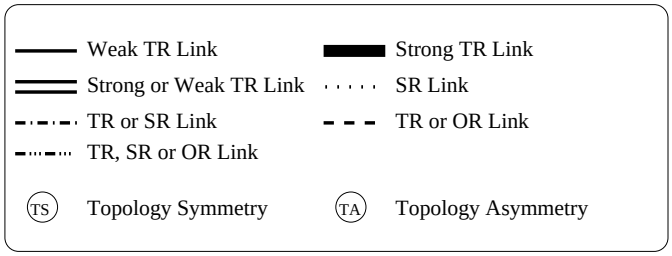

Fig. 3. Legend For All Models

3) Robust Flows: We now consider the scenarios in which one or both flows are robust to noise. The following three models represent the cases where the flows behave differently from the previous models.

a) UM_8: One Robust Flow, Senders OR: If only one flow is robust, it dominates the other flow whenever senders are out of sensing range. Any link between two flows indicates that there is the chance that the robust flow would corrupt the packets of the weak flow, and therefore the robust flow always wins.

b) UM_9: Two Robust Flows, Asymmetric (TR, SR), Senders OR: Just as model UM_7, the flow whose sender senses loses since it unnecessarily uses EIFS to back off whenever it senses the MAC ACKs from the other flow. But there is no interference in this model.

c) UM_{10}: Two Robust Flows, Asymmetric (TR/SR, OR), Senders OR: In this model, the flow whose sender can capture or sense the MAC-ACK packets of the other flow will lose. This is because the sender will delay its transmission unnecessarily while the other sender sends at full speed.

d) UM_{11}: Two Robust Flows, Symmetric, Senders OR: Symmetric topology leads to fair network allocation in this model. Moreover, since both flows are robust to interference, there is no packet collision either.

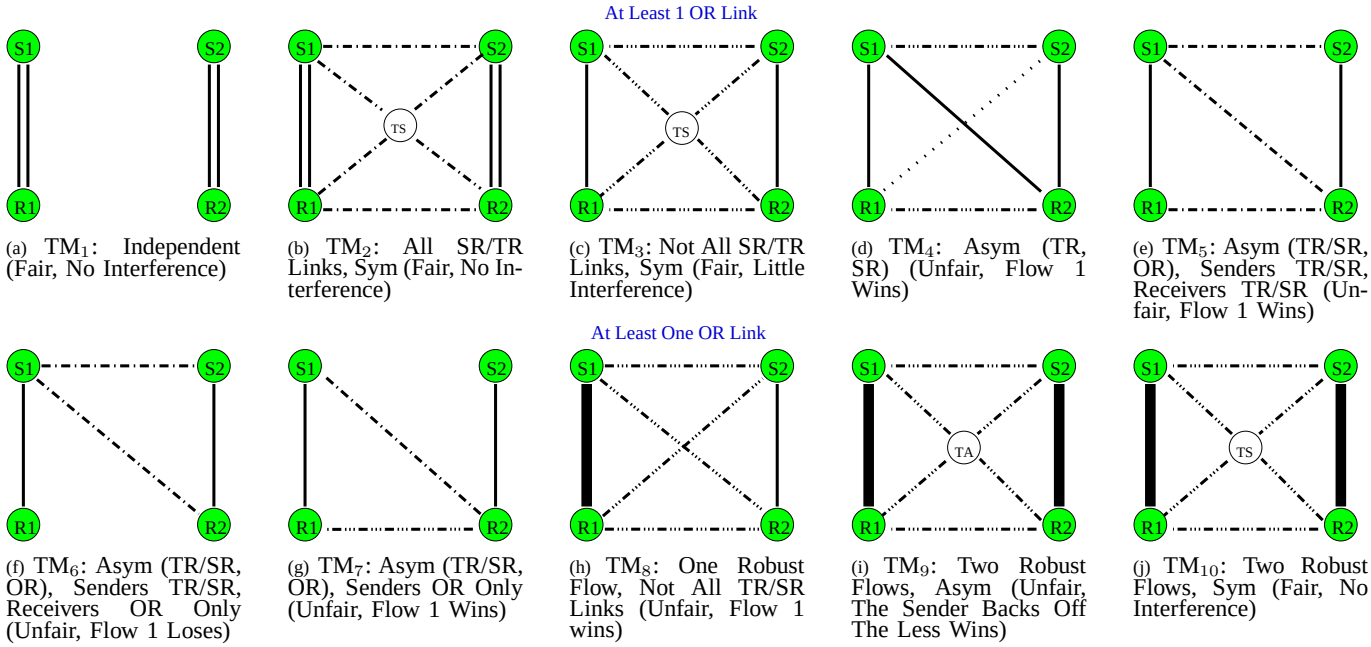

(a) TM_1: Independent (Fair, No Interference)

(b) TM_2: All SR/TR Links, Sym (Fair, No Interference)

(c) TM_3: Not All SR/TR Links, Sym (Fair, Little Interference)

(d) TM_4: Asym (TR, SR) (Unfair, Flow 1 Wins)

(e) TM_5: Asym (TR/SR, OR), Senders TR/SR, Receivers TR/SR (Unfair, Flow 1 Wins)

(f) TM_6: Asym (TR/SR, OR), Senders TR/SR, Receivers OR Only (Unfair, Flow 1 Loses)

(g) TM_7: Asym (TR/SR, OR), Senders OR Only (Unfair, Flow 1 Wins)

(h) TM_8: One Robust Flow, Not All TR/SR Links (Unfair, Flow 1 wins)

(i) TM_9: Two Robust Flows, Asym (Unfair, The Sender Backs Off The Less Wins)

(j) TM_{10}: Two Robust Flows, Sym (Fair, No Interference)

Fig. 4. Models for TCP flows

B. TCP Models

Modeling TCP is more complex than UDP due to the fact that each TCP flow consists of two sub-flows: DATA packets sent from sender to receiver and TCP-ACK packets sent from receiver to sender. TCP-ACK packets differ from MAC-ACK packets in the way that they are initiated. A MAC-ACK packet follows reception of a DATA packet after a bounded interval, but TCP-ACK packets are simply DATA packets to the MAC layer and are thus sent only when the channel is sensed to be clear. However, the TCP sender plays a more important role than the receiver since it triggers the TCP-ACK subflow and generates twice as much packets.

We will first introduce the three fair models. We then present four unfair models based on their different causes, followed by three models that consider robust flows.

1) Fair Models:

a) TM_1: Independent: We again begin with the trivial model in which two flows that are sufficiently distant from each other that neither flow affects the other.

b) TM_2: All SR/TR Links, Symmetric: Model TM_2 covers the symmetric topologies in which all pairs of nodes are at least within sensing range. Since the topology is symmetric, the two flows will split the network bandwidth evenly. Also, considering any two TCP subflows in this model, the two senders can at least sense each other and thus will back off sufficiently to avoid packet collisions.

c) TM_3: Incomplete and Symmetric: When not all pairs of nodes are connected, two TCP flows can still achieve fairness as long as the topology is symmetric. However, if two nodes cannot either decode or sense each other's packets, there is a good chance that their packets will collide, which distinguishes this model from TM_2.

It is worth pointing out that TCP treats transmission failure as signal of network congestion and consequently reduces its sending rate. This behaviour is often not desirable, because wireless links are usually lossy, but it indeed alleviates the problems of congestion and signal interference. For example, given a symmetric topology in which only the link between the sender of flow 1 and sender of flow 2 is not present, interference causes the collective 802.11b UDP throughput to drop by 30% , while TCP flows suffer only 7% degradation.

2) Unfair Models: There are four asymmetric topologies that result in unfair network allocation. One is attributed to the fact of EIFS versus DIFS; the others are due to packet collisions.

a) TM_4: Asymmetric (TR, SR): The only cause of asymmetry in this model is that the sender of flow 1 and the receiver of flow 2 are within transmission range, but the sender of flow 2 and the receiver of flow 1 are within sensing range. This asymmetry leads flow 2 to lose because the sender 2 has to use EIFS to schedule its next transmission after the receiver 1 sends a packet (either a MAC- or TCP-layer ACK) while the sender 1 uses DIFS after receiver 2 finishes transmitting.

b) TM_5: Asymmetric (TR/SR, OR), Senders TR/SR, Receivers TR/SR: When the only link missing is the link between the TCP sender of flow 2 and TCP receiver of flow 1, most of the packet collisions are TCP-DATA/TCP-ACK collisions. Flow 1 wins in this model because sending a TCP-ACK packet usually takes less time than sending a TCP-DATA packet. Therefore, when such collisions occur, the probability of successfully retransmitting an ACK packet is higher than that of a DATA packet. Moreover, sender 1 can take the chance to send more data packets during the time that sender 2 is backing off.

c) TM_6: Asymmetric (TR/SR, OR), Senders TR/SR, Receivers OR Only: In this model, the receiver of flow 1 is likely to send at the same time as the two nodes of flow 2. When this happens, its packets will be corrupted but the flow

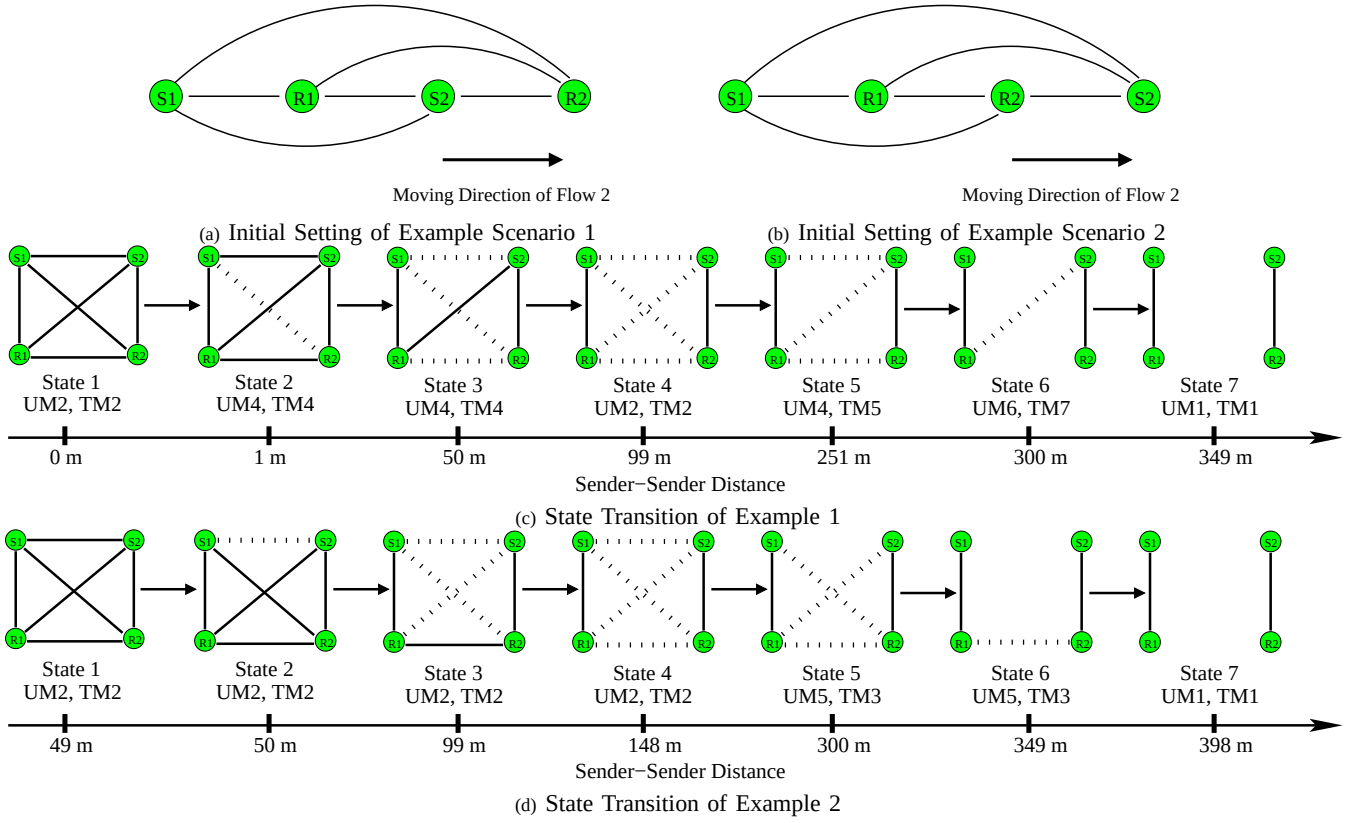

(a) Initial Setting of Example Scenario 1 (b) Initial Setting of Example Scenario 2

(c) State Transition of Example 1

(d) State Transition of Example 2

Fig. 5. Model Transition in Two Scenarios

2's packets are intact and thus flow 2 dominates.

d) TM_7: Asymmetric (TR/SR, OR), Senders OR Only: When the senders are out of range, the main problem is collision between two TCP-DATA subflows. The flow whose data packets are not corrupted dominates. For example, if sender 2 and receiver 1 are out of range, then flow 1 wins.

3) Robust Flows: We now consider the models in which one or both flows are robust to any interference.

a) TM_8: One Robust Flow, Not All TR/SR Links: If flow 1 is the only robust flow, but not all nodes are within transmission or sensing range of each other, then packet collisions will occur. Flow 1 is not affected because of its strong signals while flow 2 has to back off and retransmit. Thus, flow 1 always wins.

b) TM_9: Two Robust Flows, Asymmetric: When both flows are immune to interference, it is unfair only when the topology is not symmetric. In this case, any backoff is a waste and thus the TCP flow whose sender holds back more loses. For example, if the TCP sender of flow 1 is out of range of the receiver of flow 2, then flow 1 wins. Also, if the TCP sender of flow 1 is within the transmission range of the receiver of flow 2 while the TCP sender of flow 2 is within sensing range of the receiver of flow 1, then flow 1 wins as well.

c) TM_{10}: Two Robust Flows, Symmetric: When both flows are immune to interference and the topology is symmetric, the network bandwidth will be evenly distributed to the two flows without packet collisions.

V. Evaluation

We simulate all 648 scenarios using Glomosim [22]. For ease of exposition, we do not provide detailed results of the simulation. Instead, we present the results from the following two example scenarios that are capable of capturing behaviours of all the weak-link models but UM_3, UM_7 and TM_6. These two scenarios also demonstrate the importance of distinguishing between sense-only and transmit range and between UDP and TCP traffic as we do. Previous work that did not make these distinctions would have missed significantly in predicting node performance.

A. Example Scenarios

The four nodes are initially placed as illustrated in Figure 5(a) and 5(b), ensuring that all pairs of nodes are within transmission range. The only difference between these two examples is the placement of nodes S_2 and R_2.

The nodes of flow 2 then gradually move away from flow 1, which weakens the signal strength between these two flows. This causes three links, i.e., (S_1, S_2), (S_1, R_2) and (R_1, R_2), to experience all three link states as the distance increases. We can see from Figure 5(c) and 5(d) that each example involves 7 state transitions until they completely move out each other's radio range.

B. Simulation

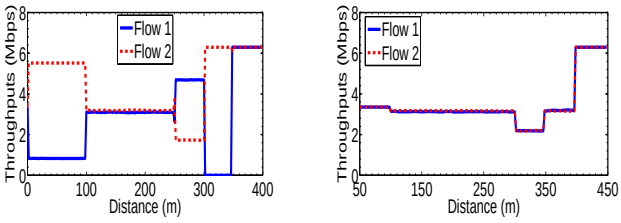

(a) 802.11b, Example 1 (b) 802.11b, Example 2

Fig. 6. UDP Performance in Simulations

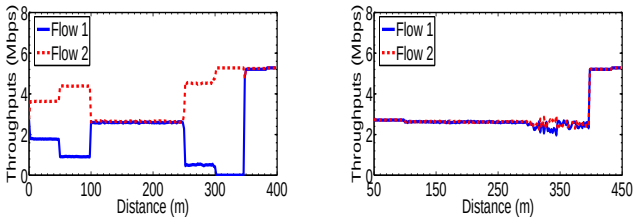

(a) 802.11b, Example 1 (b) 802.11b, Example 2

Fig. 7. TCP Performance in Simulations

1) The Simulator: Glomosim [22] is a scalable wireless simulator. However, it does not accurately model sense-only packets; a wireless node uses EIFS when packet collisions occur but not after sensing [3]. We have fixed this problem.

The transmission rates are set to the highest rates 11Mbps in our simulations. The minimum receiving signal strength is set to -58 dBm and minimum sensing signal strength -76 dBm, corresponding to 50 meters and 300 meters in the two-way propagation model. The minimum signal-to-noise ratio is set to 18 dBm (receiving signal strength - sensing signal strength) so that nodes out of sensing range are also out of interference range. It is worth noting that, since our model only assumes three link states, our choices of distances of transmission range and sensing range are rather arbitrary. Each simulation lasts 900 seconds and repeats 10 times with different seeds.

2) Simulation Results: The simulation shows that UDP and TCP performance differs significantly as seen by comparing Figures 6 and 7. This results confirm the importance of modeling TCP separately from UDP.

Similarly, an example of the importance of distinguishing sensed traffic that can be read from that cannot is seen in scenario 1 by comparing state 1 with 2 and state 5 with 6. This difference has led the performance of these states to vary noticeably. The previous models that assume that all sensed traffic is readable would have mispredicted the performance of states 2 and 5.

3) UDP flows: We present the mean values of two UDP flows' throughputs in Figure 6; the standard deviations are low and thus not shown.

a) Scenario 1: We can see from Figure 6(a) that, in the first scenario, flow 1 loses to flow 2 when the distance between two senders is between 1 and 99 meters. As we have explained in model UM_4, this is due to the fact that the asymmetric states 2 and 3 force the sender S_1 to use EIFS to schedule its next packet when it senses signals from the receiver R_2. Due to the same EIFS impact, the network bandwidth is unfairly distributed to these two flows when the distance is between 251 and 299 meters, except that this time flow 1 wins. Beyond 300 meters, sender S_1 and receiver R_2 move out of range. This asymmetric topology belongs to model UM_6 in which flow 2 will completely dominate flow 1 since it can garble flow 1's packets but not vice versa.

b) Scenario 2: We can see from Figure 6(b) that two flows fairly share the network bandwidth throughout all the states in scenario 2 even if interference occurs. This is expected because all the topologies shown in Figure 5(d) are symmetric and thus fall into models that promise fairness.

4) TCP flows: The TCP performance results are presented in Figure 7. Flow 2 wins in scenario 1 when interference starts to occur while both flows achieve fairness throughout scenario 2.

a) Scenario 1: When the distance between the senders is between 1 and 99 meters, all topologies belong to model TM_4, in which the two flows share the network bandwidth unfairly and the flow whose sender senses more loses. In other words, flow 2 wins. Beyond 250 meters, sender S_1 and receiver R_2 move out of sensing range. The topology belongs to model TM_5 and thus, unlike the corresponding UDP scenario, flow 1 loses this time because the TCP data packets of node S_1 collide with the TCP ACK packets sent from node R_2. When the distance increases to beyond 300 meters, the two TCP senders move out of sensing range and flow 1 is shut out due to the TCP DATA-DATA collisions at node R_1.

b) Scenario 2: Two flows fairly share the network throughout all topologies in scenario 2. As in the UDP scenarios, after the two senders move beyond 300 meters, the impact of interference starts to show. However, the performance penalty, compared to that of UDP flows, is much less, because TCP's sending rate is regulated by its ACKs. Lowering the sending rate can significantly alleviate congestion and thus reduce packet collisions. On the flip side, the performance of TCP flows fluctuate due to TCP's congestion avoidance scheme; the standard deviation can be as high as 800Kbps and starvations lasting for tens of seconds are not uncommon.

C. Sensing Range

Finally, we conducted a different set of experiments to understand how frequently nodes sense traffic that is too weak for them to read. The difficulty in collecting this information experimentally is that 802.11 CSMA is implemented in firmware and we are thus not able to directly determine in software when a wireless adaptor is sensing traffic.

We thus used an indirect approach to measure sense-only traffic. We configured a machine with two Dlink DWL-G520 wireless adaptors, which are based on the Atheros chipset.

One card is used to sense traffic by sending 1-byte messages at roughly 2-second intervals. We carefully measure the latency of each packet send to determine whether the wireless adaptor sensed traffic and thus backed off before sending the probe packet.

[3]NS2 has a similar inaccurate MAC model as well.

The other card operates in *monitor mode* to passively capture all traffic readable by the card, regardless of its destination address. We carefully log the start and end time of every packet received by the card and correlate these times with the log generated by the first card. If we see that a probe packet was delayed by backoff at a time when the second card was not receiving a packet, we conclude that this backoff is due to traffic that can be sensed but not read (i.e., in sense-only range).

We collected two 48-hour traces in two university labs in two buildings. We conservatively set the backoff-delay threshold to 350 μs — the longest possible 802.11g first-try back-off time without any optimization — 50 μs + 15 * 20 μs; the average delay in our traces is around 120 μs. We consider delays longer than this threshold to indicate a packet sense or capture. The trace files show an average of 75% of the delays are due to sensing instead of packet receiving. Even in the worst hour in our trace, at least 42% of the backoffs are due to signal sensing. Lowering the delay threshold actually increases the percentage of sensing delays. Among the entire data collected delays as large as 9.6 ms were observed, during which 18 packets were received.

VI. Conclusion

This paper analyzes the scenarios of two competing flows and provides 21 concrete 802.11b models that predict the performance and fairness based on node topology. These models consider three factor absent from previous work: (1) sensing state, (2) TCP flows and (3) weak or strong signals. Our testbed validates these models and the results show that they are indeed accurate.

References

[1] V. Bharghavan, A. Demers, S. Shenker, and L. Zhang, "Macaw: A media access protocol for wireless lans," in *SIGCOMM '04*, Aug. 1994, pp. 212–225.

[2] P. Karn, "Maca: A new channel access method for packet radio," in *ARRL/CRRL Amateur Radio 9th Computer Networking Conference*, Apr. 1990.

[3] V. Bharghavan, "Performance evaluation of algorithms for wireless medium access," in *IPDS '98: IEEE Internation Computer Performance and Dependability Symposium*, Mar. 1998.

[4] C. cheng Chen and H. Luo, "The case for heterogeneous wireless macs," in *HotNets '05: Proceedings of the 4th Workshop on Hot Topics in Networks*, Nov. 2005.

[5] T. Henderson, D. Kotz, and I. Abyzov, "The changing usage of a mature campus-wide wireless network," in *Proceedings of the Tenth Annual International Conference on Mobile Computing and Networking (MobiCom)*, September 2004, pp. 187–201.

[6] M. Garetto, J. Shi, and E. W. Knightly, "Modeling media access in embedded two-flow topologies of multi-hop wireless networks," in *MobiCom '05: Proceedings of the 11th annual international conference on Mobile computing and networking*, 2005, pp. 200–214.

[7] J. Deng, B. Liang, and P. K. Varshney, "Tuning the carrier sensing range of ieee 802.11 mac," in *GLOBECOM '04: Global Telecommunications Conference*, Dec. 2004.

[8] K. Xu, M. Gerla, and S. Bae, "How effective is the ieee 802.11 rts/cts handshake in ad hoc networks," in *GLOBECOM '02: Global Telecommunications Conference*, Nov. 2002, pp. 72–76.

[9] K. Jamieson, B. Hull, A. Miu, and H. Balakrishnan, "Understanding the real-world performance of carrier sense," in *E-WIND '05: Proceeding of the 2005 ACM SIGCOMM workshop on Experimental approaches to wireless network design and analysis*, 2005, pp. 52–57.

[10] H. Zhai and Y. Fang, "Physical carrier sensing and spatial reuse in multirate and multihop wireless ad hoc networks," in *Infocom '06*, Apr. 2006.

[11] X. Yang and N. H. Vaidya, "On the physical carrier sense in wireless ad-hoc networks," in *Infocom '05*, 2005.

[12] G. Tan and J. Guttag, "Long-term time-share guarantees are necessary for wireless lans," in *Proc. of SIGOPS European Workshop*, Leuven, Belgium, Sept. 2004.

[13] N. H. Vaidya, P. Bahl, and S. Gupta, "Distributed fair scheduling in a wireless lan," in *MobiCom '00: Proceedings of the 6th annual international conference on Mobile computing and networking*, 2000.

[14] H.-J. Ju, I. Rubin, and Y.-C. Kuan, "An adaptive rts/cts control mechanism for ieee 802.11 mac protocol," in *VTC '03*, Apr. 2003.

[15] F.Talucci, M.Gerla, and L.Fratta, "Macabi (maca by invitation): A receiver oriented access protocol for wireless multiple networks," in *PIMRC '97*, 1994, pp. 1–4.

[16] H. Kim and J. C. Hou, "Improving protocol capacity with model-based frame scheduling in ieee 802.11-operated wlans," in *MobiCom '03: Proceedings of the 9th annual international conference on Mobile computing and networking*, 2003, pp. 190–204.

[17] B. Sadeghi, V. Kanodia, A. Sabharwal, and E. Knightly, "Opportunistic media sccess for multirate ad hoc networks," in *MobiCom '02: Proceedings of the 8th annual international conference on Mobile computing and networking*, 2002, pp. 24–35.

[18] G. Bianchi, "Performance analysis of the ieee 802.11 distributed coordination function," *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 3, pp. 535–547, Mar. 2000.

[19] H. Zhai, Y. Kwon, and Y. Fang, "Performance analysis of ieee 802.11 mac protocols in wireless lans," *Wireless Communications and Mobile Computing, Special Issue on Emerging WLAN Technologies and Applications*, vol. 4, pp. 917–931, Dec. 2004.

[20] N. Gupta and P. R. Kuman, "A performance analysis of the 802.11 wireless lan medium access control," *Communications in Information and Systems*, vol. 3, no. 4, pp. 279–304, Sept. 2004.

[21] H. Kim and J. C. Hou, "A fast simulation framework for ieee 802.11-operated wireless lans," in *SIGMETRICS '04: Proceedings of the joint international conference on Measurement and modeling of computer systems*, 2004, pp. 143–154.

[22] X. Zeng, R. Bagrodia, and M. Gerla, "Glomosim: A library for parallel simulation of large-scale wireless networks," in *PADS '98*, May 1998.

[23] K. Cai, M. Feeley, and S. George, "Understanding 802.11 performance for two competing flows," unpublished, in submission to MASS 2007.